Oracle® Exalogic Elastic Cloud

Machine Owner's Guide

Release 2.0.6.4

E18478-24

February 2018

ORACLE®

Oracle Exalogic Elastic Cloud Machine Owner's Guide, Release 2.0.6.4

E18478-24

Primary Author: Salvador Esparza, Ashish Thomas

Contents

6 Understand Network Requirements and Configuration

7 Set Up ILOM on the Compute Nodes

13 Manage the InfiniBand Network Using Subnet Manager

14 Use the Sun Datacenter InfiniBand Switch 36 in Multirack Configurations

15 Monitor and Control the InfiniBand Fabric

16 Use InfiniBand Partitions in Exalogic Physical Environments

17 Monitoring the Exalogic Machine Using Oracle Enterprise Manager Ops Center

18 Install the ASR Software

A Configure the Exalogic Machine Using ECU

B Site Checklists

C Cabling Diagrams

D Replacement Units

E Cabling Tables

F Manage Solaris Zones on Exalogic

G Customize Linux on the Compute Nodes

List of Figures

List of Tables

Preface

This guide describes the Oracle Exalogic machine, which is an integrated cloud machine comprising hardware and software. It includes information about hardware operations and site planning, as well as physical, electrical, and environmental specifications.

> **Note:**
>
> All hardware-related specifications in this guide are based on information for a typical deployment provided by Oracle at the time this guide was written. Oracle is not responsible for hardware problems that may result from following the average specifications in this document. For detailed information about preparing your site for Oracle Exalogic machine deployment, consult your hardware specification.

This preface contains the following sections:

- Audience
- Revision History
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This guide is intended for Oracle Exalogic machine customers and those responsible for data center site planning.

It is assumed that the readers of this manual have knowledge of the following:

- System administration concepts
- Hardware and networking concepts

Revision History

- **E18478-18: April 2014**
- **E18478-17: February 2014**
  – Hardware Components of Exalogic: Added generic link that points to data sheets.

  - Environmental Requirements: Added generic link that points to data sheets.
  - What Next?: Added note with a link to MOS document for changing passwords.
  - Cabling Diagrams : Removed sections about multirack cabling and added a reference to the *Exalogic Elastic Cloud Multirack Cabling Guide.*
- **E18478-16: December 2013**
  - Hardware Components of Exalogic: Added generic link that points to X4-2 data sheet.
  - Hardware Components of Exalogic: Added precabling information for X4-2.
  - Environmental Requirements: Added generic link that points to X4-2 data sheet.
  - Cabling Tables: Added precabling information for X4-2.
  - Rebranded the document to cater to X4-2 in addition to X3-2 and X2-2.
- **E18478-15: September 2013**
  - Added Guidelines for Managing pkey Allocation in a Hybrid Rack.
- **E18478-14: August 2013**
  - Default Port Assignments: Added default ports assignments of the Exalogic rack.
  - Initial Configuration: Updated note about the initial configuration of the storage appliance.
  - Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1: Added Oracle Solaris 11.1 procedure for setting up Ethernet over InfiniBand.
  - Oracle Solaris: Creating VNICs and Associating Them with VLANs: Added Oracle Solaris 11.1 procedure for creating VLAN-tagged VNICs.
  - Activate the Storage Appliance: Updated the procedure and screenshot in this section.
  - Manage Solaris Zones on Exalogic: Created appendix describing the use of Solaris zones on Exalogic.
  - Customize Linux on the Compute Nodes: Created appendix describing how to install, upgrade, and remove RPMs on Linux compute nodes.
- **E18478-13: March 2013**
  - Hardware Components of Exalogic: Added direct links to the data sheets for X2-2 and X3-2.
  - Hardware Components of Exalogic: Added information about the spares kit for Exalogic.
  - Cabling Diagrams : Added additional Exalogic wiring configurations.
  - Replacement Units: Added note about the spares kit for Exalogic.
  - Cabling Tables: Corrected errors in the networking tables.
- **E18478-12: February 2013**
  - Parts for Sun Server X3-2 Compute Nodes: Added the replacement part numbers of the new BBU.

  - Create an IPoIB Partition and Adding Ports: Corrected the `smpartition` command syntax in step 3.
- **E18478-11: December 2012**
  - Hardware Components of Exalogic: Added generic link that points to both X2-2 and X3-2 data sheets.
  - Environmental Requirements: Added generic link that points to both X2-2 and X3-2 data sheets.
- **E18478-10: October 2012**
  - Environmental Requirements: Added environmental requirements for X3-2.
  - Create an IPoIB Partition and Adding Ports: Added steps for configuring the storage appliance when creating IB partitions.
  - Chapter 16, "Monitoring the Exalogic Machine Using Oracle Enterprise Manager Ops Center": Added information, up front, to clarify the scope of the content in this chapter.
  - Replacement Units: Added replacement parts list for the new hardware components in X3-2.
  - Rebranded the document cater to both X3-2 and X2-2.
- **E18478-09: September 2012**
  - Configure Ethernet Over InfiniBand:

    Updated the procedure to set up EoIB on Oracle Linux and Oracle Solaris, to correct certain errors and inconsistencies.

    Added information about the recommended naming convention for VNICs on Oracle Linux.
  - Subnet Manager Operation in Different Rack Configurations:

    Updated Table 13-1 to make the information clearer.

    Added information for running the SM in rack configurations with varying switch-firmware versions.
  - Create a Partition for EoIB and Associating the pkey with a VNIC and VLAN: Added information about the recommended naming convention for VNICs on Oracle Linux.
  - Configuration Tasks: Updated with information about contacting Oracle Advanced Customer Support.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information, see the following documents:

- *Oracle Exalogic Release Notes*
- *Oracle Exalogic Machine Multirack Cabling Guide*

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.
italic	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.

1 Learn About the Exalogic Machine

This chapter describes the features and hardware components of an Exalogic machine. It also includes usage information related to the Exalogic machine.
This chapter contains these topics:

- About Oracle Exalogic
- Hardware Components of Exalogic
- Exalogic Machine Rack Layout
- Operational Procedures for Exalogic Machines

1.1 About Oracle Exalogic

Oracle Exalogic is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and widely varied workloads. Exalogic is intended for large-scale, performance-sensitive, mission-critical application deployments. It combines Oracle Fusion Middleware software and industry-standard Sun hardware to enable a high degree of isolation between concurrently deployed applications, which have varied security, reliability, and performance requirements. Exalogic enables customers to develop a single environment that can support end-to-end consolidation of their entire applications portfolio.

Exalogic is designed to fully leverage an internal InfiniBand fabric that connects all of the processing, storage, memory and external network interfaces within an Exalogic machine to form a single, large computing device. Each Exalogic machine is connected to the customer's data center networks via 10 GbE (traffic) and GbE (management) interfaces.

Customers can integrate Exalogic machines with an Exadata machine or additional Exalogic machines by using the available InfiniBand expansion ports and optional data center switches. The InfiniBand technology used by Exalogic offers significantly high bandwidth, low latency, hardware-level reliability, and security. If you are using applications that follow Oracle's best practices for highly scalable, fault-tolerant systems, you do not need to make any application architecture or design changes to benefit from Exalogic. You can connect many Exalogic machines or a combination of Exalogic machines and Oracle Exadata Database Machines to develop a single, large-scale environment.

1.1.1 Features of Exalogic

Exalogic includes the following features:

- complete, pre-assembled, and certified system for maximum performance
- best platform for end-to-end consolidation of large-scale application deployments
- ready to deploy
- linearly scalable

- enterprise-ready
- enterprise-level software and hardware support

1.2 Hardware Components of Exalogic

The Exalogic machines consist of compute nodes, a storage appliance, and InfiniBand and Ethernet networking components. The number of these components in each machine varies based on the hardware configuration.

> **Note:**
>
> The hardware components listed in this chapter are for Exalogic machines shipped from the factory.

Table 1-1 Hardware Components of Exalogic

Component	Full Rack	Half Rack	Quarter Rack	Eighth Rack
Sun Rack II 1242	1	1	1	1
Compute nodes: • **Oracle Server X6–2** • Previous versions of the hardware: Sun Fire X4170 M2 (X2–2), Sun Server X3-2, Sun Server X4-2, and Oracle Server X5-2.	30	16	8	4
Storage Appliance: • **X6-2**: Oracle ZS5–ES storage appliance • **X2-2** and **X3-2**: Dual controller Sun ZFS Storage 7320 appliance (60 TB) • **X4-2** and **X5-2**: Oracle ZS3-ES storage appliance	1	1	1	1
Sun Network QDR InfiniBand Gateway Switches (NM2-GW)	4	2	2	2
Sun Datacenter InfiniBand Switch 36 (NM2-36P). This spine switch is not included in Exalogic machine configurations and must be purchased separately.	0	0	0	0
Cisco Ethernet management switch • **X2-2**: 48-port Cisco Catalyst 4948 switch (part number: 371-4784-N) • **X6-2**, **X5-2**, **X4-2**, and **X2-2**: 48-port Cisco Catalyst 4948E-F-S switch (part number: 7023685)	1	1	1	1
Redundant Power Distribution Units (PDU). See PDUs in Exalogic for the available PDU options.	2	2	2	2

You can connect up to eight Exalogic machines, or a combination of Exalogic and Exadata on the same InfiniBand fabric, without the need for any external switches.

> **Note:**
>
> If more than eight racks of Exalogic or Exadata are required to be connected on the same InfiniBand fabric, Oracle offers a choice of several high-capacity datacenter switches, which allow the creation of Exalogic clouds consisting of hundreds of racks and tens of thousands of processors.

1.2.1 PDUs in Exalogic

Each Exalogic rack contains **two redundant PDUs** that are available in **both high and low voltage options**. The available PDU options can be found in the data sheets available at the following location:

`http://www.oracle.com/technetwork/server-storage/engineered-systems/exalogic/index.html`

> **Note:**
>
> When you upgrade to a full rack, you must upgrade to a three-phase 24 kVA PDU.

1.2.2 Spares Kit for Exalogic

The Exalogic machine includes a spares kit containing additional parts and accessories (cables, for example) that Oracle Services personnel will use to replace non-working and broken parts quickly. When Oracle Services personnel visit your site for repairing and replacing hardware, you should make the spares kit available to them. Note that any parts used by Oracle Service personnel from the spares kit will be replenished by Oracle.

1.2.3 Pre-cabling for Exalogic X4-2 and newer Machines

The eighth-, quarter-, and half-rack X4-2 and newer machines are pre-cabled with a complete set (equivalent in number to a full rack) of InfiniBand (92), Ethernet (42), and power cables. All the cables are connected to the appropriate switches and routed to the correct rack-unit location. The unconnected ends of the cables are tied off to lacer bars. At a later time, if the machine is upgraded to a larger rack, the filler panels and lacer bars are removed, the X4-2 and newer servers are installed with rack rails and cable-management arms, and the cables (power, InfiniBand, and Ethernet) are connected to the newly installed compute nodes.

1.3 Exalogic Machine Rack Layout

The Exalogic machines are available in the following hardware configurations:

- full rack
- half rack

- quarter rack
- eighth rack

The following table displays the location of the hardware components on the 42 rack units of an Exalogic machine. The table includes the information for all the available rack configurations. The dashes within the rack configuration columns of the table, represent empty rack units that are covered with solid one-unit fillers. The rack unit 42 is at the top of the Exalogic machine and the rack unit 1 is at the bottom.

See a summary of the hardware components for the different Exalogic machine configurations in Hardware Components of Exalogic.

> **Note:**
>
> Do not alter the rack layout. Do not use the free space in the rack for anything other than a rack upgrade kit.

Table 1-2 Layout for all Exalogic machine configurations

Rack Unit	Full Rack	Half Rack	Quarter Rack	Eighth Rack	Front View	Rear View
U29 to U42	X	-	-	-	Compute Node	Compute Node
U27, U28	X	X	-	-	Compute Node	Compute Node
U26	X	-	-	-	Vented one-unit filler	NM2-GW InfiniBand Switch / Gateway
U25	X	X	X	X	Vented one-unit filler	Cisco Ethernet Management Switch
U24	X	-	-	-	Vented one-unit filler	NM2-GW InfiniBand Switch / Gateway
U23	X	X	X	X	Vented one-unit filler	NM2-GW InfiniBand Switch / Gateway
U21, U22	X	X	X	X	Storage Appliance Server Head	Storage Appliance Server Head
U17 to U20	X	X	X	X	Storage Appliance Disk Shelf	Storage Appliance Disk Shelf
U16	X	X	X	X	Vented one-unit filler	NM2-GW InfiniBand Switch / Gateway
U10 to U15	X	X	-	-	Compute Node	Compute Node
U06 to U09	X	X	X	-	Compute Node	Compute Node
U02 to U05	X	X	X	X	Compute Node	Compute Node
U01	-	-	-	-	Vented one-unit filler	NM2-36P InfiniBand Switch (not included)

1.4 Operational Procedures for Exalogic Machines

This section includes the following topics:

- Non-emergency Power Procedure

- Emergency Power-off Considerations
- Cautions and Warnings

1.4.1 Non-emergency Power Procedure

Compute nodes and controllers in the storage device are powered on by either pressing the power button on the front of the machine, or by logging in to the ILOM interface and applying power to the system.

Compute nodes and controllers can also be powered off by using operating system commands.

For example, you can use the following command on the Oracle Linux operating system:

```
shutdown -h -y now
```

On the Oracle Solaris operating system, you can use the following command:

```
shutdown -i 5 now
```

> **Note:**
>
> You can use this command after logging in to the compute nodes.
>
> The network switches do not have power switches. They power off when power is removed, by way of the PDU or at the breaker in the data center.

1.4.1.1 Power On Sequence

The power on sequence is as follows:

1. Rack, including switches.

   Ensure that the switches have had power applied for a few minutes to complete power-on configuration before starting the storage controllers and compute nodes.
2. Storage server heads attached to the chassis of the storage device.

   Wait for a few minutes for the storage appliance to boot and start NFS services and daemons.
3. Compute nodes.

> **Note:**
>
> After power is applied, the LEDs on all compute nodes and storage server heads will start blinking after a few minutes. From the rear of the rack, you can see the green LEDs on the PSUs turn on instantly after power is applied. In addition, from the rear of the rack, you can see the display on the PDUs that lights up once power is available.

1.4.1.2 Power Off Sequence

The power off sequence is as follows:

1. Compute nodes

   > **Note:**
   >
   > Once the compute nodes are down, you can proceed to shut down the storage appliance.

2. Storage server heads attached to the chassis of the storage device
3. Rack, including switches

1.4.2 Emergency Power-off Considerations

If there is an emergency, then power to the Exalogic machine should be halted immediately. The following emergencies may require powering off the Exalogic machine:

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

1.4.2.1 Emergency Power-off Procedure

To perform an emergency power-off procedure for the Exalogic machine, 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 machine.

1.4.2.2 Emergency Power-off Switch

Emergency power-off (EPO) switches are required when computer equipment contains batteries capable of supplying more than 750 volt-amperes for more than five minutes. Systems that have these batteries include internal EPO hardware for connection to a site EPO switch or relay. Use of the EPO switch will remove power from the Exalogic machine.

1.4.3 Cautions and Warnings

The following cautions and warnings apply to Exalogic machines:

- Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.
- Do not power off Exalogic machines unless there is an emergency. In that case, follow the Emergency Power-off Procedure.
- Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to hardware components.

- Keep the top, front, and back of the cabinets clear to allow proper airflow and prevent overheating of components.
- Use only the supplied hardware.

2 Site Requirements

This chapter describes the site requirements for the Exalogic machine.
This chapter contains the following topics:

- Environmental Requirements
- Space Requirements
- Flooring Requirements
- Electrical Power Requirements
- Temperature and Humidity Requirements
- Ventilation and Cooling Requirements

> **See Also:**
>
> Site Checklists

2.1 Environmental Requirements

The environmental requirements for an Exalogic machine depend on the purchased hardware configuration. The environmental requirements for an Exalogic rack can be found in the data sheets at the following location:

`http://www.oracle.com/technetwork/server-storage/engineered-systems/exalogic/index.html`

2.2 Space Requirements

All Exalogic machines use the same rack, and have the same space requirements. The space requirements are as follows:

- Height: 1998 mm (78.66 inches)
- Width: 600 mm with side panels (23.62 inches)
- Depth (front door handle to rear door handle): 1200 mm (47.24 inches)
- Depth (doors removed): 1112 mm (43.78 inches)

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

2.2.1 Receiving and Unpacking Requirements

Before your Exalogic machine arrives, ensure that the receiving area is large enough for the package. The following are the package dimensions for the Exalogic machine:

- Shipping height: 2159 mm (85 inches)
- Shipping width: 1219 mm (48 inches)
- Shipping depth: 1575 mm (62 inches)
- For the shipping weight, see the data sheets at the following location:

  `http://www.oracle.com/technetwork/server-storage/engineered-systems/exalogic/index.html`

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

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

Allow enough space for unpacking it from its shipping cartons. Ensure that there is enough clearance and clear pathways for moving the Exalogic machine from the unpacking location to the installation location. Table 2-1 lists the access route requirements for the Exalogic machine.

Table 2-1 Access Route Requirements

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 inches)	2040 mm (80.32 inches)
Minimum door width	1220 (48 inches)	600 mm (23.62 inches)
Minimum elevator depth	1575 mm (62 inches)	1200 mm (47.24 inches)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lbs)	1134 kg (2500 lbs)

2.2.2 Maintenance Access Requirements

The maintenance area must be large enough for the Exalogic machine, and have the required access space. For example, the required space to remove the side panels is 675. 64 mm (26.6 inches). Table 2-2 lists the maintenance access requirements for the Exalogic machine.

Table 2-2 Maintenance Access Requirements for Exalogic Machine

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

Table 2-2 (Cont.) Maintenance Access Requirements for Exalogic Machine

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

2.3 Flooring Requirements

Oracle recommends that the Exalogic machine be installed on raised flooring. The site floor and the raised flooring must be able to support the total weight of the Exalogic machine as specified in Environmental Requirements.

Table 2-3 lists the floor load requirements.

Table 2-3 Floor Load Requirements for Exalogic Machine

Description	Requirement
Maximum allowable weight of installed rack equipment	952.54 kg (2100 lbs)
Maximum allowable weight of installed power distribution units	52.16 kg (115 lbs)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1004.71 kg (2215 lbs)

> **Note:**
>
> Open tiles are required for electrical access.

2.4 Electrical Power Requirements

Exalogic Machine 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 Exalogic Machine:

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

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

Table 2-4 lists the PDU low voltage requirements.

Table 2-4 PDU Requirements for Low Voltage

Specification	15 kVA, 1 ph	15 kVA, 3 ph	22 kVA, 1 ph	24 kVA, 3 ph
Phase	1 ph	3 ph	1 ph	3 ph
Market Part Number	6442A	6440A	7100873	XSR-24K-IEC309-4P
Manufacturing Part Number	597-0566-01	597-0564-01	7018123	594-5596-01
Voltage	200-240 VAC 1ph	200-240 VAC 3ph	200-240 VAC 1ph	200-240 VAC 3ph
Amps Per PDU	72A (3 × 24A)	69A (3 × 23A)	110.4 (3x36.8A)	120A (6 × 20A)
Outlets	42 C13, 6 C19	42 C13, 6 C19	42 C13, 6 C19	42 C13, 6 C19
Number of Inputs	3x30A, 1 ph	1x60A, 3 ph	3 x50A 1 ph	2x60A, 3 ph
Current	24A max. per input	40A max. per phase	36.8A per input	34.6A max. per phase
Data Center Receptacle	NEMA L6-30R	IEC309-3P4W-IP67 (60A, 250V, AC, 3 ph) IEC309 60A 3ph 4 Wire Hubbell HBL460R/C9W or equivalent.	Hubbell CS8265C	IEC309-3P4W-IP67 (60A, 250V, AC, 3 ph) IEC309 60A 3ph 4 Wire Hubbell HBL460R/C9W or equivalent.
Number of Outlets Per Rack	6	2	6	4

Table 2-5 lists the PDU high voltage requirements.

Table 2-5 PDU Requirements for High Voltage

Specification	15 kVA, 3 ph	22 kVA, 1 ph	24 kVA, 3 ph
Phase	3 ph	1 ph	3 ph
Market Part Number	6441A	7100874	XSR-24K-IEC309-5P
Manufacturing Part Number	597-0565-01	7018124	594-5600-01
Voltage	220/380-240/415 VAC 3ph	200-240 VAC 1ph	220/380-240/415 VAC 3ph
Amps Per PDU	62.7 A (3 × 20.9A)	96A (3x32A)	109A (6 × 18.1A)
Outlets	42 C13, 6 C19	42 C13, 6 C19	42 C13, 6 C19
Number of Inputs	1x25A, 3 ph	3 x 32A 1 ph	2x25A, 3 ph
Current	24A max. per input	32 A per input	18 A max. per input
Data Center Receptacle	IEC 309-4P5W-IP44 (32A, 400V, AC, 3ph) IEC309 32A 3ph 5 Wire Hubbell HBL532R/C9W or equivalent.	IEC 309-2P3W-IP44 (32A, 250V, AC, 3ph) IEC309 32A 1ph 3 Wire Hubbell HBL332R/C9W or equivalent.	IEC 309-4P5W-IP44 (32A, 400V, AC, 3ph) IEC309 32A 3ph 5 Wire Hubbell HBL532R/C9W or equivalent.
Number of Outlets Per Rack	2	6	4

2.4.1 Facility Power Requirements

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

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 m (13.12 feet) long, and 1 to 1.5 m (3.3 to 4.9 feet) of the cord will be routed within the rack cabinet. The installation site AC power receptacle must be within 2 m (6.6 feet) of the rack.

2.4.2 Circuit Breaker Requirements

To prevent catastrophic failures, the design of your power system must ensure that adequate power is provided to all of the compute nodes. Use dedicated AC breaker panels for all power circuits that supply power to the compute nodes. Electrical work and installations must comply with applicable local, state, or national electrical codes. Compute nodes require electrical circuits to 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:**
>
> Circuit breakers are supplied by the customer. One circuit breaker is required for each power cord.

2.4.3 Grounding Guidelines

The cabinets for the Exalogic machine are shipped with grounding-type power cords (three-wire). 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 IEC documents, for the correct grounding method. Ensure that the facility administrator or qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

2.5 Temperature and Humidity Requirements

Airflow through Exalogic machines is from front to back. For cooling and airflow requirements, see Environmental Requirements.

> **Note:**
>
> Studies have shown that temperature increases of 10 degrees Celsius (15 degrees Fahrenheit) above 20 degrees Celsius (70 degrees Fahrenheit) reduce long-term electronics reliability by 50 percent.
>
> Excessive internal temperatures may result in full or partial shutdown of Exalogic machines.

Table 2-6 lists the temperature, humidity and altitude requirements for operating and nonoperating machines.

Table 2-6 Temperature, Humidity, and Altitude Requirements

Condition	Operating Requirement	Nonoperating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (59 to 89.6 degrees Fahrenheit)	-40 to 70 degrees Celsius (-40 to 158 degrees Fahrenheit).	For optimal rack cooling, data center temperatures from 21 to 23 degrees Celsius (70 to 47 degrees Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, non-condensing	Up to 93 percent relative humidity.	For optimal data center rack cooling, 45 to 50 percent, non-condensing
Altitude	3048 meters (10000 feet) maximum	12000 meters (40000 feet).	Ambient temperature is reduced by 1 degree Celsius per 300 m above 900 m altitude above sea level

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

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

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

- Optimal range helps protect computer systems from corrosion problems associated with high humidity levels.
- Optimal range provides the greatest operating time buffer in the event of air conditioner control failure.

- This range helps avoid failures or temporary malfunctions caused by intermittent interference from static discharges that may occur when relative humidity is too low.

> **Note:**
>
> Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent. It is not difficult to maintain humidity in a data center because of the high-efficiency vapor barrier and low rate of air changes normally present.

2.6 Ventilation and Cooling Requirements

Always provide adequate space in front of 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 warm air out the rear of the rack. There is no air flow requirement for the left and right sides due to front-to-back cooling.

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

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

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

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

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

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

- Four floor tiles for Exalogic machine full rack.
- Three floor tiles for Exalogic machine half rack.
- One floor tile for Exalogic machine quarter rack and for Exalogic machine eighth rack.

Figure 2-1 shows a typical installation of the floor tiles for an Exalogic machine full rack in a typical data center.

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

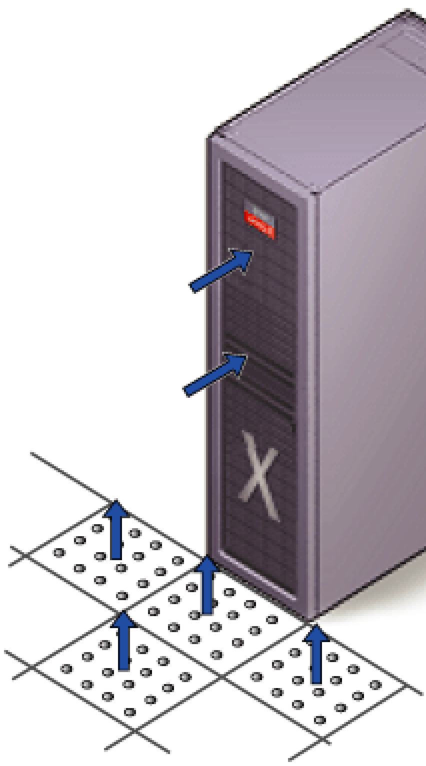

3 Install Planning and Procedure

This chapter explains how to prepare for the installation of Exalogic machines, to install an Exalogic machine at the site, and to power on the rack and its components. It contains the following topics:

3.1 Installation Overview

The following list describes the process of installing an Exalogic machine at the site. Note that you must commission the Exalogic machine after installing it at the site.

1. Review the safety precautions, guidelines, site checklists, and site requirements.
2. Ensure that the site is prepared for the installation of the Exalogic Machine.
3. Unpack the Exalogic machine.
4. Place the Exalogic machine in its allocated space.
5. Perform preliminary checks before connecting the power cords.
6. Perform a visual inspection of the hardware.
7. Supply rack power and perform the power-on self test device checks.
8. Switch on the six PDU circuit breakers located on the rear of the PDU A.
9. Wait three to five minutes for all ILOM service processors to boot.
10. Verify that the server standby power is on for each compute node.
11. Verify that the main power is on for each compute node.
12. Press the soft switches located on the front of the two storage heads of the storage appliance.
13. Wait three to five minutes for the storage appliance to start the NFS services, daemons, and basic services.
14. Ping the IP address assigned to the storage appliance to verify if the system is up and running.
15. Press the soft switches located on the front of the compute nodes.
16. Verify that power is applied to the Ethernet switch.
17. Verify that power is applied to the NM2-GW InfiniBand Gateway switches.
18. (Only for half and full racks.) Verify that power is supplied to the NM2-36 InfiniBand Switch.

19. Proceed to configure the Exalogic machine.

3.2 Prepare to Install an Exalogic Machine

This section contains the following topics:

- Before You Begin
- Exalogic Installation Safety Guidelines
- Unpack the Exalogic Machine Rack
- Tools for Installation
- Prepare the Installation Site

3.2.1 Before You Begin

Before installing the Exalogic machine, or installing any server or equipment into the rack, read the *Important Safety Information for Sun Hardware Systems* (816-7190) document included with the rack.

Observe all safety notices printed on the packaging and listed in the *Sun Rack II Safety and Compliance Guide* (820-4762) and the *Sun Rack II Power Distribution Units Users Guide* (820-4760). Go to `http://download.oracle.com/docs/cd/E19657-01/index.html` to download these guides.

Exalogic machine cabinets can enclose a variety of rack-mountable Sun servers, storage products, and other third-party equipment.

3.2.2 Exalogic Installation Safety Guidelines

Before the Exalogic machine arrives, the following safety precautions should be reviewed to ensure the site is safe, as well as ready for delivery. Failing to observe these precautions can result in injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not install the Exalogic machine in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install the Exalogic machine in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install the Exalogic machine in a location that is exposed to frequent vibrations. Install the Exalogic machine on a flat, level surface.
- Use a power outlet that uses proper grounding. When using shared grounding, the grounding resistance must not be greater than 10 ohms. Ensure that your facility administrator or a qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.
- Be sure that each grounding wire used for the Exalogic machine is used exclusively for the Exalogic machine. Also be sure to observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.

- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.
- Do not place anything on top of the Exalogic machine 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 the Exalogic machine. Allow for a sufficient warm-up period prior to server operation.
- Do not install the Exalogic machine near a photocopy machine, air conditioner, welding machine, or any other equipment that generates loud, electronic noises.
- Avoid static electricity at the installation location. Static electricity transferred to the Exalogic machine can cause malfunctions. Static electricity is often generated on carpets.
- Confirm that the supply voltage and frequency match the electrical ratings indicated on the Exalogic machine.
- Do not insert anything into any Exalogic machine opening, unless doing so is part of a documented procedure. The Exalogic machine contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in the Exalogic machine, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.

> **See Also:**
>
> - *Important Safety Information for Sun Hardware Systems* (816-7190) document that is included with the rack
> - All safety notices printed on the packaging and listed in *Sun Rack II Safety and Compliance Guide* (820-4762), and *Sun Rack II Power Distribution Units Users Guide* (820-4760)

3.2.3 Unpack the Exalogic Machine Rack

Refer to the unpacking instructions included with the packaging when unpacking the rack from the shipping carton. After unpacking the rack, follow local laws and guidelines to recycle the packaging properly.

> **Caution:**
>
> Carefully unpack the rack from the packaging and shipping pallet. Rocking or tilting the rack can cause it to fall over and cause serious injury or death. You should always use professional movers when unpacking and installing this rack.

> **Note:**
>
> After unpacking the rack from the packaging, save the mounting brackets used to secure the rack to the shipping pallet. You can use these mounting brackets to secure the rack permanently to the installation site floor. Do not dispose of these brackets, as you will not be able to order replacement brackets.

Figure 3-1 Unpacking Exalogic Machine Rack

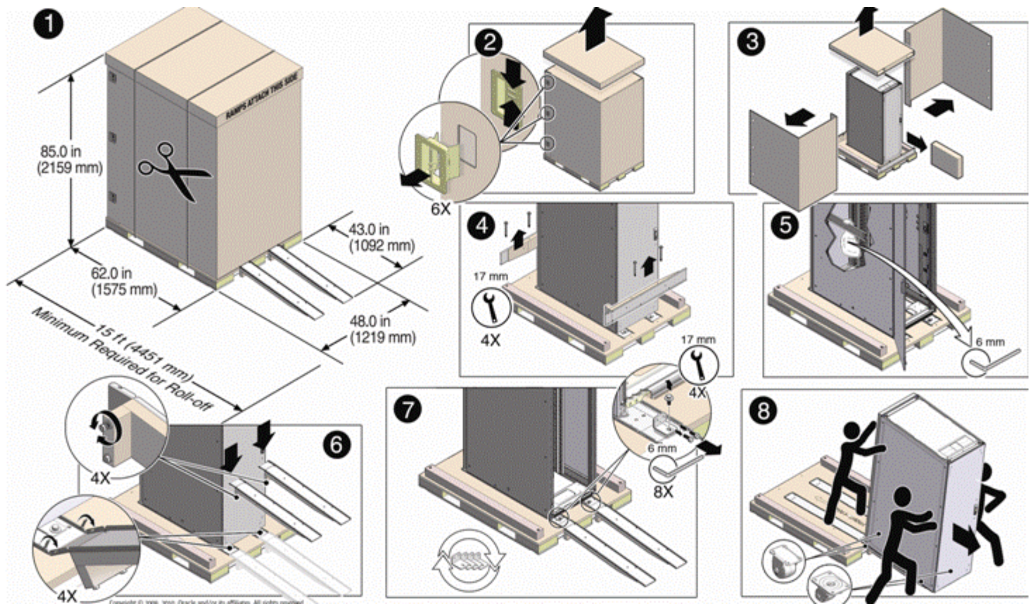

3.2.4 Tools for Installation

The following tools are required for installation:

- Screwdriver handle (magnetic)
- T-30 Torx wrench key
- T-25 Torx wrench key
- 6-mm hexagon Allen wrench key
- SW 12-mm single-headed wrench
- Side panel removal tool
- Keys to the front door, rear door, and side panel locks
- Cage nut mounting tool
- 32 M6 cage nuts
- 32 M6 screws
- Straight tip

3.2.5 Prepare the Installation Site

The following procedure describes how to prepare the site prior to unpacking and situating the Exalogic machine:

1. Thoroughly clean and vacuum the area in preparation for the installation.
2. Note problems or peculiarities at the site that require special equipment.
3. Verify that the installation site flooring has a strength rating to withstand the combined weight of the Exalogic machine and any other installed equipment.

   > **Note:**
   >
   > For more information, see Environmental Requirements.

4. Install all necessary electrical equipment and ensure that sufficient power is provided.

   > **See Also:**
   >
   > The *Sun Rack II Power Distribution Units User's Guide* for the Sun Rack II Power Distribution Unit (PDU) power requirements

5. Ensure that the installation site provides adequate air conditioning.
6. Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.

3.2.6 Place the Exalogic Machine in Its Allocated Space

This section contains the following topics:

- Move the Exalogic Machine
- Stabilize the Exalogic Machine
- Attach a Ground Cable (Optional)

3.2.6.1 Move the Exalogic Machine

The following procedure describes how to move an Exalogic machine:

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

   When moving the Exalogic machine to the installation site, the front casters do not roll; you must steer the unit by moving the rear casters. You can safely maneuver the Exalogic machine by carefully pushing it.

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

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

WARNING:

- Never attempt to move the Exalogic machine by pushing on the side panels. Pushing on the side panels can tip the rack over. This action can cause serious personal injury or death as well as damage to the equipment.
- Never tip or rock the Exalogic machine because the rack can fall over.

3.2.6.2 Stabilize the Exalogic Machine

After moving the Exalogic machine 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. After installation, use feet and the brackets to stabilize the Exalogic machine.

3.2.6.3 Stabilize the Exalogic Machine with Leveling Feet

The rack contains four leveling feet that can be lowered to stabilize the rack. The leveling feet can be used even when the rack is permanently secured to the floor. To adjust the leveling feet, do the following:

1. Locate the four leveling feet located at the bottom corners of the Exalogic machine. Figure 3-2 shows the location of the leveling feet on the bottom of the Exalogic machine.

Figure 3-2 Location of Leveling Feet on Exalogic Machine

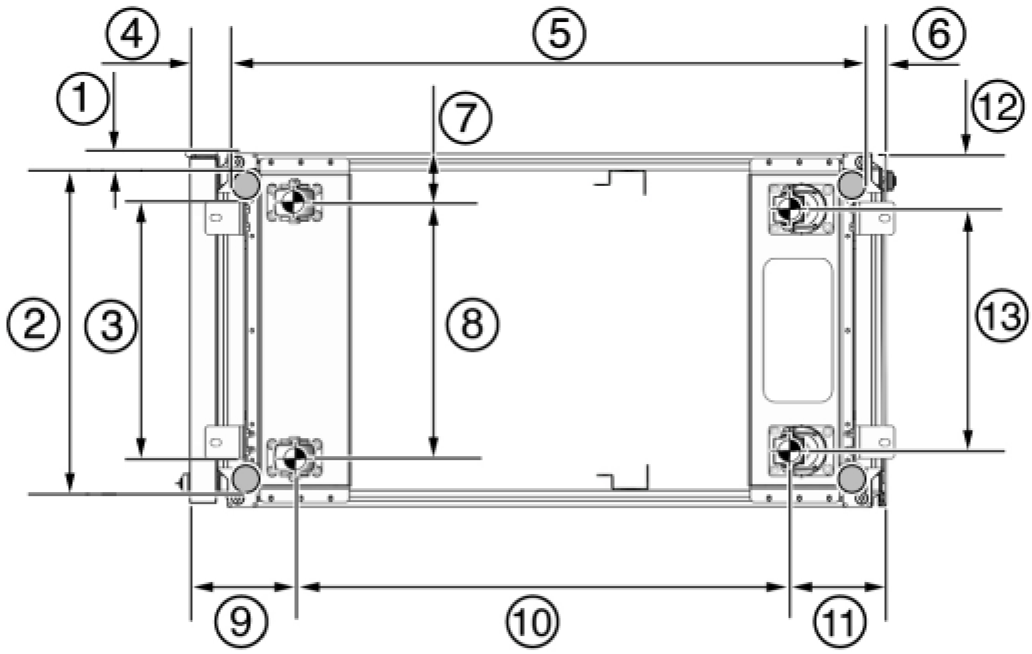

- 1: Distance from the edge of the mounting feet to the side of the rack is 33.75 mm (1.33 inches)
- 2: Width from the outside edges of the leveling feet is 532.5 mm (20.96 inches)
- 3: Width from the inside edges of the leveling feet is 429 mm (16.89 inches)
- 4: Distance from the edge of the feet to the front rack surface is 73.75 m (2.90 inches)
- 5: Depth of the outside edges of the leveling feet is 1058.5 mm (41.67 inches)
- 6: Distance from the edge of the leveling feet to the rear rack surface is 33.75 mm (1.33 inches)
- 7: Distance from the center of front casters to the side of the rack is 86.7 mm (3.41 inches)
- 8: Width between the center of the front casters is 426.6 mm (16.80 inches)
- 9: Distance from the center of the rear casters to the rear of the rack is 173.7 mm (6.83 inches)
- 10: Depth between the front and rear casters is 828.6 mm (32.62 inches)
- 11: Distance between the rear casters and the rear of the rack is 162.4 mm (6.39 inches)
- 12: Distance from the center of rear casters to the side of the rack is 96.4 mm (3.80 inches)
- 13: Width between the center of the rear casters is 407.2 mm (16.03 inches)

2. Lower the leveling feet to the floor as shown in Figure 3-3 using the SW 12 mm wrench. When lowered correctly, the four leveling feet should support the full weight of the Exalogic machine.

Figure 3-3 Securing Exalogic Machine Using the Leveling Feet

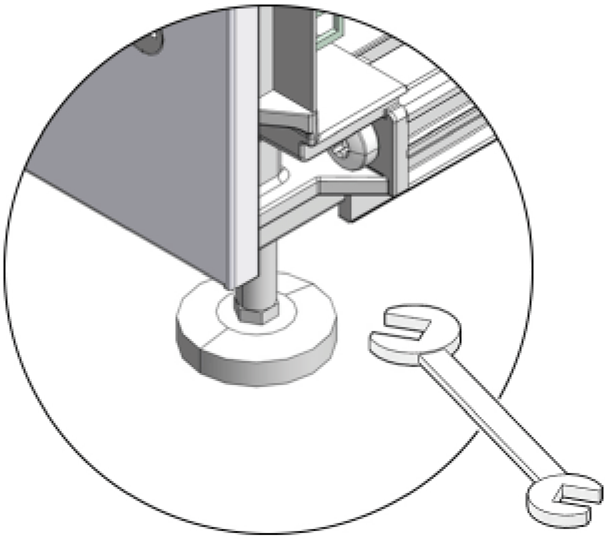

3.2.6.4 Stabilize the Exalogic Machine with Mounting Brackets

The rack can be permanently mounted to the installation site floor using the same mounting brackets that secured the rack to the shipping pallet. The rack is secured to the pallet with four mounting brackets. Use the front and rear brackets to stabilize the rack to the installation floor. Prepare the installation site by drilling four holes into the floor. Before permanently stabilizing the Exalogic machine with the mounting brackets, pre-drill the mounting holes. Figure 3-4 shows the location and dimensions of the mounting brackets.

Figure 3-4 Location of Mounting Brackets on Rack

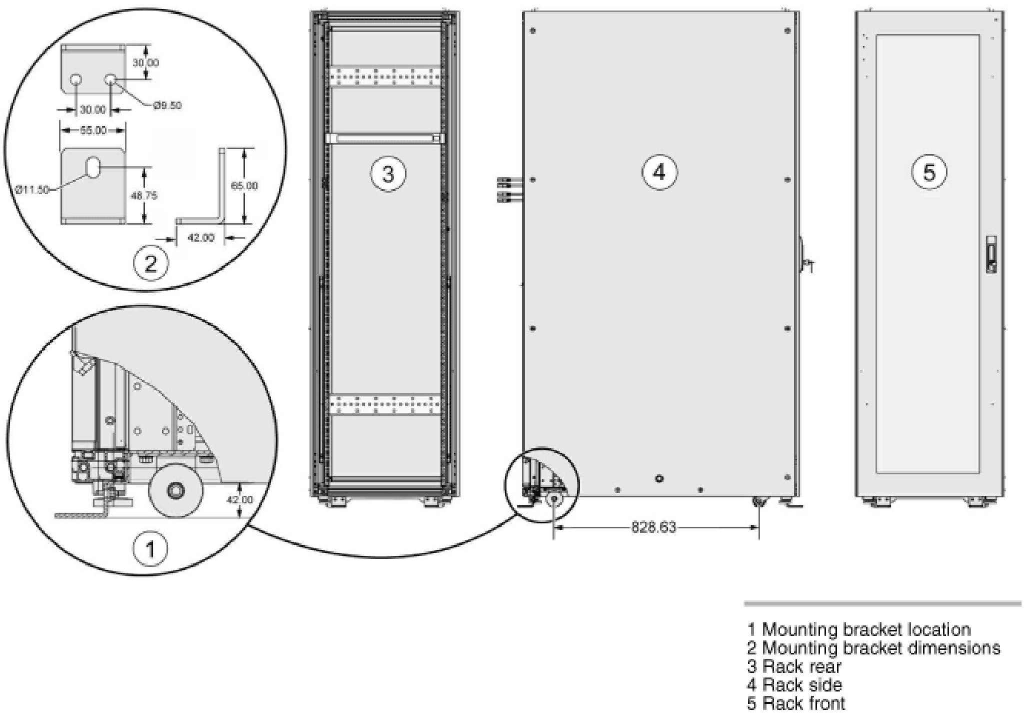

1. Obtain four bolts and washers to mount the Exalogic machine to the floor. The bolt holes in the mounting brackets have a 10.0 mm diameter. Oracle does not provide mounting bolts because different floors require different bolt types and strengths.
2. Position the Exalogic machine over the pre-drilled holes. Figure 3-5 shows the bottom view of the Exalogic machine, and the location for the mounting hole and floor cut dimensions.

Figure 3-5 Bottom View of Exalogic Machine Showing Mounting Hole and Floor Cutout Dimensions

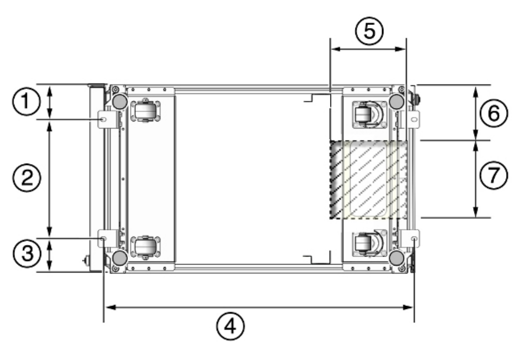

- 1: Distance from mounting bracket to the edge of the rack is 113 mm (4.45 inches)
- 2: Width between the centers of the mounting hole slots is 374 mm (14.72 inches)

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

If you plan to route data or PDU power cords down through the bottom of the rack, then you need to cut a hole in the installation floor site. Cut a rectangular hole below the rear portion of the rack, between the two rear casters and behind the rear RETMA (Radio Electronics Television Manufacturers Association) rails. Figure 3-6 shows the base positions of the rack from the bottom.

Figure 3-6 Base Position to Measure

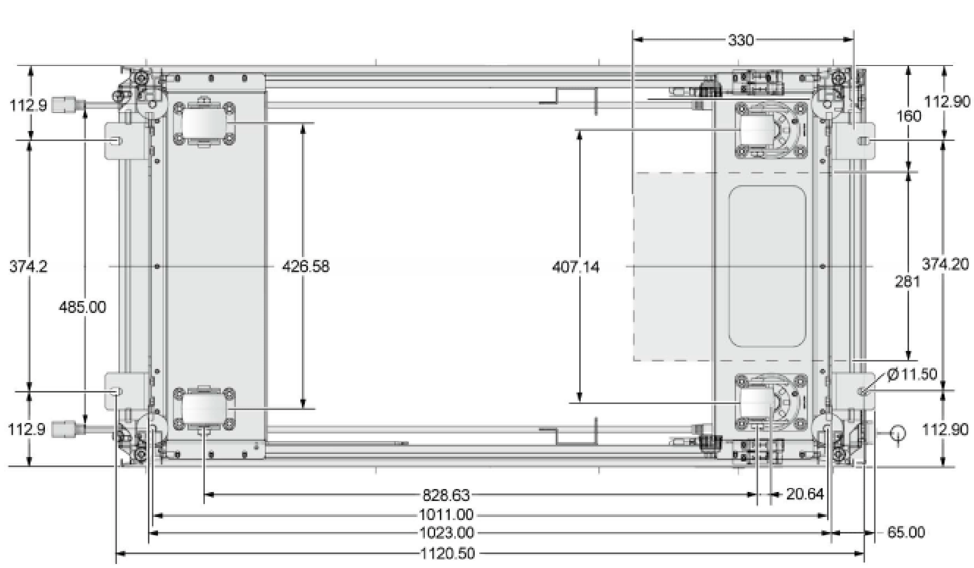

> **Caution:**
>
> Do not create a hole where the rack casters or leveling feet brackets will be placed.

3. Open the front and rear Exalogic machine doors.
4. Install the mounting brackets to the rack as shown in Figure 3-7 using a 6 mm hexagon Allen wrench key.

Figure 3-7 Securing Exalogic Machine Using Mounting Brackets

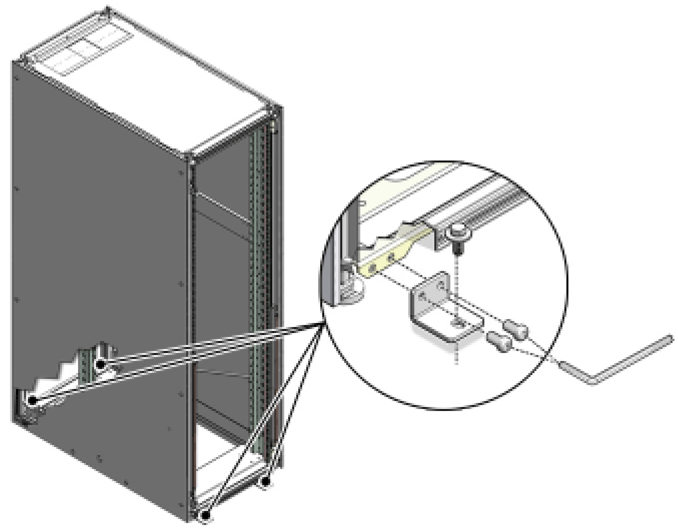

5. Using bolts and washers that are appropriate for your installation site, permanently mount your system to the floor using the four mounting brackets as shown in Figure 3-7.

> **Note:**
>
> The bolts required for securing the Exalogic machine to the floor vary depending on the installation location. Select bolts that are appropriate for your location.

6. Firmly tighten all of the bolts that secure the mounting brackets to the Exalogic machine and to the floor.

3.2.6.5 Attach a Ground Cable (Optional)

The Exalogic machine 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 Exalogic machine. 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.

Note:

A grounding cable is not shipped with the system.

1. Ensure the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
2. Ensure all grounding points, such as raised floors and power receptacles, reference the facility ground.
3. During manufacturing, the ground cable attachment area might be a painted or coated surface. Ensure that direct, metal-to-metal contact is made for this installation.
4. Attach the ground cable to one of the attachment points located at the bottom rear of the system frame as shown in Figure 3-8. The attachment point is an adjustable bolt that is inside the rear of the Exalogic machine cabinet on the right side.

Figure 3-8 Earth Ground Attachment Bolt Location

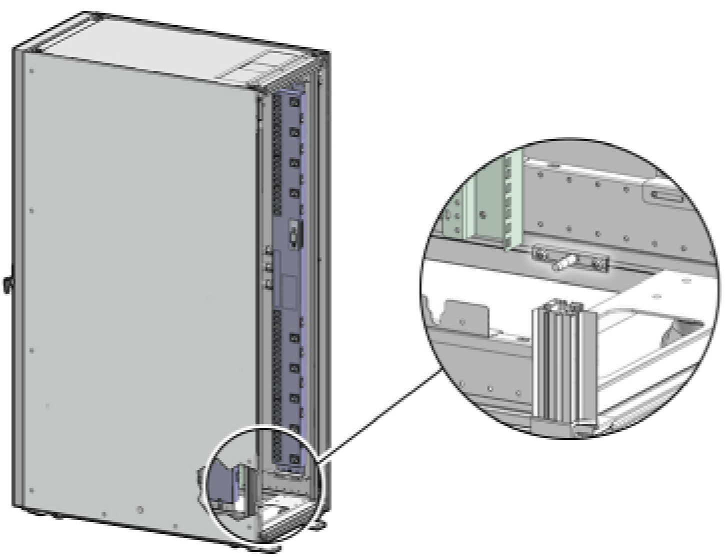

3.3 Power on the System the First Time

Before powering on the system for the first time, it is necessary to inspect the machine, and connect the power cords. This section contains the following topics:

- Inspect the Machine After It Is in Place
- Connect the Power Cords
- Power On the Exalogic Machine

3.3.1 Inspect the Machine After It Is in Place

The following procedure describes how to visually examine the Exalogic machine physical system after it is in place, but before power is supplied:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check your Exalogic machine 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 in to the front of the Exalogic machine.

> **See Also:**
>
> Ventilation and Cooling Requirementsfor more information

3.3.2 Connect the Power Cords

The following procedure describes how to connect power cords to the Exalogic machine:

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.
4. Route the power cords to the facility receptacles either above the rack or below the flooring, as shown in Figure 3-9 and Figure 3-10.
5. Secure the power cords in bundles, as shown in Figure 3-9.
6. Plug the PDU power cord connectors into the facility receptacles.

Figure 3-9   Power Cord Routing from the Bottom of the Rack

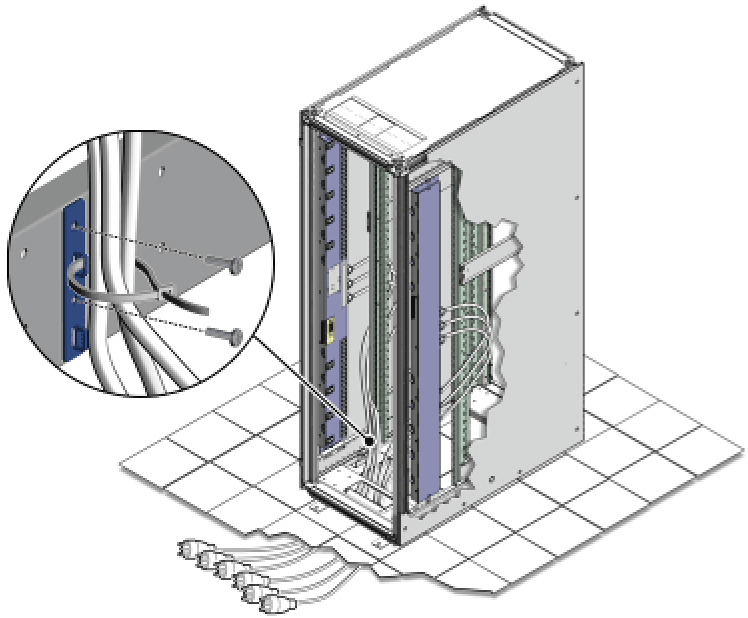

Figure 3-10   Power Cord Routing Example from the Top of the Rack

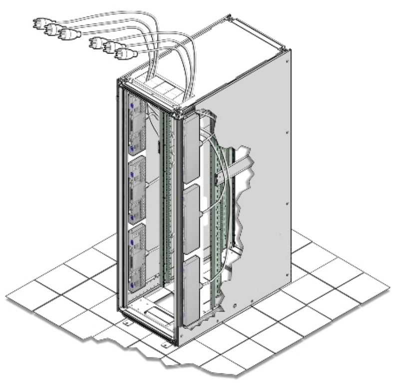

3.3.3 Power On the Exalogic Machine

The following procedure describes how to power on the Exalogic machine:

1. Ensure that each of the three main power cords is connected.
2. Switch on the six power distribution unit (PDU) circuit breakers located on the rear of the main PDU (PDU A) inside the Exalogic machine. The circuit breakers are on the rear of the Exalogic machine cabinet as shown in Figure 3-11. Press the ON (|) side of the toggle switch.

> **Note:**
>
> The 24 kVA PDU has 4 plugs, and the 15 kVA PDU has 2 plugs. In either case, you can use the instructions described in this section to power on your Exalogic machine.

Figure 3-11 PDU Switch Locations

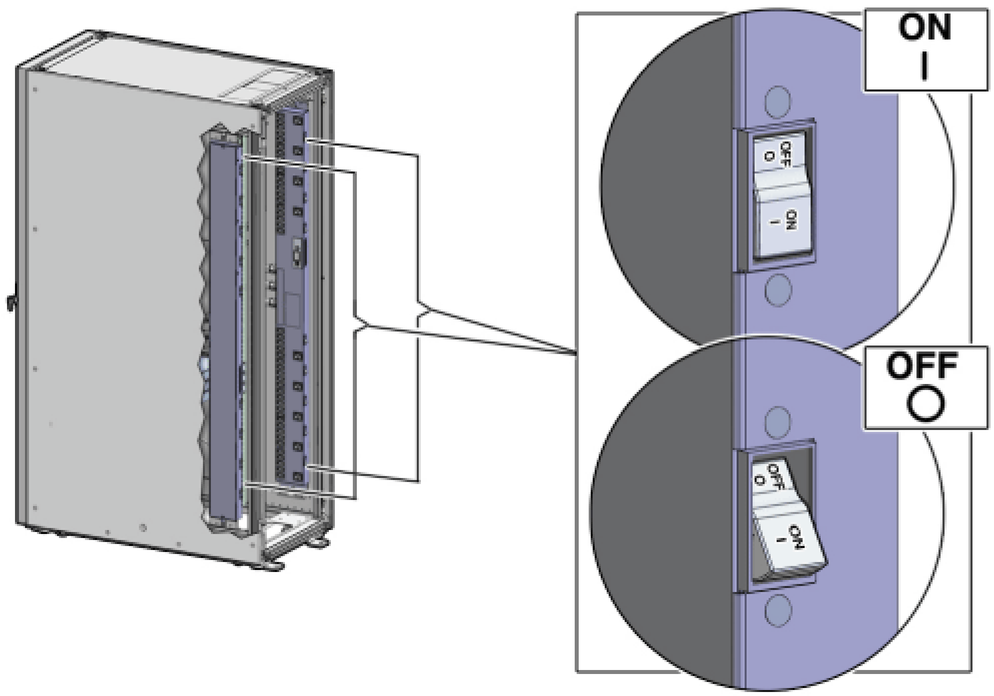

3. Wait 3 to 5 minutes for all Oracle Integrated Lights Out Manager (ILOM) service processors to boot.
4. Open the front cabinet door.
5. Verify that server standby power is on for compute nodes, InfiniBand gateway switches, and storage appliance in the Exalogic machine.

   When power is delivered to the receptacles at the rear of the server chassis, standby power is made available by the power supplies. When standby power is distributed to the chassis, the service processor (SP) powers on and starts up to manage the system. The main power is supplied for the remainder of the chassis and fans when the power button on the front of the server chassis is pressed. The power button is disabled while the SP is starting. The start-up state is indicated by a steady one second on, one second off blinking pattern of the Power/OK LED on the front of the system. After the SP has started, the power button is enabled and the system is placed in standby power mode.

   In standby power mode, the Power/OK LED on the front panel blinks green in a 0.1 second on, 2.9 seconds off pattern, as shown in Figure 3-12.

Figure 3-12   Exalogic Compute Node LED Lights

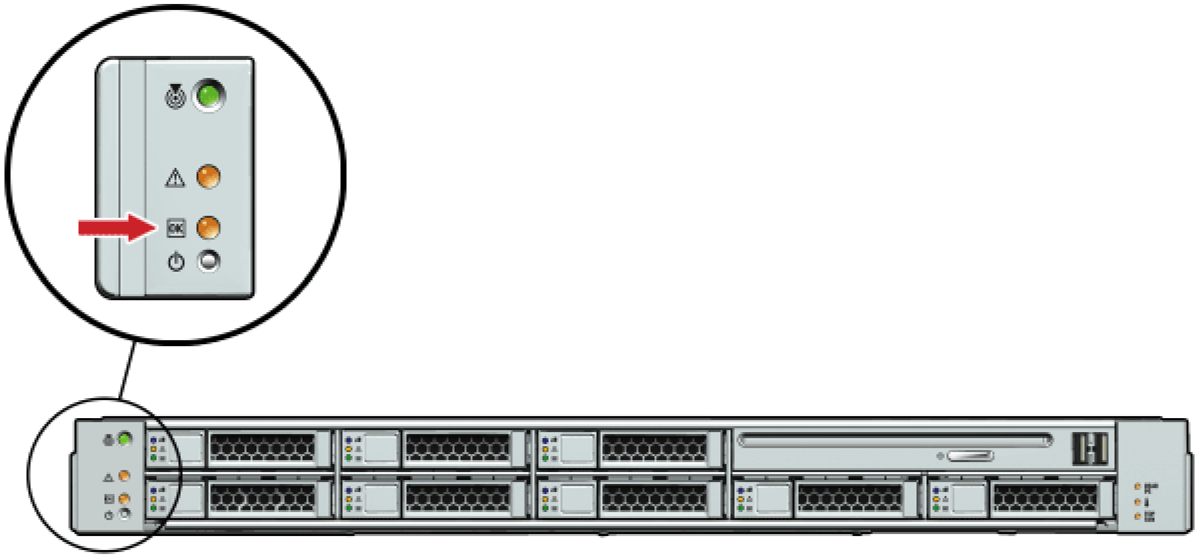

6. Verify that the main power is on for each compute node.

   In main power mode, the Power/OK LED on the front panel blinks in a one second on, one second off pattern while the system BIOS starts. After BIOS initialization completes and the operating system begins to start, the Power/OK LED illuminates and remains a steady green.

7. Press the soft switches located on the front of the two storage heads in the storage appliance, as shown in Figure 3-13. In addition, see Figure 3-14.

Figure 3-13   Soft Switches on the Server Heads in the Sun ZFS Storage 7320 appliance

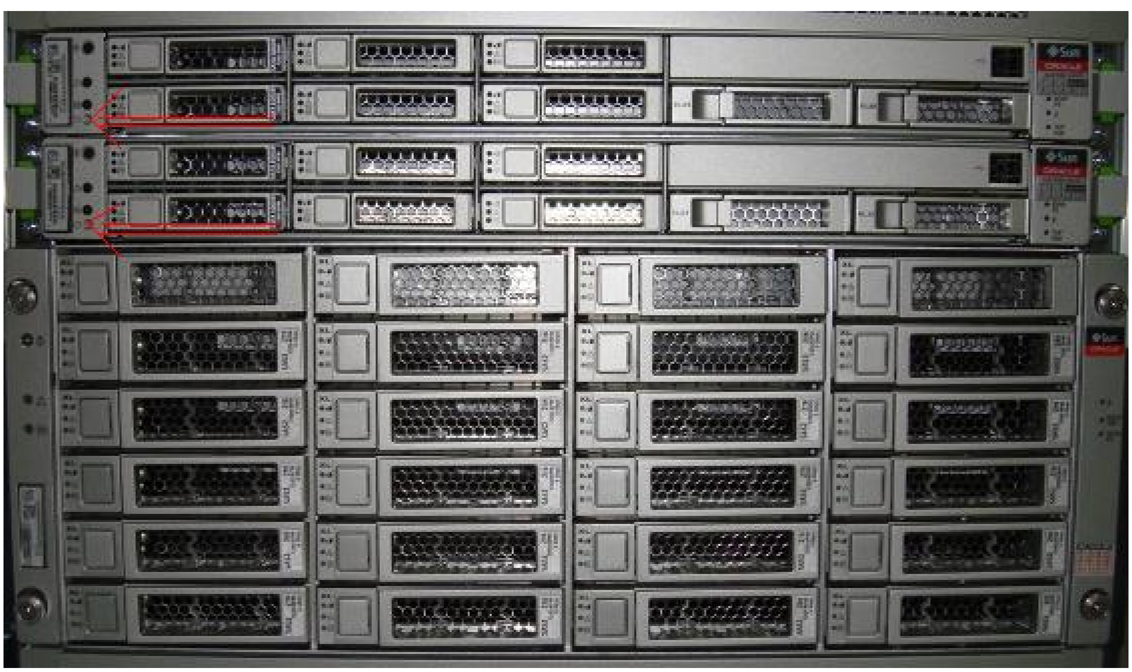

Figure 3-14 shows the soft switch and LED lights on each of the server heads in the storage appliance.

Figure 3-14 Soft Switch and LED Lights on Server Heads

8. Wait 3-5 minutes for the storage appliance to initiate NFS services, daemons, and basic services.

> **Tip:**
>
> You can ping the IP address assigned to the storage appliance to verify whether the system is up and running. For the default `NET0` IP addresses, see Default IP Addresses and Ports.
>
> Alternatively, you can try to launch the administration console for the storage appliance. Before you can ping the IP address or launch the administration console, you must connect a laptop to the rack, as described in Connect a Laptop to the Exalogic Machine.

9. After making sure that the storage appliance is up and running, press the soft switches located on the front of the compute nodes in your Exalogic machine, as shown in Figure 3-13. After power-on, the LED indicators turn green, as shown in Figure 3-14.

> **Note:**
>
> If you are using an Exalogic machine full rack, you must switch on thirty compute nodes. If you are using an Exalogic machine half rack, you must switch on sixteen compute nodes. If you are using an Exalogic machine quarter rack, you must switch on eight compute nodes.

10. Verify that power is applied to the Cisco Ethernet switch.

> **Note:**
>
> The Ethernet switch is inside a vented filler panel in Unit 25 of the rack.

11. Verify that power is applied to the Sun Network QDR InfiniBand Gateway Switches.

> **Note:**
>
> The gateway switches are inside vented filler panels. If you are using an Exalogic machine full rack, verify that power is supplied to the four gateway switches included in the Exalogic machine. If you are using an Exalogic machine half rack, verify that power is supplied to the two gateway switches included in the Exalogic machine. If you are using an Exalogic machine quarter rack, verify that power is supplied to the two gateway switches included in the Exalogic machine.

12. Optional: Verify that power is supplied to the Sun Datacenter InfiniBand Switch 36.

> **Note:**
>
> The Sun Datacenter InfiniBand Switch 36 is inside a vented filler panel when deployed. This spine switch is used in multirack configuration scenarios only and it must be purchased separately. A multirack configuration consists of an Exalogic machine connected to another Exalogic machine or an Oracle Exadata Database Machine.

3.4 Connect a Laptop to the Exalogic Machine

You can connect a laptop to the Exalogic machine as follows:

1. Ensure that you have a laptop with functional USB and network ports.
2. Ensure that you have a Category 5E patch cable of maximum length 25 feet and a serial cable of maximum length 15 feet.
3. Open the rear cabinet door of the rack.
4. Connect the network port of your laptop into an unused input port in the Cisco Ethernet switch. This switch is inside a vented filler panel in Unit 25 of your Exalogic machine rack. Note that you should not connect to any of the management or console ports on the switch. The ports are labeled on the switch.

> **Note:**
>
> If you require serial connectivity, you can use a USB-to-Serial adapter to connect from the USB port of your laptop to the Cisco switch.
>
> A USB-to-Serial adapter is installed in the rack on all of the gateway switches (Sun Network QDR InfiniBand Gateway Switches).
>
> An extra adapter is included in the shipping kit in the Exalogic machine full rack and half rack configurations.

5. If you have not booted the operating system on your laptop, start the operating system now.

   - If you are using the Windows operating system on your laptop, do the following:

     a. Go to **Control Panel** > **Network Connections**. Select your wired network adapter in the list of network connections, right-click and select **Properties**. The network properties screen is displayed.

     b. Click the **General** tab, and select **Internet Protocol (TCP/IP)**. Click **Properties**. The Internet Protocol (TCP/IP) Properties screen is displayed.

     c. Select the **Use the following IP address:** option, and enter a static IP address for your laptop. Although a default gateway is not necessary, enter the same IP address in the **Default Gateway** field. Click **OK** to exit the network connections screen.

     > **Note:**
     >
     > This static IP should be on the same subnet and address range as the network on which the Cisco Ethernet switch resides. You can use the default `NET0` IP addresses of compute nodes assigned at the time of manufacturing or the custom IP address that you reconfigured using the Oracle `Exalogic Configuration Utility`. For the list of default `NET0` IP addresses, see Default IP Addresses and Ports.

   - If you are using a Linux operating system on your laptop, do the following:

     a. Log in as a `root` user.

     b. At the command prompt, enter the following command to display the network devices, such as `ETH0`, attached to the Exalogic machine:

     ```
     # ifconfig -a
     ```

     The list of network devices or adapters attached to the Exalogic machine is displayed.

     c. To set up the desired network interface, run the `ifconfig` command at the command prompt, as in the following example:

     ```
     # ifconfig eth0 192.168.1.150 netmask 255.255.255.0 up
     ```

     In this example, the `ifconfig` command assigns the IPv4 address `192.168.1.150`, with a network mask of `255.255.255.0`, to the `eth0` interface.

6. For laptop connectivity, open any `telnet` or `ssh` client program, such as PuTTY. Connect to one of the service processor IP addresses or to the IP address of a compute node, which is up and running.

> **Note:**
>
> After you cable your laptop to the Cisco Ethernet switch, you can use the `NET0` IP addresses of Exalogic machine components to communicate with them. For a list of default IP addresses assigned at the time of manufacturing, see Default IP Addresses and Ports.
>
> If you have not run the Oracle `Exalogic Configuration Utility` set of tools and scripts to reconfigure IP addresses for the Exalogic machine, you can use a set of default IP addresses. If you have already run the Oracle `Exalogic Configuration Utility` set of tools and scripts, you can use the network IP address that you provided as input to the *Exalogic Configurator* spreadsheet. For more information about Oracle `Exalogic Configuration Utility`, see Configure the Exalogic Machine Using ECU.

3.5 What Next?

After powering on the Exalogic machine, including compute nodes and storage appliance, proceed to configure the Exalogic machine.

> **Note:**
>
> For information about changing passwords of Exalogic machine components, see MOS document 1594316.1 at:
>
> `https://support.oracle.com/epmos/faces/DocumentDisplay?id=1594316.1`

4

Default IP Addresses and Ports

This chapter lists the default ILOM, NET0, InfiniBand Bonded IP addresses, and ports assigned to Exalogic machine components during manufacturing.
This chapter contains the following sections:

> **Tip:**
>
> For more information about how these interfaces are used, see Figure 6-1.

4.1 Exalogic Machine Full Rack

Table 4-1 lists the default ILOM, NET0, and InfiniBand Bonded IP addresses assigned during manufacturing to Exalogic compute nodes and other hardware components in an Exalogic machine Full Rack.

Table 4-1 Default ILOM, NET0, and IB Bonded IP for Exalogic Full Rack

Unit	Component (Front View)	ILOM IP Address	NET0 IP Address	IB Bonded IP Address
42	Compute Node	192.168.1.132	192.168.1.32	192.168.10.32
41	Compute Node	192.168.1.131	192.168.1.31	192.168.10.31
40	Compute Node	192.168.1.130	192.168.1.30	192.168.10.30
39	Compute Node	192.168.1.129	192.168.1.29	192.168.10.29
38	Compute Node	192.168.1.128	192.168.1.28	192.168.10.28
37	Compute Node	192.168.1.127	192.168.1.27	192.168.10.27
36	Compute Node	192.168.1.126	192.168.1.26	192.168.10.26
35	Compute Node	192.168.1.125	192.168.1.25	192.168.10.25
34	Compute Node	192.168.1.124	192.168.1.24	192.168.10.24
33	Compute Node	192.168.1.123	192.168.1.23	192.168.10.23
32	Compute Node	192.168.1.122	192.168.1.22	192.168.10.22
31	Compute Node	192.168.1.121	192.168.1.21	192.168.10.21
30	Compute Node	192.168.1.120	192.168.1.20	192.168.10.20

Table 4-1 (Cont.) Default ILOM, NET0, and IB Bonded IP for Exalogic Full Rack

Unit	Component (Front View)	ILOM IP Address	NET0 IP Address	IB Bonded IP Address
29	Compute Node	192.168.1.119	192.168.1.19	192.168.10.19
28	Compute Node	192.168.1.118	192.168.1.18	192.168.10.18
27	Compute Node	192.168.1.117	192.168.1.17	192.168.10.17
26	Sun Network QDR InfiniBand Gateway Switch	192.168.1.204	Not applicable	Not applicable
25	Cisco Management Switch	Not applicable	192.168.1.200	Not applicable
24	Sun Network QDR InfiniBand Gateway Switch	192.168.1.203	Not applicable	Not applicable
23	Sun Network QDR InfiniBand Gateway Switch	192.168.1.202	Not applicable	Not applicable
22	Server head for the storage appliance	192.168.1.116	192.168.1.16	192.168.10.16
21	Server head for the storage appliance	192.168.1.115	192.168.1.15	192.168.10.15
20	Disk storage tray for the storage appliance	Not applicable	Not applicable	Not applicable
19				
18				
17				
16	Sun Network QDR InfiniBand Gateway Switch	192.168.1.201	Not applicable	Not applicable
15	Compute Node	192.168.1.114	192.168.1.14	192.168.10.14
14	Compute Node	192.168.1.113	192.168.1.13	192.168.10.13
13	Compute Node	192.168.1.112	192.168.1.12	192.168.10.12
12	Compute Node	192.168.1.111	192.168.1.11	192.168.10.11
11	Compute Node	192.168.1.110	192.168.1.10	192.168.10.10
10	Compute Node	192.168.1.109	192.168.1.9	192.168.10.9
9	Compute Node	192.168.1.108	192.168.1.8	192.168.10.8
8	Compute Node	192.168.1.107	192.168.1.7	192.168.10.7
7	Compute Node	192.168.1.106	192.168.1.6	192.168.10.6
6	Compute Node	192.168.1.105	192.168.1.5	192.168.10.5
5	Compute Node	192.168.1.104	192.168.1.4	192.168.10.4
4	Compute Node	192.168.1.103	192.168.1.3	192.168.10.3
3	Compute Node	192.168.1.102	192.168.1.2	192.168.10.2
2	Compute Node	192.168.1.101	192.168.1.1	192.168.10.1
1	Sun Datacenter InfiniBand Switch 36	192.168.1.205	Not applicable	Not applicable
-	PDU-A (left from rear view)	192.168.1.210	Not applicable	Not applicable
-	PDU-B (right from rear view)	192.168.1.211	Not applicable	Not applicable

Note:

The following table shows the hardware that is included in different Exalogic versions.

Exalogic Rack	Compute Node's Server	Cisco Management Switch	Storage Appliance
X6–2	Oracle Server X6-2	Cisco 4948E	Oracle ZS5-ES
X5–2	Oracle Server X5-2	Cisco 4948E	Oracle ZS3-ES
X4–2	Sun Server X4-2	Cisco 4948E-F-S	Oracle ZS3-ES
X3–2	Sun Server X3-2	Cisco 4948E-F-S	Sun ZFS 7320
X2–2	Sun Fire X4170 M2	Cisco 4948	Sun ZFS 7320

4.2 Exalogic Machine Half Rack

Table 4-2 lists the default ILOM, NET0, and InfiniBand Bonded IP addresses assigned during manufacturing to Exalogic compute nodes and other hardware components in an Exalogic machine half rack.

Table 4-2 Default ILOM, NET0, and IB Bonded IP for Exalogic Half Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	IB Bonded IP Address
42	4U Solid Filler	Not applicable	Not applicable	Not applicable
41		Not applicable	Not applicable	Not applicable
40		Not applicable	Not applicable	Not applicable
39		Not applicable	Not applicable	Not applicable
38	4U Solid Filler	Not applicable	Not applicable	Not applicable
37		Not applicable	Not applicable	Not applicable
36		Not applicable	Not applicable	Not applicable
35		Not applicable	Not applicable	Not applicable
34	4U Solid Filler	Not applicable	Not applicable	Not applicable
33		Not applicable	Not applicable	Not applicable
32		Not applicable	Not applicable	Not applicable
31		Not applicable	Not applicable	Not applicable

Table 4-2 (Cont.) Default ILOM, NET0, and IB Bonded IP for Exalogic Half Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	IB Bonded IP Address
30	2U Solid Filler	Not applicable	Not applicable	Not applicable
29		Not applicable	Not applicable	Not applicable
28	Compute Node	192.168.1.118	192.168.1.18	192.168.10.18
27	Compute Node	192.168.1.117	192.168.1.17	192.168.10.17
26	1U Solid Filler	Not applicable	Not applicable	Not applicable
25	Cisco Management Switch	Not applicable	192.168.1.200	Not applicable
24	1U Solid Filler	Not applicable	Not applicable	Not applicable
23	Sun Network QDR InfiniBand Gateway Switch	192.168.1.202	Not applicable	Not applicable
22	Server head for the storage appliance	192.168.1.116	192.168.1.16	192.168.10.16
21	Server head for the storage appliance	192.168.1.115	192.168.1.15	192.168.10.15
20 19 18 17	Disk storage tray for the storage appliance	Not applicable	Not applicable	Not applicable
16	Sun Network QDR InfiniBand Gateway Switch	192.168.1.201	Not applicable	Not applicable
15	Compute Node	192.168.1.114	192.168.1.14	192.168.10.14
14	Compute Node	192.168.1.113	192.168.1.13	192.168.10.13
13	Compute Node	192.168.1.112	192.168.1.12	192.168.10.12
12	Compute Node	192.168.1.111	192.168.1.11	192.168.10.11
11	Compute Node	192.168.1.110	192.168.1.10	192.168.10.10
10	Compute Node	192.168.1.109	192.168.1.9	192.168.10.9
9	Compute Node	192.168.1.108	192.168.1.8	192.168.10.8
8	Compute Node	192.168.1.107	192.168.1.7	192.168.10.7
7	Compute Node	192.168.1.106	192.168.1.6	192.168.10.6
6	Compute Node	192.168.1.105	192.168.1.5	192.168.10.5
5	Compute Node	192.168.1.104	192.168.1.4	192.168.10.4
4	Compute Node	192.168.1.103	192.168.1.3	192.168.10.3
3	Compute Node	192.168.1.102	192.168.1.2	192.168.10.2
2	Compute Node	192.168.1.101	192.168.1.1	192.168.10.1
1	Sun Datacenter InfiniBand Switch 36	192.168.1.205	Not applicable	Not applicable
-	PDU-A (left from rear view)	192.168.1.210	Not applicable	Not applicable
-	PDU-B (right from rear view)	192.168.1.211	Not applicable	Not applicable

Note:

The following table shows the hardware that is included in different Exalogic versions.

Exalogic Rack	Compute Node's Server	Cisco Management Switch	Storage Appliance
X6–2	Oracle Server X6-2	Cisco 4948E	Oracle ZS5-ES
X5–2	Oracle Server X5-2	Cisco 4948E	Oracle ZS3-ES
X4–2	Sun Server X4-2	Cisco 4948E-F-S	Oracle ZS3-ES
X3–2	Sun Server X3-2	Cisco 4948E-F-S	Sun ZFS 7320
X2–2	Sun Fire X4170 M2	Cisco 4948	Sun ZFS 7320

4.3 Exalogic Machine Quarter Rack

Table 4-3 lists the default ILOM, NET0, and InfiniBand Bonded IP addresses assigned during manufacturing to Exalogic compute nodes and other hardware components in an Exalogic machine quarter rack.

Table 4-3 Default ILOM, NET0, and IB Bonded IP for Exalogic Quarter Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	InfiniBand Bonded IP Address
42	4U Solid Filler	Not applicable	Not applicable	Not applicable
41		Not applicable	Not applicable	Not applicable
40		Not applicable	Not applicable	Not applicable
39		Not applicable	Not applicable	Not applicable
38	4U Solid Filler	Not applicable	Not applicable	Not applicable
37		Not applicable	Not applicable	Not applicable
36		Not applicable	Not applicable	Not applicable
35		Not applicable	Not applicable	Not applicable
34	4U Solid Filler	Not applicable	Not applicable	Not applicable
33		Not applicable	Not applicable	Not applicable
32		Not applicable	Not applicable	Not applicable

Table 4-3 (Cont.) Default ILOM, NET0, and IB Bonded IP for Exalogic Quarter Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	InfiniBand Bonded IP Address
31		Not applicable	Not applicable	Not applicable
30	2U Solid Filler	Not applicable	Not applicable	Not applicable
29		Not applicable	Not applicable	Not applicable
28	2U Solid Filler	Not applicable	Not applicable	Not applicable
27		Not applicable	Not applicable	Not applicable
26	1U Solid Filler	Not applicable	Not applicable	Not applicable
25	Cisco Management Switch	Not applicable	192.168.1.200	Not applicable
24	1U Solid Filler	Not applicable	Not applicable	Not applicable
23	Sun Network QDR InfiniBand Gateway Switch	192.168.1.202	Not applicable	Not applicable
22	Server head for the storage appliance	192.168.1.116	192.168.1.16	192.168.10.16
21	Server head for the storage appliance	192.168.1.115	192.168.1.15	192.168.10.15
20	Disk storage tray for the storage appliance	Not applicable	Not applicable	Not applicable
19				
18				
17				
16	Sun Network QDR InfiniBand Gateway Switch	192.168.1.201	Not applicable	Not applicable
15	2U Solid Filler	Not applicable	Not applicable	Not applicable
14		Not applicable	Not applicable	Not applicable
13	4U Solid Filler	Not applicable	Not applicable	Not applicable
12		Not applicable	Not applicable	Not applicable
11		Not applicable	Not applicable	Not applicable
10		Not applicable	Not applicable	Not applicable
9	Compute Node	192.168.1.108	192.168.1.8	192.168.10.8
8	Compute Node	192.168.1.107	192.168.1.7	192.168.10.7
7	Compute Node	192.168.1.106	192.168.1.6	192.168.10.6
6	Compute Node	192.168.1.105	192.168.1.5	192.168.10.5
5	Compute Node	192.168.1.104	192.168.1.4	192.168.10.4
4	Compute Node	192.168.1.103	192.168.1.3	192.168.10.3
3	Compute Node	192.168.1.102	192.168.1.2	192.168.10.2
2	Compute Node	192.168.1.101	192.168.1.1	192.168.10.1
1	1U Solid Filler	Not applicable	Not applicable	Not applicable
-	PDU-A (left from rear view)	192.168.1.210	Not applicable	Not applicable
-	PDU-B (right from rear view)	192.168.1.211	Not applicable	Not applicable

Note:

The following table shows the hardware that is included in different Exalogic versions.

Exalogic Rack	Compute Node's Server	Cisco Management Switch	Storage Appliance
X6–2	Oracle Server X6-2	Cisco 4948E	Oracle ZS5-ES
X5–2	Oracle Server X5-2	Cisco 4948E	Oracle ZS3-ES
X4–2	Sun Server X4-2	Cisco 4948E-F-S	Oracle ZS3-ES
X3–2	Sun Server X3-2	Cisco 4948E-F-S	Sun ZFS 7320
X2–2	Sun Fire X4170 M2	Cisco 4948	Sun ZFS 7320

4.4 Exalogic Machine Eighth Rack

Table 4-4 lists the default ILOM, NET0, and InfiniBand Bonded IP addresses assigned during manufacturing to Exalogic compute nodes and other hardware components in an Exalogic machine eighth rack.

Table 4-4 Default ILOM, NET0, and IB Bonded IP for Exalogic Eighth Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	InfiniBand Bonded IP Address
42	4U Solid Filler	Not applicable	Not applicable	Not applicable
41		Not applicable	Not applicable	Not applicable
40		Not applicable	Not applicable	Not applicable
39		Not applicable	Not applicable	Not applicable
38	4U Solid Filler	Not applicable	Not applicable	Not applicable
37		Not applicable	Not applicable	Not applicable
36		Not applicable	Not applicable	Not applicable
35		Not applicable	Not applicable	Not applicable
34	4U Solid Filler	Not applicable	Not applicable	Not applicable
33		Not applicable	Not applicable	Not applicable
32		Not applicable	Not applicable	Not applicable

Table 4-4 (Cont.) Default ILOM, NET0, and IB Bonded IP for Exalogic Eighth Rack

Unit	Rack Component (Front View)	ILOM IP Address	NET0 IP Address	InfiniBand Bonded IP Address
31		Not applicable	Not applicable	Not applicable
30	2U Solid Filler	Not applicable	Not applicable	Not applicable
29		Not applicable	Not applicable	Not applicable
28	2U Solid Filler	Not applicable	Not applicable	Not applicable
27		Not applicable	Not applicable	Not applicable
26	1U Solid Filler	Not applicable	Not applicable	Not applicable
25	Cisco Management Switch	Not applicable	192.168.1.200	Not applicable
24	1U Solid Filler	Not applicable	Not applicable	Not applicable
23	Sun Network QDR InfiniBand Gateway Switch	192.168.1.202	Not applicable	Not applicable
22	Server head for the storage appliance	192.168.1.116	192.168.1.16	192.168.10.16
21	Server head for the storage appliance	192.168.1.115	192.168.1.15	192.168.10.15
20	Disk storage tray for storage appliance	Not applicable	Not applicable	Not applicable
19				
18				
17				
16	Sun Network QDR InfiniBand Gateway Switch	192.168.1.201	Not applicable	Not applicable
15	2U Solid Filler	Not applicable	Not applicable	Not applicable
14		Not applicable	Not applicable	Not applicable
13	4U Solid Filler	Not applicable	Not applicable	Not applicable
12		Not applicable	Not applicable	Not applicable
11		Not applicable	Not applicable	Not applicable
10		Not applicable	Not applicable	Not applicable
9	4U Solid Filler	Not applicable	Not applicable	Not applicable
8		Not applicable	Not applicable	Not applicable
7		Not applicable	Not applicable	Not applicable
6		Not applicable	Not applicable	Not applicable
5	Compute Node	192.168.1.104	192.168.1.4	192.168.10.4
4	Compute Node	192.168.1.103	192.168.1.3	192.168.10.3
3	Compute Node	192.168.1.102	192.168.1.2	192.168.10.2
2	Compute Node	192.168.1.101	192.168.1.1	192.168.10.1
1	1U Solid Filler	Not applicable	Not applicable	Not applicable
-	PDU-A (left from rear view)	192.168.1.210	Not applicable	Not applicable
-	PDU-B (right from rear view)	192.168.1.211	Not applicable	Not applicable

Note:

The following table shows the hardware that is included in different Exalogic versions.

Exalogic Rack	Compute Node's Server	Cisco Management Switch	Storage Appliance
X6–2	Oracle Server X6-2	Cisco 4948E	Oracle ZS5-ES
X5–2	Oracle Server X5-2	Cisco 4948E	Oracle ZS3-ES
X4–2	Sun Server X4-2	Cisco 4948E-F-S	Oracle ZS3-ES
X3–2	Sun Server X3-2	Cisco 4948E-F-S	Sun ZFS 7320
X2–2	Sun Fire X4170 M2	Cisco 4948	Sun ZFS 7320

4.5 Default Port Assignments

Table 4-5 lists the default ports assigned to Exalogic compute nodes and other hardware components in an Exalogic machine.

Table 4-5 Default Ports

Source	Target	Protocol	Port	Application
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	SSH over TCP	22	SSH
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	HTTP over TCP	80	Web (user configurable)
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	SNMP over UDP	161	SNMP (Simple Network Management Protocol) (user configurable)

Table 4-5 (Cont.) Default Ports

Source	Target	Protocol	Port	Application
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	LDAP over UDP/TCP	389	Outgoing LDAP (Lightweight Directory Access Protocol) (user configurable)
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	HTTPS over TCP	443	Web (user configurable)
Any	Compute nodes, storage appliance heads, and InfiniBand ILOMs	IPMI over UDP	623	IPMI (Intelligent Platform Management Interface)
Any	Compute nodes, and storage appliance heads ILOMs	TCP	5120	ILOM remote console: CD
Any	Compute nodes, and storage appliance heads ILOMs	TCP	5121	ILOM remote console: keyboard and mouse
Any	Compute nodes, and storage appliance heads ILOMs	TCP	5123	ILOM remote console: diskette
Any	Compute nodes, and storage appliance heads ILOMs	TCP	5555	ILOM remote console: encryption
Any	Compute nodes, and storage appliance heads ILOMs	TCP	6481	ILOM remote console: Servicetag daemon
Any	Compute nodes, and storage appliance heads ILOMs	TCP	5556	ILOM remote console: authentication
Any	Compute nodes, and storage appliance heads ILOMs	TCP	7578	ILOM remote console: video
Any	Compute nodes, and storage appliance heads ILOMs	TCP	7579	ILOM remote console: serial
ASR Manager	Compute nodes, and storage appliance heads ILOMs	HTTP	6481	Service tags listener for asset activation
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	TFTP over UDP	69	Outgoing TFTP (Trivial File Transfer Protocol)
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	NTP over UDP	123	Outgoing NTP

Table 4-5 (Cont.) Default Ports

Source	Target	Protocol	Port	Application
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	IPMI over UDP	162	Outgoing IPMI (Intelligent Platform Management Interface) Platform Event Trap (PET)
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	Syslog over UDP	514	Outgoing Syslog
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	DHCP over UDP	546	Client DHCP (Dynamic Host Configuration Protocol)
Compute nodes, storage appliance heads, and InfiniBand ILOMs	Any	RADIUS over UDP	1812	Outgoing RADIUS (Remote Authentication Dial In User Service) (user configurable)
Compute nodes, and storage appliance heads ILOMs	ASR Manager	SNMP	162	Telemetry messages sent to ASR Manager
Any	Storage appliance	HTTP over TCP	215	Browser interface
Storage management	Any	NTP over UDP	123	Outgoing NTP
Any	PDU	HTTP over TCP	80	Browser interface
Any	PDU	SNMP over UDP	161	SNMP (user configurable)
Any	PDU	HTTPS over TCP	443	Browser interface
PDU	Any	SNMP over UDP	162	Outgoing SNMPv2 traps
PDU	Any	Syslog over UDP	514	Outgoing Syslog
PDU	Any	DHCP over UDP	546	DHCP (Dynamic Host Configuration Protocol) client

5 Add the Exalogic Machine to Your Network

This chapter discusses the factory configuration and initial network configuration for the Exalogic machine.
This chapter contains the following topics:

- Default State of the Exalogic Machine Network Configuration
- Verify the Factory Configuration
- Prerequisites for Connecting Exalogic Machine to External Ethernet Network
- Initial Network Configuration of Exalogic Machine

5.1 Default State of the Exalogic Machine Network Configuration

At the time of manufacturing in the factory, an Exalogic machine has the following network topology:

- Compute nodes configured with `NET0`, `ILOM`, and `BOND0` interfaces

> **Note:**
>
> Exalogic compute nodes are not pre-configured with the `BOND1` interface for the Ethernet over InfiniBand (EoIB) connectivity. You can configure `BOND1` for compute notes during the initial configuration of the Exalogic machine.

- Storage heads configured with `NET0`, `ILOM`, and `BOND0` interfaces
- Sun Network QDR InfiniBand Gateway Switches, referred to as leaf switches, configured with the `ILOM` interface
- Sun Datacenter InfiniBand Switch 36, referred to as the spine switch, not configured

> **Note:**
>
> This switch is used in multirack configuration scenarios only.

- Cisco Ethernet Management Switch, which is not configured

5.2 Verify the Factory Configuration

You can verify the factory configuration of your Exalogic machine by viewing the `/etc/exalogic.conf` file on the compute nodes.

The file provides default configuration information, such as the following:

```
RACK_TYPE=3
NODE_INDEX=19
RACK_NAME=el01
CNODE_PREFIX=cn
SNODE_PREFIX=sn
DNS_SERVER_LIST=
DOMAIN_SEARCH_ORDER=
DOMAIN_NAME=abc.example.com
eth0_BOOTPROTO=static
eth0_NETMASK=255.255.252.0
eth0_GATEWAY=10.11.12.13
bond0_NETMASK=255.255.255.0
has_bond1=0
SPNET0_NETMASK=
SPNET0_GATEWAY=
cnode1_host_name=el01cn01
cnode1_private_host_name=el01cn01-priv
cnode1_eth0_IPADDR=192.168.1.1
cnode1_bond0_IPADDR=192.168.10.1
cnode1_SPNET0_IPADDR=192.168.1.101
cnode2_host_name=el01cn02
cnode2_private_host_name=el01cn02-priv
cnode2_eth0_IPADDR=192.168.1.2
cnode2_bond0_IPADDR=192.168.10.2
cnode2_SPNET0_IPADDR=192.168.1.102
cnode3_host_name=el01cn03
cnode3_private_host_name=el01cn03-priv
cnode3_eth0_IPADDR=192.168.1.3
cnode3_bond0_IPADDR=192.168.10.3
cnode3_SPNET0_IPADDR=192.168.1.103
cnode4_host_name=el01cn04
cnode4_private_host_name=el01cn04-priv
cnode4_eth0_IPADDR=192.168.1.4
cnode4_bond0_IPADDR=192.168.10.4
cnode4_SPNET0_IPADDR=192.168.1.104
cnode5_host_name=el01cn05
cnode5_private_host_name=el01cn05-priv
cnode5_eth0_IPADDR=192.168.1.5
cnode5_bond0_IPADDR=192.168.10.5
cnode5_SPNET0_IPADDR=192.168.1.105
cnode6_host_name=el01cn06
cnode6_private_host_name=el01cn06-priv
cnode6_eth0_IPADDR=192.168.1.6
cnode6_bond0_IPADDR=192.168.10.6
cnode6_SPNET0_IPADDR=192.168.1.106
cnode7_host_name=el01cn07
cnode7_private_host_name=el01cn07-priv
cnode7_eth0_IPADDR=192.168.1.7
cnode7_bond0_IPADDR=192.168.10.7
cnode7_SPNET0_IPADDR=192.168.1.107
cnode8_host_name=el01cn08
cnode8_private_host_name=el01cn08-priv
cnode8_eth0_IPADDR=192.168.1.8
cnode8_bond0_IPADDR=192.168.10.8
cnode8_SPNET0_IPADDR=192.168.1.108
cnode9_host_name=el01cn09
cnode9_private_host_name=el01cn09-priv
cnode9_eth0_IPADDR=192.168.1.9
cnode9_bond0_IPADDR=192.168.10.9
```

```
cnode9_SPNET0_IPADDR=192.168.1.109
cnode10_host_name=el01cn10
cnode10_private_host_name=el01cn10-priv
cnode10_eth0_IPADDR=192.168.1.10
cnode10_bond0_IPADDR=192.168.10.10
cnode10_SPNET0_IPADDR=192.168.1.110
cnode11_host_name=el01cn11
cnode11_private_host_name=el01cn11-priv
cnode11_eth0_IPADDR=192.168.1.11
cnode11_bond0_IPADDR=192.168.10.11
cnode11_SPNET0_IPADDR=192.168.1.111
cnode12_host_name=el01cn12
cnode12_private_host_name=el01cn12-priv
cnode12_eth0_IPADDR=192.168.1.12
cnode12_bond0_IPADDR=192.168.10.12
cnode12_SPNET0_IPADDR=192.168.1.112
cnode13_host_name=el01cn13
cnode13_private_host_name=el01cn13-priv
cnode13_eth0_IPADDR=192.168.1.13
cnode13_bond0_IPADDR=192.168.10.13
cnode13_SPNET0_IPADDR=192.168.1.113
cnode14_host_name=el01cn14
cnode14_private_host_name=el01cn14-priv
cnode14_eth0_IPADDR=192.168.1.14
cnode14_bond0_IPADDR=192.168.10.14
cnode14_SPNET0_IPADDR=192.168.1.114
cnode15_host_name=el01cn15
cnode15_private_host_name=el01cn15-priv
cnode15_eth0_IPADDR=192.168.1.17
cnode15_bond0_IPADDR=192.168.10.17
cnode15_SPNET0_IPADDR=192.168.1.117
cnode16_host_name=el01cn16
cnode16_private_host_name=el01cn16-priv
cnode16_eth0_IPADDR=192.168.1.18
cnode16_bond0_IPADDR=192.168.10.18
cnode16_SPNET0_IPADDR=192.168.1.118
cnode17_host_name=el01cn17
cnode17_private_host_name=el01cn17-priv
cnode17_eth0_IPADDR=192.168.1.19
cnode17_bond0_IPADDR=192.168.10.19
cnode17_SPNET0_IPADDR=192.168.1.119
cnode18_host_name=el01cn18
cnode18_private_host_name=el01cn18-priv
cnode18_eth0_IPADDR=192.168.1.20
cnode18_bond0_IPADDR=192.168.10.20
cnode18_SPNET0_IPADDR=192.168.1.120
cnode19_host_name=el01cn19
cnode19_private_host_name=el01cn19-priv
cnode19_eth0_IPADDR=192.168.1.21
cnode19_bond0_IPADDR=192.168.10.21
cnode19_SPNET0_IPADDR=192.168.1.121
cnode20_host_name=el01cn20
cnode20_private_host_name=el01cn20-priv
cnode20_eth0_IPADDR=192.168.1.22
cnode20_bond0_IPADDR=192.168.10.22
cnode20_SPNET0_IPADDR=192.168.1.122
cnode21_host_name=el01cn21
cnode21_private_host_name=el01cn21-priv
cnode21_eth0_IPADDR=192.168.1.23
cnode21_bond0_IPADDR=192.168.10.23
cnode21_SPNET0_IPADDR=192.168.1.123
```

```
cnode22_host_name=el01cn22
cnode22_private_host_name=el01cn22-priv
cnode22_eth0_IPADDR=192.168.1.24
cnode22_bond0_IPADDR=192.168.10.24
cnode22_SPNET0_IPADDR=192.168.1.124
cnode23_host_name=el01cn23
cnode23_private_host_name=el01cn23-priv
cnode23_eth0_IPADDR=192.168.1.25
cnode23_bond0_IPADDR=192.168.10.25
cnode23_SPNET0_IPADDR=192.168.1.125
cnode24_host_name=el01cn24
cnode24_private_host_name=el01cn24-priv
cnode24_eth0_IPADDR=192.168.1.26
cnode24_bond0_IPADDR=192.168.10.26
cnode24_SPNET0_IPADDR=192.168.1.126
cnode25_host_name=el01cn25
cnode25_private_host_name=el01cn25-priv
cnode25_eth0_IPADDR=192.168.1.27
cnode25_bond0_IPADDR=192.168.10.27
cnode25_SPNET0_IPADDR=192.168.1.127
cnode26_host_name=el01cn26
cnode26_private_host_name=el01cn26-priv
cnode26_eth0_IPADDR=192.168.1.28
cnode26_bond0_IPADDR=192.168.10.28
cnode26_SPNET0_IPADDR=192.168.1.128
cnode27_host_name=el01cn27
cnode27_private_host_name=el01cn27-priv
cnode27_eth0_IPADDR=192.168.1.29
cnode27_bond0_IPADDR=192.168.10.29
cnode27_SPNET0_IPADDR=192.168.1.129
cnode28_host_name=el01cn28
cnode28_private_host_name=el01cn28-priv
cnode28_eth0_IPADDR=192.168.1.30
cnode28_bond0_IPADDR=192.168.10.30
cnode28_SPNET0_IPADDR=192.168.1.130
cnode29_host_name=el01cn29
cnode29_private_host_name=el01cn29-priv
cnode29_eth0_IPADDR=192.168.1.31
cnode29_bond0_IPADDR=192.168.10.31
cnode29_SPNET0_IPADDR=192.168.1.131
cnode30_host_name=el01cn30
cnode30_private_host_name=el01cn30-priv
cnode30_eth0_IPADDR=192.168.1.32
cnode30_bond0_IPADDR=192.168.10.32
cnode30_SPNET0_IPADDR=192.168.1.132
snode1_host_name=el01sn01
snode1_private_host_name=el01sn01-priv
snode1_eth0_IPADDR=192.168.1.15
snode1_bond0_IPADDR=192.168.10.15
snode1_SPNET0_IPADDR=192.168.1.115
snode2_host_name=el01sn02
snode2_private_host_name=el01sn02-priv
snode2_eth0_IPADDR=192.168.1.16
snode2_bond0_IPADDR=192.168.10.16
snode2_SPNET0_IPADDR=192.168.1.116
```

5.3 Prerequisites for Connecting Exalogic Machine to External Ethernet Network

Depending on the type of your Ethernet device, you should use cables and transceivers to connect an Exalogic machine to your existing network. For more information, see Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch.

5.4 Initial Network Configuration of Exalogic Machine

After you connect your Exalogic machine to the external Ethernet network using cables and transceivers, you must configure the following Exalogic components to add them to the network:

1. Mandatory: Cisco Ethernet Switch
2. Optional: vNICs on Sun Network QDR InfiniBand Gateway Switches for Ethernet over InfiniBand (EoIB) connectivity, if necessary
3. Mandatory: Compute nodes
4. Mandatory: Storage heads
5. Mandatory: Sun Network QDR InfiniBand Gateway Switches
6. Optional: Sun Datacenter InfiniBand Switch 36, if you are connecting your Exalogic machine to another Exalogic machine rack or to an Oracle Exadata Database Machine rack

> **Note:**
>
> You can use the Oracle **Exalogic Configuration Utility** set of tools and scripts to configure the compute nodes, storage heads, and Sun Network QDR InfiniBand Gateway Switches. For more information about **Exalogic Configuration Utility**, see Configure the Exalogic Machine Using ECU.
>
> You must configure the Cisco Ethernet Switch manually. For more information, see Configuring the Cisco Ethernet Switch.
>
> For information about configuring VNICs for Ethernet connectivity, see Configure Ethernet Over InfiniBand.
>
> For information about configuring the Sun Datacenter InfiniBand Switch 36 manually in multirack scenarios, see Use the Sun Datacenter InfiniBand Switch 36 in Multirack Configurations.

5.4.1 Configuring the Cisco Ethernet Switch

The host name, IP address, DNS and NTP configurations must be configured for the Cisco Ethernet switch. Before configuring the switch, note the following:

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

The following procedure describes how to configure the Cisco Ethernet switch:

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An RJ45 to DB9 serial cable is included in the Cisco documentation package.
2. Ensure that the terminal session is logged on the laptop by scripting the output. The data can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

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

3. Change to enable mode using the following command. The Cisco switch prompts for the default password. For the default password, contact Oracle Support.

   ```
   Switch> enable
   ```

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# *Sep 15 14:12:06.309:%SYS-5-CONFIG_I:Configured from console by console
   Switch# write memory
   Building configuration...
   Compressed configuration from 2474 bytes to 1066 bytes [OK ]
   ```

5. (Optional) Disable the default IP routing setting, and configure the default gateway as follows. This step is required if IP routing will not be used on the switch.

   ```
   Switch# configure terminal
   Enter configuration commands,one per line.End with CNTL/Z.
   Switch(config)# no ip routing
   Switch(config)# ip default-gateway 10.7.7.1
   Switch(config)# end
   *Sep 15 14:12:46.309:%SYS-5-CONFIG_I:Configured from console by console
   Switch# write memory
   Building configuration...
   Compressed configuration from 2492 bytes to 1070 bytes [OK ]
   ```

Note:

If IP routing is required on the switch, then leave the IP routing setting as the default, 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
*Sep 15 14:13:26.013:%SYS-5-CONFIG_I:Configured from console by console
Switch# write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

6. Set the host name of the switch as follows:

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# hostname corxsw-ip
burxsw-ip(config)# end
burxsw-ip# write memory
Building configuration...
Compressed configuration from 3789 bytes to 1469 bytes [OK ]
```

   The system host name will be used as the prompt name.

7. Set the password as follows:

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# enable password Your_Password
Switch(config)# enable secret Your_Password
The enable secret you have chosen is thesame as your enable password.
This is not recommended.Re-enter the enable secret.
Switch(config)# end
Switch#write memory
*Sep 15 14:25:05.893:%SYS-5-CONFIG_I:Configured from console by console
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

8. Set the password for telnet network access. The following is an example:

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# line vty 0 15
Switch(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
...
%Login disabled on line 15,until 'password'is set
%Login disabled on line 16,until 'password'is set
Switch(config-line)# password Your_Password
Switch(config-line)# login
Switch(config-line)# end
Switch#write memory
Building configuration...
Compressed configuration from 3786 bytes to 1468 bytes [OK ]
Switch#
```

In the preceding example, the first `login` output shows the password has not been set, and telnet access is disabled. If the `login` command returns nothing, then the password is set and telnet access is available.

9. Configure up to three DNS servers. The following is an example:

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

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

    - To use UTC use the following command:

      ```
      no clock timezone global configuration
      ```

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

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

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

    - To set summer time hours, use the following command:

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

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

    - To manually set the clock to any time, use the following command:

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

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

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

    ```
    Switch# configure terminal
    Enter configuration commands,one per line.End with CNTL/Z.
    Switch(config)# clock timezone EST -5
    Switch(config)# clock summer-time EDT recurring
    Switch(config)# end
    Switch# clock set 21:00:00 December 09 2009
    Switch#write memory
    Building configuration...
    Compressed configuration from 3784 bytes to 1465 bytes [OK ]
    ```

```
Switch# show clock
21:00:06.643 EST Wed Dec 9 2009
```

11. Configure up to two NTP servers. The following is an example:

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# ntp server 10.7.7.32 prefer
Switch(config)# ntp server 129.148.9.19
Switch(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by
console
Switch# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
Switch# show ntp status
<output will vary per network>
Switch# show clock
20:59:06.643 EST Wed Dec 9 2009
```

The preceding should show the NTP server synchronized to local time if the Cisco switch is connected to the network and has access to NTP.

12. Verify the configuration using the following command:

```
Switch# 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
!
hostname Switch
!
boot-start-marker
boot-end-marker
!
enable secret 5 $1$mS8h$EaJrIECUxavfGH6vLZg1T.
enable password Your_Password
!
no aaa new-model
ip subnet-zero
ip domain-name sodm.com
ip name-server 10.7.7.3
ip name-server 172.16.5.5
ip name-server 10.8.160.1
!
ip vrf mgmtVrf
!
vtp mode transparent
!
power redundancy-mode redundant
!
!
spanning-tree mode pvst
spanning-tree extend system-id
```

```
no spanning-tree vlan 1
!
vlan internal allocation policy ascending
!
!
interface FastEthernet1
ip vrf forwarding mgmtVrf
no ip address
speed auto
duplex auto
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
...
!
interface GigabitEthernet1/44
!
interface GigabitEthernet1/45
media-type rj45
!
interface GigabitEthernet1/46
media-type rj45
!
interface GigabitEthernet1/47
media-type rj45
!
interface GigabitEthernet1/48
media-type rj45
!
interface Vlan1
ip address 10.7.7.34 255.255.255.0
!
interface Vlan48
no ip address
shutdown
!
ip default-gateway 10.7.7.1
ip http server
!
!
control-plane
!
!
line con 0
stopbits 1
line vty 0 4
password Your_Password
login
line vty 5 15
password Your_Password
login
!
ntp server 10.7.7.32 prefer
ntp server 172.16.9.1
end
```

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, the following commands would be entered:

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

13. Save the current configuration using the following command:

```
Switch# copy running-config startup-config
```

14. Exit from the session using the following command:

```
Switch# exit
```

15. Disconnect the cable from the Cisco console.

> **Note:**
>
> To check the configuration, attach a laptop computer to port 48, and ping the IP address of the internal management network.

6 Understand Network Requirements and Configuration

This chapter describes the network requirements for the Exalogic machine.
It contains the following topics:

6.1 Overview of Network Requirements

An Exalogic machine includes compute nodes, storage appliance, and equipment to connect the compute nodes to your network. The network connections allow the servers to be administered remotely, enable clients to connect to the compute nodes, and enable client access to the storage appliance.

The following table describes the network components and interfaces for each compute node and the storage appliance:

Table 6-1 Available network components and interfaces on the compute nodes and storage appliance

	Compute Node	Storage Appliance (two server heads)
Gigabit Ethernet (GbE) ports: • Oracle Linux: `NET0`, `NET1`, `NET2`, and `NET3` • Oracle Solaris: `igb0`, `igb1`, `igb2`, and `igb3`	4 (only NET0 or igb0 is connected and used)	4 per server head (1 and 10 GbE ports available for Exalogic X4 and newer systems, 1 GbE for earlier systems)
Dual-port QDR InfiniBand Host Channel Adapter: • Oracle Linux: `ib0` and `ib1` • Oracle Solaris: `ibp0` and `ibp1`	1 (this port is not connected or used)	1 per server head

Table 6-1 (Cont.) Available network components and interfaces on the compute nodes and storage appliance

	Compute Node	Storage Appliance (two server heads)
Ethernet Port for ILOM remote management	1	4 per head (the ETH0 and ETH1 interfaces are used for active and passive clustering support; the dedicated ILOM port is not used, sideband management is used instead through the igb0 port)

Note:

These ports are pre-wired in the Exalogic machine at the time of manufacturing. Do not touch or modify the ports.

The Cisco Ethernet switch supplied with the Exalogic machine is minimally configured during installation. The minimal configuration disables IP routing, and sets the following:

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

Additional configuration, such as defining multiple virtual local area networks (VLANs) or enabling routing, may be required for the switch to operate properly in your environment and is beyond the scope of the installation service.

To deploy the Exalogic machine, verify that you meet the minimum network requirements. There are up to five networks for an Exalogic machine. Each network must be on a distinct and separate subnet from the others. The network descriptions are as follows:

- **Management network**: This required network connects to your existing management network, and is used for administrative work for all components of the Exalogic machine. It connects ILOM, compute nodes, server heads in the storage appliance, switches connected to the Ethernet switch in the Exalogic machine rack. This management network is in a single subnet. ILOM connectivity uses the `NET0` (on Oracle Solaris, `igb0`) sideband interface.

  For multirack configurations, you may have any of the following:

  - A single subnet per configuration
  - A single subnet per rack in the multirack configuration
  - Multiple subnets per configuration

  Oracle recommends that you configure a single subnet per configuration.

  With sideband management, only the `NET0` (on Oracle Solaris, `igb0`) interface of each compute node is physically connected to the Ethernet switch on the rack. For the server heads in the storage appliance, `NET0` and `NET1` interfaces (on Oracle Solaris, `igb0` and `igb1`) are physically connected to support active-passive clustering.

  > **Note:**
  >
  > Do not use the management network interface (`NET0` on Oracle Linux, and `igb0` on Oracle Solaris) on compute nodes for client or application network traffic. Cabling or configuration changes to these interfaces on Exalogic compute nodes is not permitted.

- **InfiniBand private network**: This required network connects the compute nodes and the storage appliance through the `BOND0` interface to the InfiniBand switches/ gateways on the Exalogic rack. It is the default IP over InfiniBand (IPoIB) subnet created automatically during the initial configuration of the Exalogic machine.

  > **Note:**
  >
  > This network is either based on the default InfiniBand partition or based on a partition allocated for the Exalogic machine. A single default partition is defined at the rack level. For more information, see Work with the Default Rack-Level InfiniBand Partition.

- **Client access network**: This required network connects the compute nodes to your existing client network through the `BOND1` interface and is used for client access to the compute nodes (this is related primarily to a physical Exalogic deployment). Each Exalogic compute node has a single default client access (edge network) to an external 10 Gb Ethernet network through a Sun Network QDR InfiniBand Gateway Switch.

  The logical network interface of each compute node for client access network connectivity is bonded. Bond1 consists of 2 vNICs (Ethernet over IB vNICs). Each vNIC is mapped to a separate Sun Network QDR InfiniBand Gateway Switch for high availability (HA) and each host EoIB vNIC is associated with a different HCA IB port (On Oracle Linux, vNIC0 -> ib0, vNIC1 -> ib1; on Oracle Solaris, vNIC0 -> ibp0, vNIC1 -> ibp1).

- **Additional networks (optional)**: Each Sun Network QDR InfiniBand Gateway Switch has eight 10 Gb Ethernet ports. The number of ports used in Exalogic deployment depends on your specific bandwidth requirements (how many 10 Gb ports can be shared per compute node) and on your specific LAN/VLAN connection requirements. A group of 16 compute nodes connects 2 Sun Network QDR InfiniBand Gateway Switches in an active-passive bond. Each compute node is connected to two separate Sun Network QDR InfiniBand Gateway Switches for HA.

Note that each compute node requires a bond for each external network (physical network or VLAN).

Figure 6-1 shows the network diagram for the Exalogic machine with Oracle Linux operating system.

Figure 6-1 Network Diagram for Exalogic Machine

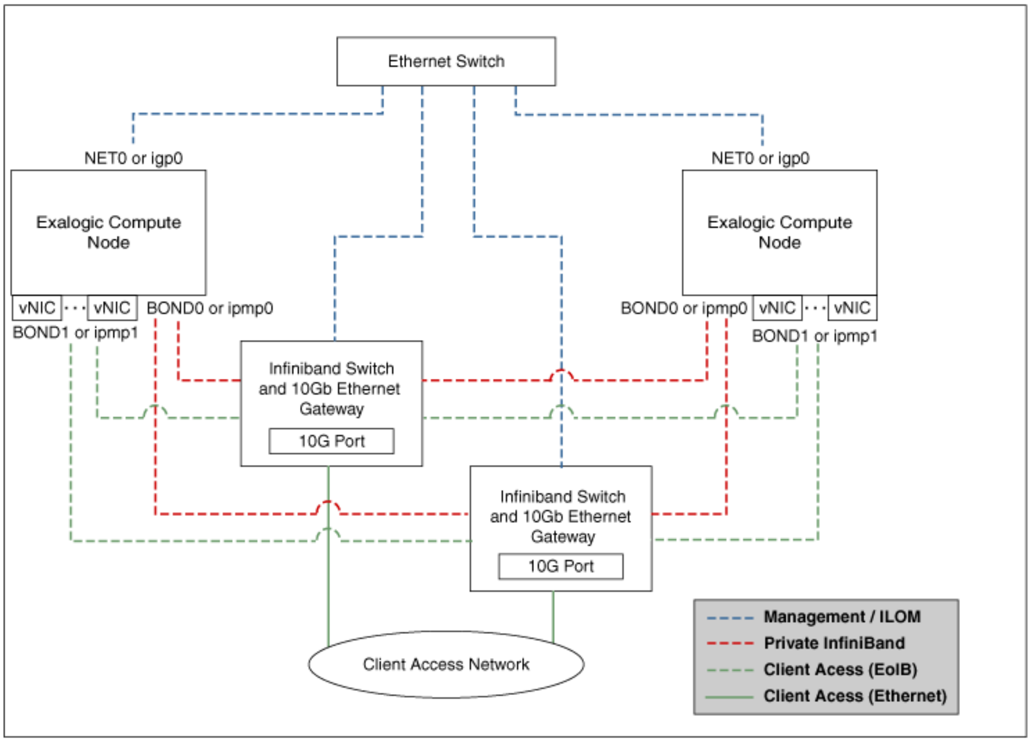

Note:

If you are using Oracle Solaris, you can assign the logical names of IPMP groups to be `ipmp0` or `BOND0`, and `ipmp1` or `BOND1` and have the name of the datalink corresponding to the `NET0` Ethernet port to be displayed as `igp0` or `net0` in the Solaris administration commands. For more information, see IPMP Overview for Oracle Solaris Users.

6.2 Naming Conventions

The Exalogic machine name is used to generate host names for network interfaces for all systems. For example, an Exalogic machine name of `el01` will result in compute node host names of `el01cn01`, `el01cn02`, `el01cn03`, and so on.

In a multirack configuration, each Exalogic rack name should be unique and identify the rack. Oracle recommends using `el01` for the first Exalogic rack, `el02` for the second, `el03` for the third, and so on.

Note:

You can connect up to eight Exalogic machine full racks, or a combination of 4 Exalogic machine full racks and 4 Oracle Exadata Database Machine full racks together on the same Infiniband fabric, without purchasing additional hardware. To connect more machines, Oracle offers a choice of several high-capacity datacenter switches which allow the creation of Exalogic clouds comprising hundreds of racks and tens of thousands of processors

6.3 Important Notes for Oracle Solaris Users

If you are using the Oracle Solaris operating system on Exalogic compute nodes, keep the following points in mind:

- `BOND0` and `BOND1`, two important terms used in this guide, refer to the default interfaces for IP over InfiniBand (IPoIB) and Ethernet over InfiniBand (EoIB), respectively, on the Oracle Linux operating system.
- Oracle Solaris uses the IP Multipathing (IPMP) technology to support **IPMP Groups** that consist of one or more physical interfaces on the same system that are configured with the same IPMP group name. This technology provides the same functionality as *Bonded Interfaces* on Oracle Linux. You can name the IPMP groups anything. In this guide, `BOND0` and `BOND1` are used as example names to keep the terminology consistent with Oracle Linux.

Note:

For an overview of IPMP, see IPMP Overview for Oracle Solaris Users.

6.4 Network Connection and IP Address Requirements

This section describes the network connections required for multiple networks of Exalogic machine. The network cables must be run from your network equipment to the location of the Exalogic machine.

The requirements to connect the Exalogic machine to your existing Ethernet network infrastructure are described in the following sections:

- Network Connection Requirements
- IP Address Requirements

6.4.1 Network Connection Requirements

Before installation, network cables must be run from your existing network infrastructure to the installation site. The requirements to connect the Exalogic machine to your existing network infrastructure are as follows:

- Management network connection requirements

  - At least one Ethernet connection for the Ethernet switch in the rack to the existing management network. The total number of connections depends on your HA requirements for the data center.

- Client access network connection requirements
  - At least one Ethernet over InfiniBand (EoIB) bond per Exalogic compute node to the client access network is required. Additional client access network interface bonds are added for each physical LAN or each virtual LAN (VLAN) in which the compute node is connected. The minimum configuration requires (single Bond to single LAN per compute node) the number of network connections per Exalogic compute node, as outlined in the subsequent table for Full Rack, Half Rack, and Quarter Rack configurations.

    The following table describes the number of required connections:

Type of Exalogic Machine	Bonded Network Configuration
Exalogic machine full rack	30
Exalogic machine half rack	16
Exalogic machine quarter rack	8
Exalogic machine eighth rack	4

  Up to eight external physical networks can be connected through EoIB using the Sun Network QDR InfiniBand Gateway Switch. You can create multiple bonded vNICs on the compute nodes.

- Additional network connection requirements

  You must determine the following:
  - How many LANs to connect to (typically, 1 or 2, maximum 8 physically isolated LANs)
  - How many 10 Gb ports per LAN are required (your network throughput requirements for client access network)

  At least two 10 Gb Ethernet ports per compute node - one active and another passive.

  The following are the connection choices based on the number of Ethernet ports:
  - 1 active 10 Gb Ethernet port shared by up to 8 compute nodes
  - 2 active 10 Gb Ethernet ports shared by up to 8 compute nodes
    * 1 port shared by 4 compute nodes for 1 LAN
    * 1 port shared by 8 compute nodes for 2 LANs (1 port per LAN per compute node)
  - 4 active 10 Gb Ethernet ports shared by up to 8 compute nodes
    * 1 port shared by 2 compute nodes for 1 LAN (1 port per LAN per compute node)
    * 1 port shared by 4 compute nodes for 2 LANs (1 port per LAN per compute node)
    * 1 port shared by 8 compute nodes for 4 LANs (1 port per LAN per compute node)

- 8 active 10 Gb Ethernet ports shared by up to 8 compute nodes
  * 1 port per 1 compute node for 1 LAN
  * 1 port shared by 2 compute nodes for 2 LANs (1 port per LAN per compute node)
  * 1 port shared by 4 compute nodes for 4 LANs (1 port per LAN per compute node)
  * 1 port shared by 8 compute nodes for 8 LANs (1 port per LAN per compute node)

6.4.2 IP Address Requirements

An Exalogic machine requires a large number of host names and IP addresses during initial configuration. The number of IP addresses required for a particular network, such as the management network, depends on the type of system. The network configuration, such as host names and IP addresses, used during installation is generated from information you supply to your Oracle technical representative in the completed configuration worksheet. See Network Configuration Worksheets.

Configure the new IP addresses in your existing networks only after you have completed the configuration worksheets, and received the installation Template from your Oracle representative. All IP addresses in the installation Template must be unassigned at the time of initial configuration. In addition, all IP addresses must be statically assigned IP addresses, not dynamically assigned (DHCP) addresses.

All InfiniBand physical IP addresses in `BOND0` must be in the same subnet, with a minimum subnet mask of 255.255.240.0 (or /20). The subnet mask chosen should be wide enough to accommodate possible future expansion of the Exalogic machine and InfiniBand network. Each device (compute nodes, switches, and storage appliance) in the Exalogic rack is assigned a separate Local Identifier (LID) by the IB Subnet Manager. Oracle Exalogic supports a single IP over IB (IPoIB) link over a logical Exalogic deployment over the IB subnet. Multiple IP subnets can be layered on top of the default IPoIB link.

Table 6-2 lists the IP address requirements for the Exalogic machine:

Table 6-2 IP Address Requirements for Exalogic Machine

Type of Exalogic Machine	Management Network Requirements	Client Access Network Requirements	InfiniBand Private Network (IPoIB) Requirements
full rack	Minimum of 68 IP addresses • At least 2 IP addresses per management network for compute nodes: 1 IP address for ILOM 1 IP address for ETH0 • 2 IP addresses per controller/ server heads in the storage appliance • 1 IP address per Sun Network QDR InfiniBand Gateway Switch	Minimum of one IP address per Bond. Additional IP addresses per Bond added as per Service or Protocol requirements. For example, HTTP may only expose a floating IP address to an external Load Balancer and not to the HTTP Client (only observes VIP of Load Balancer), while JMS may expose the virtual IP directly to the JMS Client. Typically, an application server supporting both HTTP and JMS will have a floating IP address for JMS (public) and a floating IP address for HTTP (private to the Load Balancer). In addition, T3 typically uses a minimum of two IP addresses. One of these IP addresses is private to the Load Balancer, and the other floating IP address is exposed to application server clients. An application server client initially accesses the application server through a well-known virtual IP address, such as `T3://floating_IP:Port`. This request is received by an external Load Balancer, which forwards the packet to a private IP address of the application server. The application server responds with a handle that consists of another floating IP address, such as `W_floatingIP:Port`. This request is forwarded through the Client. For any new requests, Client accesses the application Server through `T3:W_floatingIP:Port`.	Minimum of one IP address per bond per compute node. Exalogic machine full rack contains 30 compute nodes. One compute node can be a member of only one private network of an application domain.

Table 6-2 (Cont.) IP Address Requirements for Exalogic Machine

Type of Exalogic Machine	Management Network Requirements	Client Access Network Requirements	InfiniBand Private Network (IPoIB) Requirements
half rack	Minimum of 38 IP addresses • At least 2 IP addresses per management network for compute nodes: 1 IP address for ILOM 1 IP address for ETH0 • 2 IP addresses per controller/ server heads in the storage appliance • 1 IP address per Sun Network QDR InfiniBand Gateway Switch	Same as above.	Minimum of one IP address per bond per compute node. Exalogic machine half rack contains 16 compute nodes. One compute node can be a member of only one private network of an application domain.
quarter rack	Minimum of 22 IP addresses • At least 2 IP addresses per management network for compute nodes: 1 IP address for ILOM 1 IP address for ETH0 • 2 IP addresses for controllers/ server heads in the storage appliance • 1 IP address per Sun Network QDR InfiniBand Gateway Switch	Same as above.	Minimum of one IP address per bond per compute node. Exalogic machine quarter rack contains 8 compute nodes. One compute node can be a member of only one private network of an application domain.

Table 6-2 (Cont.) IP Address Requirements for Exalogic Machine

Type of Exalogic Machine	Management Network Requirements	Client Access Network Requirements	InfiniBand Private Network (IPoIB) Requirements
eighth rack	Minimum of 20 IP addresses • At least 2 IP addresses per management network for compute nodes: 1 IP address for ILOM 1 IP address for ETH0 • 2 IP addresses for controllers/ server heads in the storage appliance • 1 IP address per Sun Network QDR InfiniBand Gateway Switch	Same as above.	Minimum of one IP address per bond per compute node. Exalogic machine eighth rack contains 4 compute nodes. One compute node can be a member of only one private network of an application domain.

6.5 Default InfiniBand Bonded IP Addresses

Default InfiniBand bonded IP addresses are assigned to Exalogic compute nodes at the time of manufacturing. However, you can reconfigure the IP addresses by using the Oracle `Exalogic Configuration Utility`. The Oracle `Exalogic Configuration Utility` and its associated scripts reconfigure InfiniBand Bonded IP addresses based on the input you provide to the *Exalogic Configurator* spreadsheet.

For a list of default InfiniBand Bonded IP addresses, see Default IP Addresses and Ports.

6.6 Introduction to Oracle Exalogic Network Configuration

This section introduces the following topics:

- InfiniBand Fabric
- InfiniBand Switches
- Default Bonded Interfaces
- IPMP Overview for Oracle Solaris Users
- Connectivity Between Exalogic Compute Nodes
- Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch
- Additional InfiniBand Network Requirements and Specifications

6.6.1 InfiniBand Fabric

Exalogic machines use a unified 40 Gb per second InfiniBand quad data rate (QDR) fabric for internal communication.

Applications running on compute nodes communicate with applications on other compute nodes using this InfiniBand network. Exalogic machines communicate with Oracle Exadata Database Machines for database connectivity via IPoIB. Exalogic machines can be connected to an external network, including a standard database hosted on a machine outside of the Exalogic machine, via the InfiniBand-to-10 Gb Ethernet gateways using Ethernet over InfiniBand (EoIB). Each Exalogic machine configuration includes at least 2 such gateways, which also act as InfiniBand switches connecting all compute nodes and the storage appliance within the Exalogic machine.

6.6.2 InfiniBand Switches

Sun Network QDR InfiniBand Gateway Switches (part number NM2-GW) are used as the leaf switches in the Exalogic machine. They connect to the Host Channel Adapters (HCAs) of Exalogic compute nodes.

These switches (NM2-GW) also act as Ethernet gateways to connect your Exalogic machine to the external LAN over Ethernet. For more information, see Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch.

Sun Datacenter InfiniBand Switch 36 (part number NM2-36P) is used in multirack configuration scenarios (an Exalogic machine to another Exalogic machine, and an Exalogic machine to an Oracle Exadata Database Machine) only. This switch is not connected or used in an Exalogic machine single rack.

> **Note:**
>
> In the Exalogic machine, InfiniBand switches (both leaf and spine switches) are automatically configured to separate the IP over InfiniBand (IPoIB) traffic and the Ethernet over InfiniBand (EoIB) traffic.

6.6.3 Default Bonded Interfaces

After the Sun Network QDR InfiniBand Gateway Switches are connected to Exalogic compute nodes, the following bonded interfaces are configured:

- IP over InfiniBand (IPoIB) - bond0 link (`ib0`/`ib1` for Oracle Linux, and `ibp0`/`ibp1` for Oracle Solaris)

  `ib0` or `ibp0` represents the HCA port 0 of compute nodes, and `ib1` or `ibp1` represents the HCA port 1 of compute nodes.

> **Note:**
>
> Depending on your application deployment and isolation requirements, you can create additional bonded IP subnet interfaces over this default IPoIB link.
>
> For more information, see the "Application Isolation by Subnetting over IPoIB" topic in the *Oracle Exalogic Enterprise Deployment Guide*.

- Ethernet over InfiniBand (EoIB) - bond1 link, which uses two vNICs, such as vNIC0 and vNIC1 for ib0 and ib1 (vNIC0 and vNIC1 for `ibp0` and `ibp1` on Oracle Solaris), respectively.

> **Note:**
>
> Oracle Solaris uses the IP Multipathing (IPMP) technology to support IPMP Groups that provide the same functionality as bonded interfaces on Oracle Linux. If you are using Oracle Solaris on Exalogic compute nodes, you can name the IPMP groups anything. In this guide, `BOND0` and `BOND1` are used as example names to keep the terminology consistent with Oracle Linux.

6.6.4 IPMP Overview for Oracle Solaris Users

On the Oracle Solaris operating system, IP network multipathing (IPMP) provides physical interface failure detection and transparent network access failover for a system with multiple interfaces on the same IP link. IPMP also provides load spreading of packets for systems with multiple interfaces.

This section discusses the following topics:

- IPMP Components
- IPMP Groups

6.6.4.1 IPMP Components

IPMP comprises the following components:

- The `in.mpathd` daemon
- The `/etc/default/mpathd` configuration file
- `ifconfig` options for IPMP configuration

> **Note:**
>
> For information about the `in.mpathd` daemon and the `mpathd` configuration file, see the *in.mpathd (1M)* man page on the Oracle Solaris operating system installed on Exalogic compute nodes. For information about `ifconfig`, see the *ifconfig (1M)* man page.

6.6.4.2 IPMP Groups

An IP multipathing group, or IPMP group, consists of one or more physical interfaces on the same system that are configured with the same IPMP group name. All interfaces in the IPMP group must be connected to the same IP link. The same (non-null) character string IPMP group name identifies all interfaces in the group. You can place interfaces from NICs of different speeds within the same IPMP group, as long as the NICs are of the same type. IPMP groups on Oracle Solaris provide the same functionality as Bonded Interfaces on Oracle Linux in the Exalogic environment. For example, the default IPMP group `ipmp0` comprises two physical interfaces that are connected to the default IPoIB link for internal communication in your Exalogic machine. The other default IPMP group `ipmp1` comprises two virtual interfaces that are connected to the default EoIB link for external data center connectivity.

> **Note:**
>
> For information about administering and configuring IPMP groups on the Oracle Solaris operating system installed on Exalogic compute nodes, see Oracle Solaris 11.1 documentation.

6.6.5 Connectivity Between Exalogic Compute Nodes

Compute nodes in the Exalogic machine are connected to one another through dual-ported InfiniBand quad data rate (QDR) host channel adapters (HCAs). Each HCA has an IP address, and active-passive bonding is configured. The active port of the HCA connects to an Sun Network QDR InfiniBand Gateway Switch, and the passive port of the HCA connects to another Sun Network QDR InfiniBand Gateway Switch in the Exalogic machine.

> **Note:**
>
> For more information about network connectivity in different Exalogic machine configurations, see Cabling Diagrams .

6.6.6 Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch

The Sun Network QDR InfiniBand Gateway Switches also act as gateways to connect to Ethernet networks, and they support eight 10 GB Ethernet ports. These ports can be accessed by Exalogic compute nodes through the InfiniBand network through EoIB. You can create multiple VLANs per each of these Ethernet ports.

Each Exalogic compute node can access one or more Ethernet ports on two Sun Network QDR InfiniBand Gateway Switches (NM2-GW), for HA purposes. An Exalogic machine full rack includes 4 gateway switches. Therefore, a group of 8 compute nodes in the Exalogic machine full rack can access one Ethernet port on both the primary gateway switch and the secondary gateway switch that the group of compute nodes is connected to. Each port is represented as an EoIB vNIC at the compute nodes. Each compute node has two bonded vNICs (active/passive).

> **Note:**
>
> You can configure up to eight compute nodes to use a single 10 GB Ethernet port.
>
> For information about creating a VNIC for Ethernet connectivity, see Configure Ethernet Over InfiniBand.

This section discusses the following topics:

- Ethernet Device Requirements
- Network Interface Configuration for Compute Nodes
- Transceiver and Cable Requirements

6.6.6.1 Ethernet Device Requirements

Before you begin, ensure that you have a 10 Gb Ethernet Switch, Router, or NIC device that supports any of the following:

- SFP+ 10G-Base-SR Module
- XFP 10G-Base-SR Module
- QSFP Optical Module

For example, here is how a QSFP module on the Exalogic's Sun Network QDR InfiniBand Gateway Switch (NM2-GW) is connected to the SFP+/XFP modules on the data center's 10 GbE switch.

Figure 6-2 Connectivity Between NM2-GW and External 10 GB Ethernet Switch

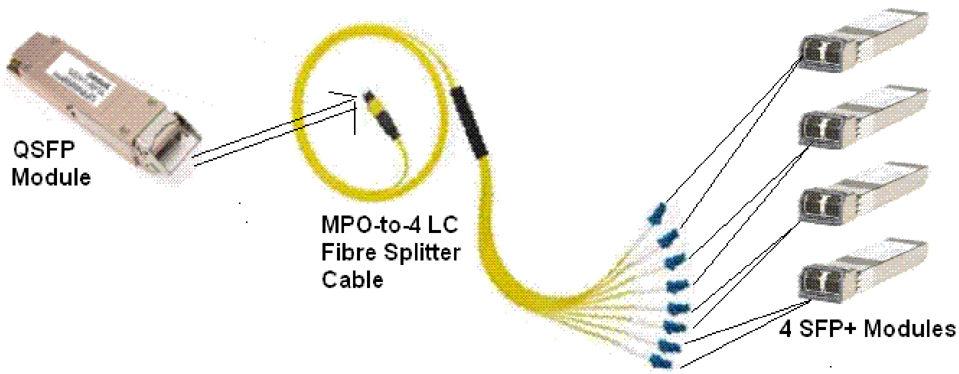

6.6.6.2 Network Interface Configuration for Compute Nodes

By default, each Exalogic compute node is configured with one bonded EoIB interface (ethX) for one external LAN. It is `BOND1` (vnic0/vnic1), which connects to one external LAN, such as LAN1.

If a vNIC is created at one of the Sun Network QDR InfiniBand Gateway Switches, the ethX interface is associated with the vNIC automatically.

> **Note:**
>
> You can configure additional EoIB network interfaces for connecting to additional LANs, as required.

6.6.6.3 Transceiver and Cable Requirements

Table 6-3 lists the transceiver and cable requirements that you must complete to connect your Exalogic machine to your data center's 10 Gb Ethernet switch.

Table 6-3 Transceivers and Cables

Optical Module on Exalogic's Sun Network QDR InfiniBand Gateway Switch	Cable Needed	Ethernet Switch Vendor	Transceiver Needed
QSFP module	QSFP MTP to 4 LC A minimum of one optical cable per NM2-GW is needed, but two cables per NM2-GW are recommended.	A Sun Oracle switch or a 10 GbE standard switch from a third-party vendor	For Sun Oracle switch: x2129/3 SFP+/XFP SR module For third-party switches: SFP+/XFP module provided by the switch vendor
QSFP module	QSFP – QSFP A minimum of one optical cable per NM2-GW is needed, but two cables per NM2-GW are recommended.	A Sun Oracle switch or a 10 GbE standard switch from a third-party vendor	For Sun Oracle switch: x2124A QSFP module For third-party switches: QSFP module provided by the switch vendor **Note**: Exalogic ships with QSFP transceivers, by default. Customers may use them on the data center switch side if they use a Sun Oracle 10GbE switch, such as the Sun Network 10 GbE Switch 72p.

6.6.7 Additional InfiniBand Network Requirements and Specifications

Table 6-4 lists additional InfiniBand specifications and cable requirements.

Table 6-4 HCA, Port Specifications and Cable Requirements

Component/Item	Exalogic Machine Full Rack	Exalogic Machine Half Rack	Exalogic Machine Quarter Rack	Two Exalogic Machines
InfiniBand quad data rate (QDR) host channel adapters (HCAs)	30	16	8	60
Unused ports in (Sun Network QDR InfiniBand Gateway Switches (NM2-GW leaf switches)	0	6	16	6
Unused ports in Sun Datacenter InfiniBand Switch (NM2-36P) spine switch **Note**: This switch is used in multirack configurations only.	Not applicable	Not applicable	Not applicable	Not applicable

6.7 Prepare to Reconfigure the Networking of Exalogic Machine

You should prepare to reconfigure the networking of your Exalogic machine as follows:

1. Identify your current management network.
2. Identify your current client access network.
3. Determine if you will connect multiple Exalogic machines or a combination of Exalogic machines and Oracle Exadata Database Machines. If you plan do to either of these, then contact your Oracle representative.
4. Provide input to the *Exalogic Configurator* spreadsheet, which will be used to generate configuration scripts required for the initial network configuration.
5. Review the values you entered in the spreadsheet.
6. Execute Oracle Exalogic Configuration Utility scripts, as described in Configure the Exalogic Machine Using ECU.
7. Run the network connections to the planned Exalogic machine location.
8. Inform your Oracle representative when you have completed these steps.

6.8 Subnet Manager Requirements for Connecting Exalogic to Exadata

See Subnet Manager Operation in Different Rack Configurations.

6.9 Network Configuration Worksheets

The following networks are used with the Exalogic machine:

- Management network
- Client access network
- InfiniBand private network (IPoIB)

The InfiniBand private network is a non-routable network fully contained in the Exalogic machine, and it does not connect to your existing network. This network is automatically configured during installation.

> **Note:**
>
> All networks must be on distinct and separate subnets from each other.
>
> All IP addresses must be statically assigned IP addresses, not dynamically assigned (DHCP) addresses.

This section contains the following worksheets:

- General Network Configuration Worksheet (Required)
- Management Network Configuration Worksheet (Required)
- Client Access Network Configuration Worksheet (Required)
- Private InfiniBand Network Configuration Worksheet (Required)

6.9.1 General Network Configuration Worksheet (Required)

Table 6-5 is the general network configuration worksheet. Fill out this worksheet and provide it to the Oracle representative, or use as a reference to complete the initial configuration of your Exalogic machine. Information entered in this worksheet is used as input to Oracle `Exalogic Configuration Utility`. This worksheet includes fields in the *Exalogic Configurator* spreadsheet that typically require data.

Table 6-5 General Network Configuration Worksheet

Item	Entry	Description and Example
Domain name		Company network domain name. Example: abc.example.com
Region		Name of the country in which the Exalogic machine resides.
Time Zone		Valid time zone.

Table 6-5 (Cont.) General Network Configuration Worksheet

Item	Entry	Description and Example
IP address of the Domain Name Server		IP address of one or more network name servers (up to four servers). Example: 10.25.45.123, 10.25.45.125
NTP Server		IP address of one or more Network Time Protocol servers (up to four servers). Example: 10.12.13.14, 10.12.13.15
Search Domains		A list of search domains for name lookup (up to four domains). Example: example.com, example.org
Default Gateway		IP address of the default gateway in your organization. Example: 10.203.72.2

6.9.2 Management Network Configuration Worksheet (Required)

The management network is used for administrative work for all components of the Exalogic machine. It connects the `NET0` (`igb0` on Oracle Solaris) network interface on all compute nodes, Integrated Lights Out Manager (ILOM), server heads of the storage appliance, and InfiniBand gateways/switches to the Cisco Ethernet switch in the rack.

The Cisco Ethernet switch supplied with the Exalogic machine is minimally configured during installation. The minimal configuration disables IP routing, and sets the following:

- First management IP address, subnet mask, and gateway for `NET0` (`igb0` for Oracle Solaris)
- First IP address, subnet mask, and gateway for ILOM

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

Table 6-6 is the management network configuration worksheet. Fill out this worksheet and provide it to the Oracle representative, or use as a reference to complete the initial configuration of your Exalogic machine. Information entered in this worksheet is used as input to Oracle `Exalogic Configuration Utility`. This worksheet includes fields in the *Exalogic Configurator* spreadsheet that typically require data.

Table 6-6 Management Network Configuration Worksheet

Item	Entry	Description and Example
First Management IP address or First ILOM IP address		First IP address in the sequential range used for management/network interfaces. This network must be distinct from all other networks on the Exalogic machine. Example: 10.204.74.100
Subnet mask		Subnet mask for the management network. Example: 255.255.248.0
Gateway IP address		Gateway IP address for the management network. Example: 10.204.72.1

6.9.3 Client Access Network Configuration Worksheet (Required)

The client access network is used for client access to the Exalogic compute nodes. Table 6-7 is the client access network configuration worksheet. Fill out this worksheet and provide it to the Oracle representative, or use as a reference to complete the initial configuration of your Exalogic machine. Information entered in this worksheet is used as input to Oracle `Exalogic Configuration Utility`. This worksheet includes fields in the *Exalogic Configurator* spreadsheet that typically require data.

Table 6-7 Client Access Network Configuration Worksheet

Item	Entry	Description and Example
First Client Access IP Address (`BOND1`)		First IP address in the sequential range used for client access network interfaces. This network must be distinct from all other networks on the Exalogic machine. For information, see IP Address Requirements.
Subnet mask		Subnet mask for client access network. Example: 255.255.252.0
Gateway IP address		Gateway IP address for the client access network. Example: 172.16.8.1

6.9.4 Private InfiniBand Network Configuration Worksheet (Required)

The private InfiniBand network is used for fabric consolidation of inter-processor communication, network and storage. It is optimized for cluster and storage traffic.

Table 6-8 is the client access network configuration worksheet. Fill out this worksheet and provide it to the Oracle representative, or use as a reference to complete the initial configuration of your Exalogic machine. Information entered in this worksheet is used as input to Oracle `Exalogic Configuration Utility`. This worksheet includes fields in the *Exalogic Configurator* spreadsheet that typically require data.

Table 6-8 Private InfiniBand Network Configuration Worksheet

Item	Entry	Description and Example
First Private InfiniBand IP Address (`BOND0`)		First IP address in the sequential range used for private InfiniBand network interfaces. This network must be distinct from all other networks on the Exalogic machine. For information, see IP Address Requirements.
Subnet mask		Subnet mask for the private InfiniBand network. Example: 255.255.192.0
Gateway IP address		Gateway IP address for private InfiniBand network. Example: 172.16.8.2

7 Set Up ILOM on the Compute Nodes

This chapter describes how to set up and access Oracle Integrated Lights Out Manager (ILOM) for Exalogic compute nodes. It contains the following topics:

7.1 ILOM Overview

Oracle Integrated Lights Out Manager (ILOM) provides advanced service processor (SP) hardware and software that you can use to manage and monitor your Exalogic machine components, such as compute nodes, gateway switches, storage appliance, and the InfiniBand switch. ILOM's dedicated hardware and software is preinstalled on these components.

ILOM enables you to actively manage and monitor compute nodes in the Exalogic machine independently of the operating system state, providing you with a reliable Lights Out Management (LOM) system.

With ILOM, you can proactively:

- Learn about hardware errors and faults as they occur
- Remotely control the power state of your compute node
- View the graphical and non-graphical consoles for the host
- 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 via IPMI PETs, SNMP Traps, or E-mail Alerts.

The ILOM service processor (SP) runs its own embedded operating system and has a dedicated Ethernet port, which together provide out-of-band management capability. In addition, you can access ILOM from the compute node's operating system. Using ILOM, you can remotely manage your compute node as if you were using a locally attached keyboard, monitor, and mouse.

ILOM automatically initializes as soon as power is applied to your compute node. It provides a full-featured, browser-based web interface and has an equivalent command-line interface (CLI).

The ILOM management interface is also integrated with Oracle Enterprise Manager Ops Center. Oracle Enterprise Manager Ops Center can discover new and existing systems on your network, update firmware and BIOS configurations, provision the operating environment with off-the-shelf distributions, manage updates and configuration changes, and remotely control key aspects of the service processor such as boot control, power status, and indicator lights.

Exalogic compute nodes are configured at the time of manufacturing to use Sideband Management. This configuration eliminates separate cables for the Service Processor (SP) NET MGT port and the NET0 Port.

7.1.1 ILOM Interfaces

ILOM supports following interfaces for accessing its features and functions:

- The *web interface* provides an easy-to-use browser interface that enables you to log in to the SP, then to perform system management and monitoring.
- The *command-line interface* enables you to operate ILOM using keyboard commands and adheres to industry-standard DMTF-style CLI and scripting protocols. ILOM supports SSH v2.0 and v3.0 for secure access to the CLI. Using the CLI, you can reuse existing scripts with Sun systems, and automate tasks using familiar interfaces.

7.2 Important Notes Before You Begin

You require at least one static IP address for Service Processor (ILOM access). For the list of default ILOM IP addresses assigned to Exalogic machine hardware components at the time of manufacturing, see Default IP Addresses and Ports.

If you reconfigure these IP addresses during the initial configuration of the Exalogic machine using Oracle `Exalogic Configuration Utility` and its associated tools, you must use those IP addresses to access ILOM.

Note that as a general convention, **root** is used as the default user name.

7.3 Management Network Diagram for Exalogic Machine

Figure 7-1 illustrates the management/ILOM network in the Exalogic machine.

Figure 7-1 Management Network in Exalogic Machine

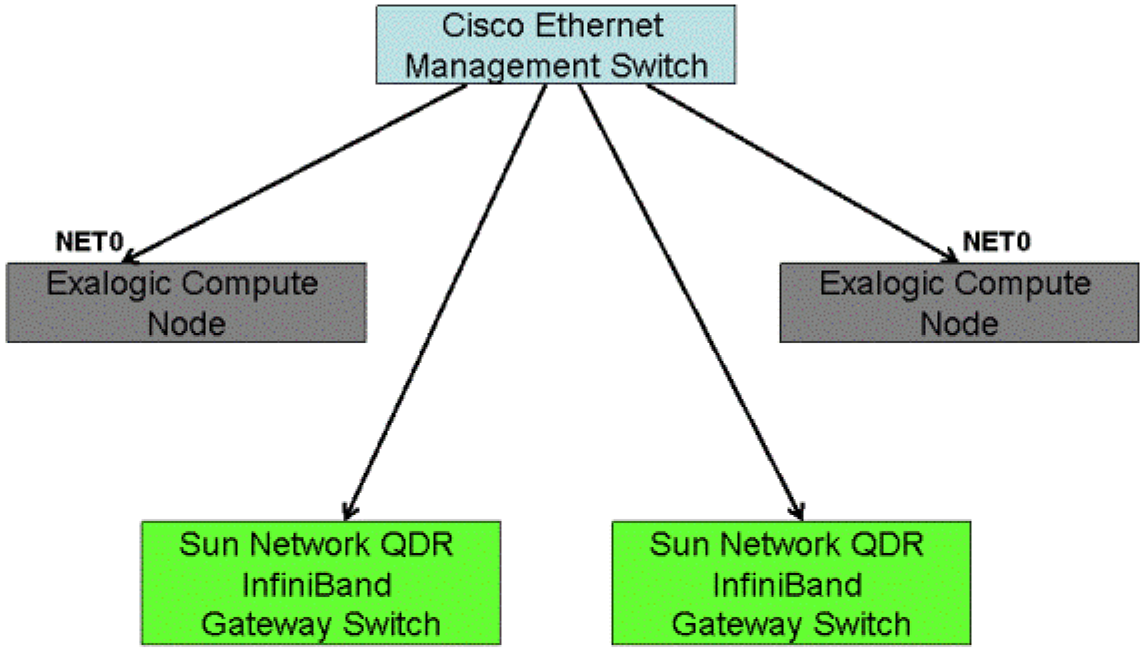

Figure 7-1 shows two compute nodes and two Sun Network QDR InfiniBand Gateway Switches as an example only. The number of compute nodes and gateway switches depends on your purchased Exalogic machine rack configuration. The Cisco Ethernet management switch is connected to the `NET0` port of compute nodes, and it is connected to the InfiniBand gateway switches.

7.4 ILOM IP Addresses for Exalogic Machine Components

ILOM IP addresses are assigned to Exalogic machine components at the time of manufacturing. For a list of default ILOM IP addresses, see Default IP Addresses and Ports.

However, you can reconfigure the IP addresses by using the Oracle `Exalogic Configuration Utility`. The Oracle `Exalogic Configuration Utility` and its associated scripts reconfigure ILOM IP addresses based on the input you provide to the *Exalogic Configurator* spreadsheet.

If the `Exalogic Configuration Utility` fails to reconfigure IP addresses or misconfigures IP addresses, you can configure IP addresses manually, as described in Configure ILOM IP Addresses Manually.

7.5 Connect to ILOM via the Network

Under normal circumstances, you can access ILOM via the network. You can use the Ethernet connection method to connect to the ILOM. You must know the ILOM's Ethernet address. For the default IP addresses assigned at the time of manufacturing, see Default IP Addresses and Ports. You can use either the CLI or the web interface to access ILOM. Alternatively, you can launch a remote KVM session to access ILOM.

> **Note:**
>
> You can use this connection method only if ILOM IP addresses are configured and can be accessed over the network. The network must be connected to the `NET0` port (with sideband) of each of the Exalogic compute nodes or hardware components.
>
> Oracle recommends that you use this connection method to access ILOM for Exalogic compute nodes.

This section discusses the following topics:

- Connect to the CLI
- Connect to the Web GUI
- Launch a Remote KVM Session

7.5.1 Connect to the CLI

1. Start your SSH client, such as PuTTY.
2. In the **Host Name (or IP address)** field, enter the ILOM IP address of the Exalogic compute node or hardware component you want to connect to.
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`.

   The CLI command prompt is displayed.

7.5.2 Connect to the Web GUI

1. In the address bar of a web browser, type the ILOM IP address of an Exalogic compute node or hardware component, and press **Enter**.

   The login screen is displayed.
2. Enter the user name and the password. The default user name is `root`.
3. Click **Log In**.

   The web GUI is displayed.

7.5.3 Launch a Remote KVM Session

To access the ILOM consoles for Exalogic machine components that are connected to ILOM, do the following:

1. Ensure that pop-up blockers are disabled in your browser before you launch the remote console.
2. Type the ILOM IP address, which you noted down in the procedure ILOM IP Addresses for Exalogic Machine Components, in the address bar of a web browser.

3. Log in to ILOM using the default user name `root`. After login, the ILOM home page is displayed.
4. Click the **Remote Control** tab, and click **Launch Remote Console**.

   The remote console is displayed.

7.6 Connect to ILOM via a Serial Connection

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

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

> **Note:**
>
> You can reconfigure network access after connecting to ILOM via a serial connection. For more information, see Reconfigure the Network Access.

7.6.1 Connect to the ILOM of a Compute Node

To connect to ILOM of a compute node in the Exalogic machine using a serial connection, complete the following steps:

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

   If there is power applied to either PSU, then ILOM will be functional regardless of the power state of compute nodes.
5. Press Enter on the terminal device. A connection between the terminal device and the ILOM is established.

   The ILOM login prompt is displayed.
6. Log in to the CLI using the default user name `root`.

   The ILOM default command prompt is displayed.

7.6.2 Connect to the ILOM of a Sun Network QDR InfiniBand Gateway Switch

To connect to ILOM of a Sun Network QDR InfiniBand Gateway Switch in the Exalogic machine using a serial connection, complete the following steps:

1. Attach a USB-to-Serial connector to the USB port of the gateway 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 as the `root` user. 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.

7.7 Reconfigure the Network Access

This step is required only if your network access must be reconfigured. You can reconfigure network access after connecting to the ILOM using a serial connection or an Ethernet connection.

> **Note:**
>
> Oracle recommends that you use an Ethernet connection if reconfiguration of network access becomes necessary.

This section contains the following topics:

- Reconfigure the Network Access Using a Serial Connection
- Reconfigure the Network Access Using the Ethernet Connection

7.7.1 Reconfigure the Network Access Using a Serial Connection

1. Connect to the ILOM using a serial connection, as described in Connect to ILOM via a Serial Connection.
2. Run the following commands at the ILOM command prompt:

```
set pendingipdiscovery=static
set pendingipaddress=<ip_address>
set pendingipnetmask=<ip_netmask>
set pendingipgateway=<ip_gateway>
set pendingmanagementport=/SYS/MB/NET0
set commitpending=true
```

7.7.2 Reconfigure the Network Access Using the Ethernet Connection

1. Connect to the ILOM using the Ethernet connection, as described in Connect to ILOM via the Network.

   > **Note:**
   >
   > An Ethernet cable is pre-wired in the Exalogic machine for this purpose.

2. Run the following commands at the ILOM command prompt:

```
set pendingipdiscovery=static
set pendingipaddress=<ip_address>
set pendingipnetmask=<ip_netmask>
set pendingipgateway=<ip_gateway>
set pendingmanagementport=/SYS/MB/NET0
set commitpending=true
```

7.7.3 Use the Ipmitool Commands when SP Network Information is Lost

If the service processor (SP) network information gets lost, you can use an SSH shell to connect to the ETH0 interface of the operating system on the compute node and run appropriate `ipmitool` commands to reconfigure network settings, as in the following example:

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

This example resets the ILOM root password.

> **Note:**
>
> For more information about `ipmitool` commands and options, navigate to the `/opt/ipmitool/bin/` directory and run `man ipmitool`.

7.7.4 Configure ILOM IP Addresses Manually

If ILOM IP addresses get misconfigured due to Oracle `Exalogic Configuration Utility` failures, you can configure ILOM IP addresses manually. Under normal circumstances, you do not need to configure ILOM IP addresses manually.

To configure IP addresses manually, complete the following steps:

1. Obtain a free static IP address from the network to which the Cisco 4948 switch is connected.
2. Configure the IP address and sideband management by using a serial port connection as follows:
   a. Establish a serial connection to the SP of a compute node by inserting a serial cable into the **SER MGT** port on the rear of the compute node.
   b. In the terminal window, configure the following settings:

      ```
      8N1: Eight data bit, no parity, one stop bit

      9600 Baud

      Disable hardware control

      Disable software control
      ```
   c. At the login prompt, log in as the `root` user.
   d. From your present working directory, use the `cd` command to move to the `/SP/network` directory, as shown in Figure 7-2.

Figure 7-2 Configuring ILOM IP Using Serial Port Connection

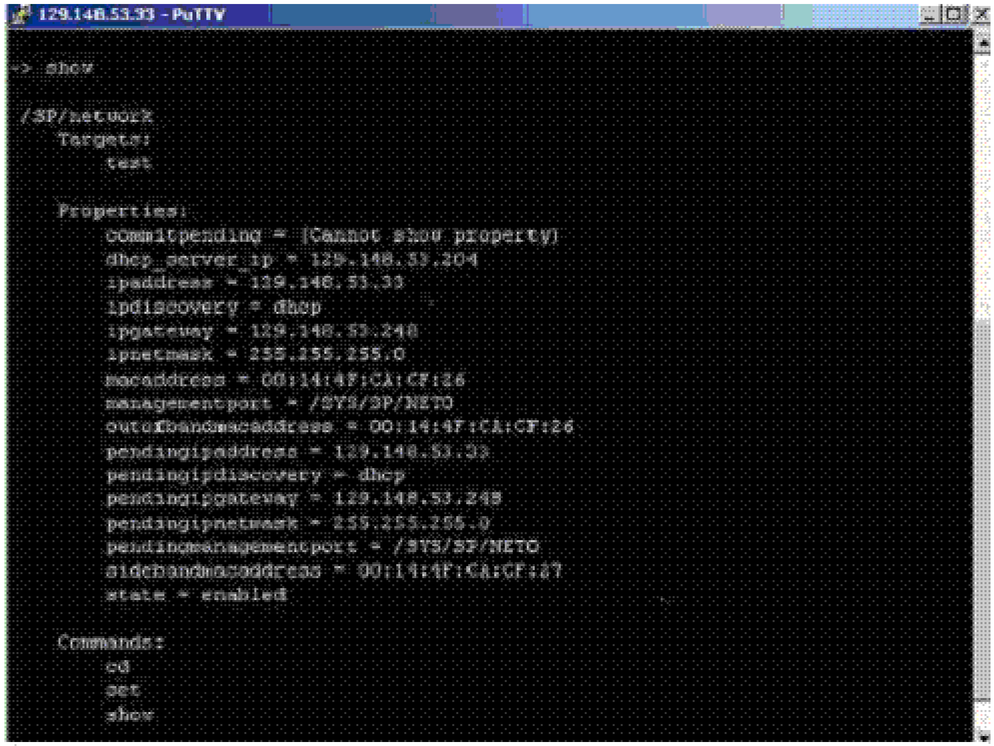

e. Run the following commands at the command prompt:

```
set pendingipdiscovery=static

set pendingipaddress=<ip_address>

set pendingipnetmask=<ip_netmask>

set pendingipgateway=<ip_gateway>

set pendingmanagementport=/SYS/MB/NET0

set commitpending=true
```

7.8 What Next?

After configuring ILOM for your Exalogic machine, verify the various configurations and proceed to complete the initial configuration of the storage appliance, which is included in the Exalogic machine.

8 Configure the Storage Appliance

This chapter describes how to configure the storage appliance, which is included in the Exalogic machine.
This chapter contains the following topics:

- Prerequisites
- Getting Started
- Storage Appliance Overview
- Configuration Overview
- Create Custom Projects
- Create Custom Shares
- Use the Phone Home Service to Manage the Storage Appliance

8.1 Prerequisites

The following are the prerequisites for configuring the storage appliance:

- Powering on the storage appliance by pressing the switches on the storage controllers, as described in Power On the Exalogic Machine
- Gathering information, such as IP address, IP net mask, Host name, Domain Name Server (DNS) domain name, DNS server IP address, Default router IP address, and Password for configuring an Ethernet interface on the storage controllers
- Running the Oracle `Exalogic Configuration Utility` to reconfigure IP addresses and other network parameters for the storage appliance

8.2 Getting Started

You can access the storage appliance over Ethernet via the Cisco Ethernet Management Switch.

The storage controllers are configured in an active-passive cluster, by default. The software propagates the configuration to the peer controller during cluster initialization. After the cluster is initialized, you can administer the system from either storage controller.

> **Tip:**
>
> Refer to the Cluster documentation in the *Oracle ZFS Storage Appliance Administration Guide* located on `http://docs.oracle.com/cd/E27998_01/html/E48433/index.html` for more information.

Complete the following steps:

1. Verify that the storage appliance is powered up and on the network.
2. Connect an Ethernet cable from your network to the `NET0` port on the back panel of the controller (storage server head).
3. Open a terminal window and use an SSH client to connect to the administrative console of the storage appliance (`ssh root@192.168.128.256`). When prompted, enter the administrative password for the storage appliance that you set when running Oracle `Exalogic Configuration Utility` to configure the Exalogic machine.
    a. After login, at the command prompt, type `start /SP/console`.
    b. Type `y` to confirm that you want to start the console.
    c. Press any key to begin configuring the appliance. The shell interface configuration screen appears. NET-0 at the top of the screen should be underlined.
    d. Verify the information on the screen, or enter values that do not appear.
    e. Apply the values by pressing `ESC-1` or the `F1` key or by pressing `Enter` after confirming the password. The final shell configuration screen appears, confirming that your appliance is ready for further configuration using the browser user interface (BUI).
4. Configure the remaining system parameters through a browser running on any client on the same network as the initial interface. The management software is designed to be fully featured and functional on the following supported web browsers: Firefox 2.x and 3.x, Internet Explorer 7, Internet Explorer 8, Safari 3.1 or later, and WebKit 525.13 or later.
5. Direct your browser to the storage system using either the IP address or host name you assigned to the NET0 port as follows:

    ```
    https://ipaddress:215
    ```

    or

    ```
    https://hostname:215
    ```

    The login screen appears.
6. Type `root` into the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen appears.
7. To begin configuring the system, click **Start** on the Welcome screen.You are guided through the Initial Configuration of the remaining network interfaces, DNS, time settings, directory service, and storage.

8.3 Storage Appliance Overview

This section introduces projects and shares.

8.3.1 Introduction to Projects

All file systems and LUNs are grouped into projects. A project defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space) can be accessed from a single point.By default, the appliance creates node-level projects based on the number of compute nodes in your Exalogic machine when a storage pool is first configured. For example, for a compute node with the host name `abc`, the default project `abc_1` is created. You can create all shares within this default project. However, Oracle recommends that you create additional projects for organizational purposes.

8.3.2 Introduction to Shares

Shares are file systems and LUNs that are exported over supported data protocols to clients of the appliance. File systems export a file-based hierarchy and can be accessed over NFS over IPoIB in the case of Exalogic machines. The project/share tuple is a unique identifier for a share within a pool. Multiple projects can contain shares with the same name, but a single project cannot contain shares with the same name. A single project can contain both file systems and LUNs, and they share the same namespace.

For a list of default shares created in the Exalogic machine, see Default Storage Configuration.

8.4 Configuration Overview

The storage appliance in the Exalogic machine is configured at different stages of the Exalogic machine setup and enterprise deployment.

The following are the configuration stages:

- Initial Configuration
- Connect Storage Heads to the Management Network and Accessing the Web Interface
- Cluster Network Configuration
- Network Configuration Options
- Default Storage Configuration
- Custom Configuration

8.4.1 Initial Configuration

The initial configuration involves networking configuration for the `NET0` interface, configuration of ILOM IP addresses, launch of service processor console, launch of several client network services, and the layout of the storage pool for standalone operation. When completed, the appliance in the Exalogic machine is ready for use, and it will have default shares configured for Exalogic compute nodes to access.

> **Note:**
>
> When you run the Oracle Exalogic Configuration Utility set of tools and scripts, the initial configuration for the storage appliance is completed.

For more information, see the Shares and Configuration sections of the *Oracle ZFS Storage Appliance Administration Guide* (`http://docs.oracle.com/cd/E27998_01/html/E48433/toc.html`). Alternatively, see the *Oracle Fusion Middleware Exalogic Enterprise Deployment Guide* for the recommended storage configuration in the Oracle Exalogic environment.

8.4.2 Connect Storage Heads to the Management Network and Accessing the Web Interface

Figure 8-1 shows the physical network connections for the storage appliance.

Figure 8-1 Network Ports on the Storage Appliance

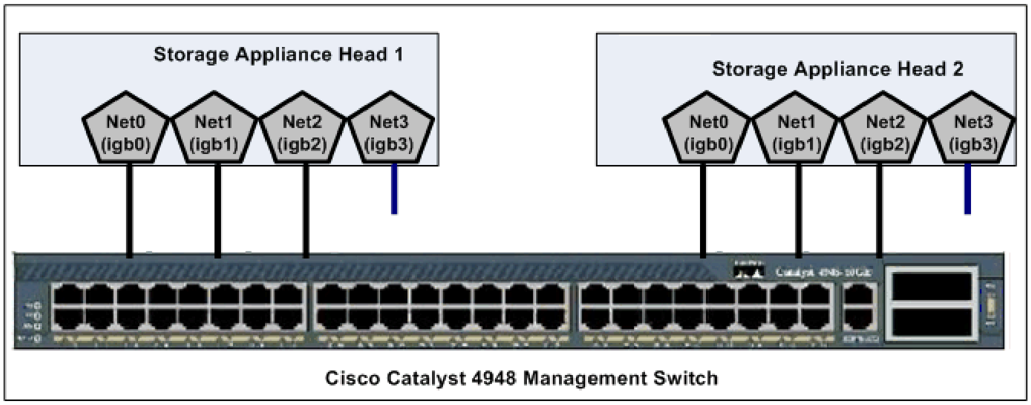

By default, the `NET0` (`igb0`), `NET1` (`igb1`), and `NET2` (`igb2`) ports on the storage heads are connected to the Cisco management switch, which is included in the Exalogic machine. The `igb0` and `igb1` interfaces are reserved for administrative access, such as access via a web browser or via command line. This configuration ensures that the storage heads are always reachable, independent of the load on the network data interfaces, and independent of which head is active. One end of a free hanging cable is connected to `NET3` (`igb3`). You can use the other end of this cable to connect to your data center network directly. Typically, for high availability purposes, this cable is connected to a data center switch other than the one that Exalogic's Cisco Management Switch is connected to.

To view the default network configuration of the storage appliance included in your Exalogic machine, do the following:

1. In a web browser, enter the IP address or host name you assigned to the `NET0` port of either storage head as follows:

   ```
   https://ipaddress:215
   ```

   or

   ```
   https://hostname:215
   ```

   The login screen appears.

2. Type `root` into the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen is displayed.

3. Click the **Configuration** tab, and click **NETWORK**. The default networking configuration is displayed, as shown in Figure 8-2.

> **Note:**
>
> The interface names and IP addresses shown on the screens in this chapter are examples only. You must verify the interface names in your environment and use them accordingly.

Figure 8-2 Network Configuration Screen

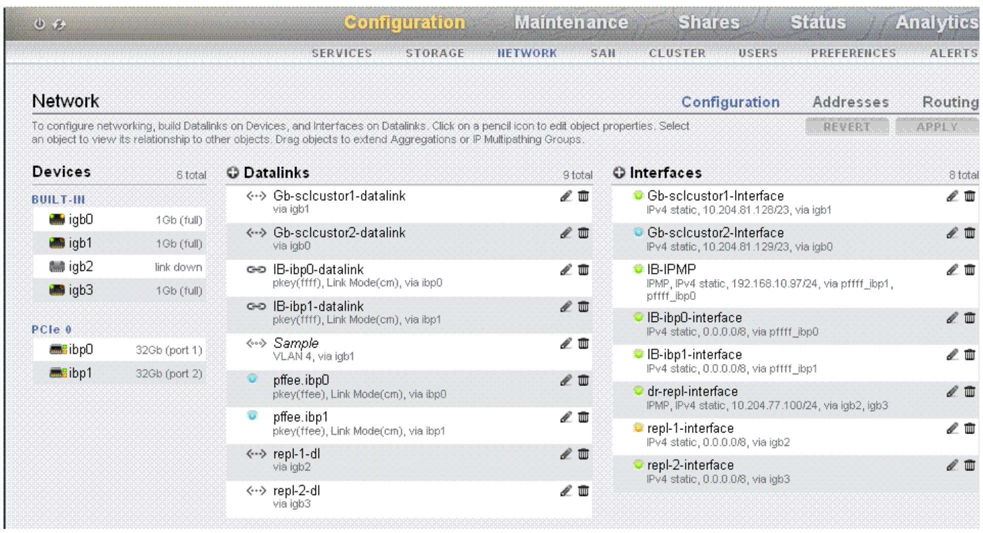

The Interfaces section shows the configured network interfaces. The green icon indicates that an interface is active on the storage head whose IP address or host name is used to access the administration console. The blue icon indicates that an interface is not active on the storage head. To view or edit the network settings for an interface, click the pencil icon. The interface settings are displayed in a screen, as in Figure 8-3.

Figure 8-3 Network Interface Settings

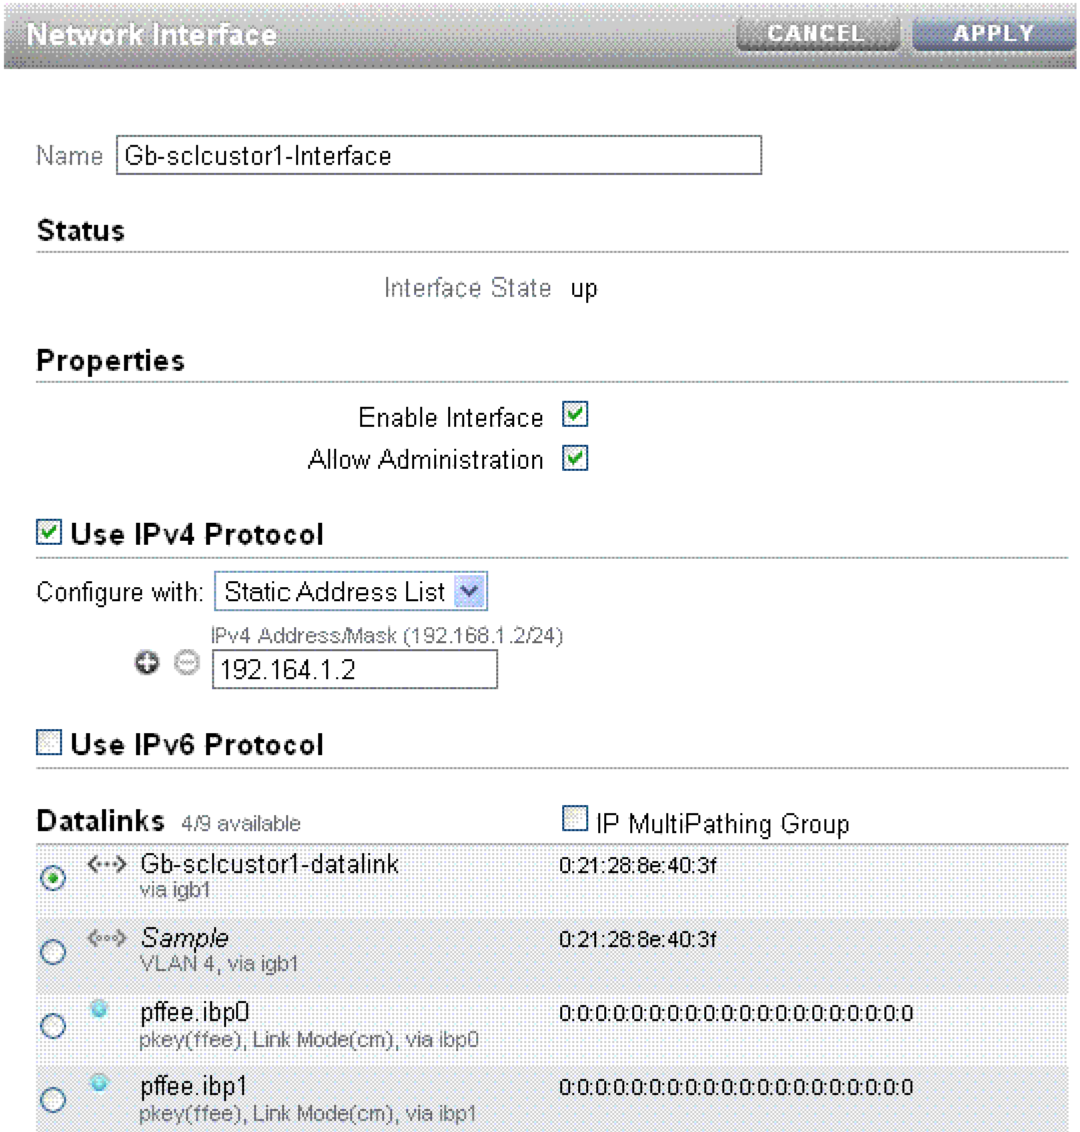

Note:

The interface names and IP addresses shown on the screens in this chapter are examples only. You must verify the interface names in your environment and use them accordingly.

8.4.3 Cluster Network Configuration

The cluster is set up in an active-passive configuration. All resources, data interface links, and storage pool are owned by the active storage head. When the active node fails, all resources (except the one that is locked to the active node) will be taken over by the passive storage head.

In the example configuration for an active head, `igb0` is used as the administrative network interface for the active storage head, such as `storagenode1`. The lock symbol indicates that `igb0` is locked to this storage head. To access this active storage head in a browser, you can use the following URL using either the host name or the IP address:

```
https://storagenode1:215
```

or

```
https://<IP_storagenode1>:215
```

In the example configuration for a passive head, `igb1` is used as the administrative network interface for the passive storage head, such as `storagenode2`. The lock symbol indicates that `igb1` is locked to this storage head. To access this passive storage head in a browser, you can use the following URL using either the host name or the IP address:

```
https://storagenode2:215
```

or

```
https://<IP_storagenode2>:215
```

> **Note:**
>
> For more information about network configuration for the storage appliance, see the *Network* topic in the *Oracle ZFS Storage Appliance Administration Guide*.

8.4.4 Network Configuration Options

You can choose any of the following network configuration options for the storage appliance, based on your specific requirements:

- Configure Option 1: ASR Support and Separate Paths for Management and Disaster Recovery
- Configure Option 2: ASR Support and Shared Path for Management and Disaster Recovery, with Single Management URL
- Configure Option 3: ASR Support and No Disaster Recovery, But with Single Management URL

8.4.4.1 Configure Option 1: ASR Support and Separate Paths for Management and Disaster Recovery

In this default configuration, the `igb0` port on your active storage head (head 1) is used, and the management option is enabled. The `igb0` port on your stand-by storage head (head 2) is not used. The `igb1` port on your stand-by storage head (head 2) is used, and the management option is disabled. The `igb2` and `igb3` ports are bonded with IP Multipathing (IPMP), and the management option is disabled on both `igb2` and `igb3`.

> **Tip:**
>
> Administrators should remember to use two different management URLs for the storage heads.

This default configuration option offers the following benefits:

- Supports Automated Service Request (ASR) for the storage appliance included in the Exalogic machine, using ports `igb0` and `igb1`
- Supports disaster recovery for the Exalogic machine, using ports `igb2` and `igb3`
- Provides Exalogic Configuration Utility, which is used to reconfigure the Exalogic machine based on your specific requirements, with ports `igb0` and `igb1`
- Separates the disaster recovery path from the management path

> **Note:**
>
> Ensure that the free hanging cable from the `igb3` port is connected to your data center network switch. Typically, for high availability purposes, this cable is connected to a data center switch other than the one that Exalogic's Cisco Management Switch is connected to.

The bonded interface is a new interface, such as `dr-repl-interface`, with `igb2` and `igb3` configured as an IPMP group. For example, the network settings of the `dr-repl-interface` is shown in Figure 8-4.

Figure 8-4   igb2 and igb3 in an IPMP Group

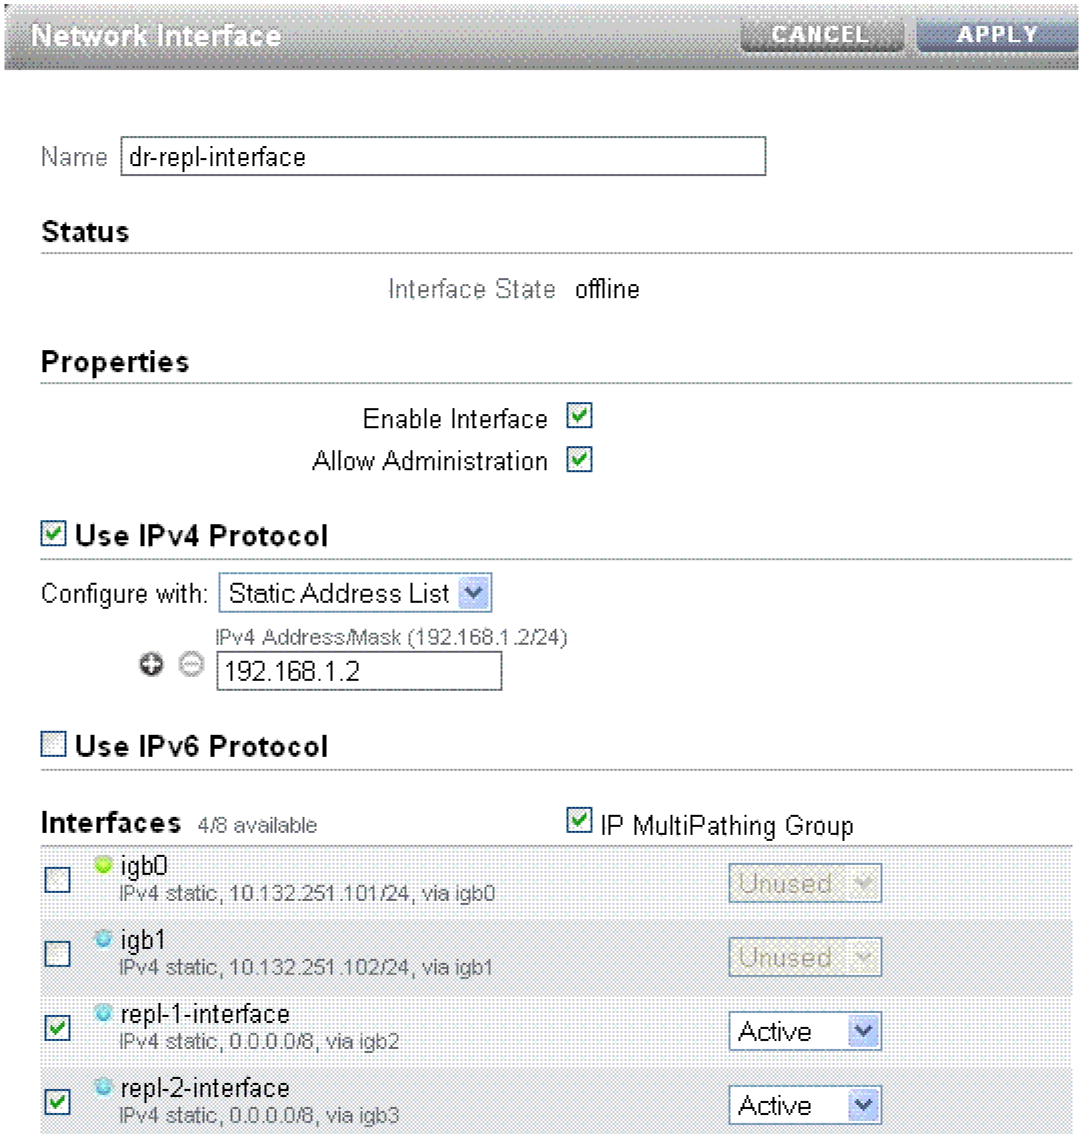

Note:

The interface names and IP addresses shown on the screens in this chapter are examples only. You must verify the interface names in your environment and use them accordingly.

In the Properties section, if you select the **Allow Administration** option, management is enabled on the interface. To create an IPMP Group with two interfaces, such as `igb2` and `igb3`, you must click the **+** icon (next to Interfaces) on the Figure 8-2. The Network Interface screen is displayed, as shown in Figure 8-5.

Figure 8-5 Creating a New IPMP Group Interface

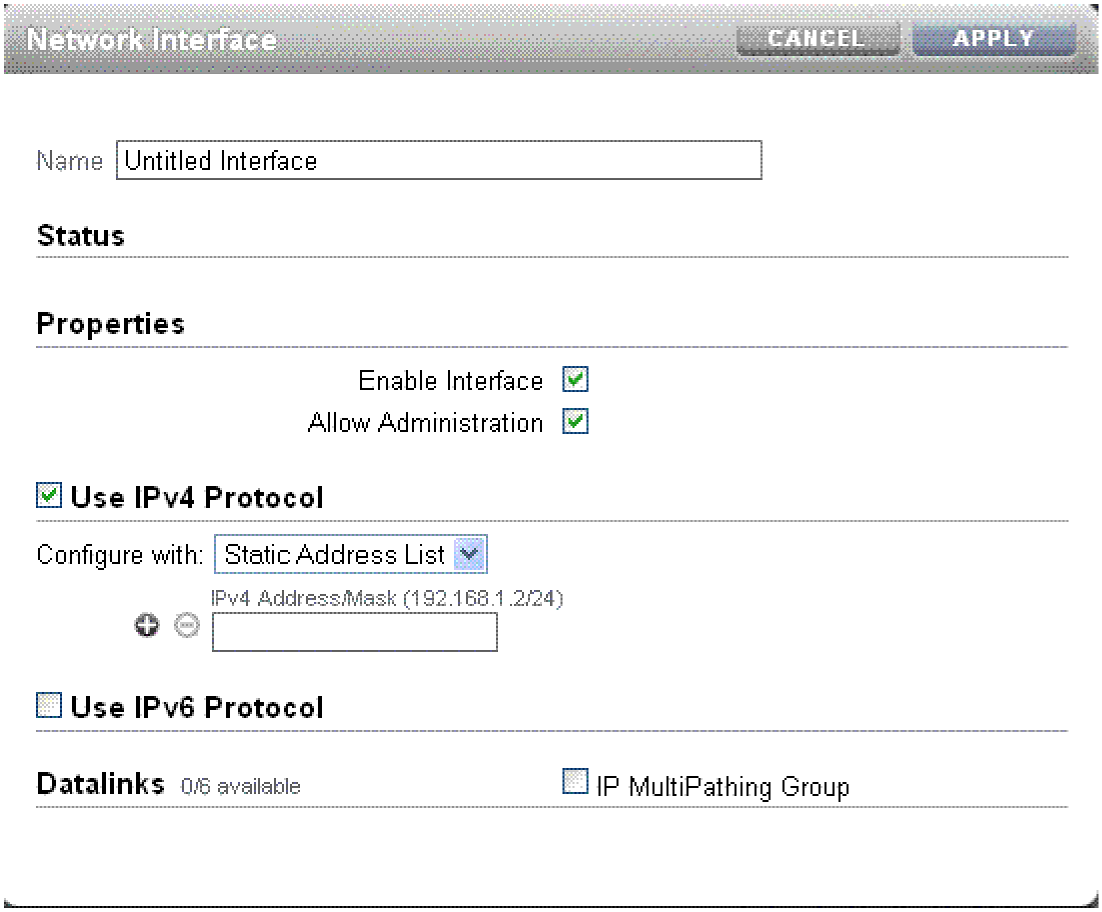

Enter a name for the new interface. In the Properties section, select the **Enable Interface** option. Select the **IP MultiPathing Group** option to configure two interfaces, such as `igb2` and `igb3`, in an IPMP group.

8.4.4.2 Configure Option 2: ASR Support and Shared Path for Management and Disaster Recovery, with Single Management URL

In this custom configuration, the `igb0` port on your active storage head (head 1) is used, and the management option is enabled. The `igb0` port on your stand-by storage head (head 2) is not used. The `igb1` port on your stand-by storage head (head 2) is used, and the management option is disabled. The `igb2` and `igb3` ports are bonded with IP Multipathing (IPMP), and the management option is enabled on both `igb2` and `igb3`.

This configuration option offers the following benefits:

- Supports Automated Service Request (ASR) for the storage appliance included in the Exalogic machine, using ports `igb0` and `igb1`
- Supports disaster recovery for the Exalogic machine, using ports `igb2` and `igb3`
- Provides Exalogic Configuration Utility, which is used to reconfigure the Exalogic machine based on your specific requirements, with ports `igb0` and `igb1`

- Provides single management URL for both storage heads, using ports `igb2` and `igb3`

> **Note:**
>
> This option does not separate the management path from the disaster recovery path.

To configure this option, complete the following steps:

1. Ensure that the physical connections are correct, as shown in Figure 8-1. Ensure that the free hanging cable from the `igb3` port is connected to your data center network switch.
2. In a web browser, enter the IP address or host name you assigned to the `NET0` port of either storage head as follows:

   `https://ipaddress:215`

   or

   `https://hostname:215`

   The login screen appears.
3. Type `root` into the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen is displayed.
4. Click the **Configuration** tab, and click **NETWORK**. The default networking configuration is displayed.
5. On the network configuration screen (Figure 8-2), click the pencil symbol next to the IPMP interface, such as `dr-repl-interface` (the bonded interface of `igb2` and `igb3`). The Network Interface screen for `dr-repl-interface` is displayed, as in Figure 8-6.

Figure 8-6 IPMP Network Interface Settings

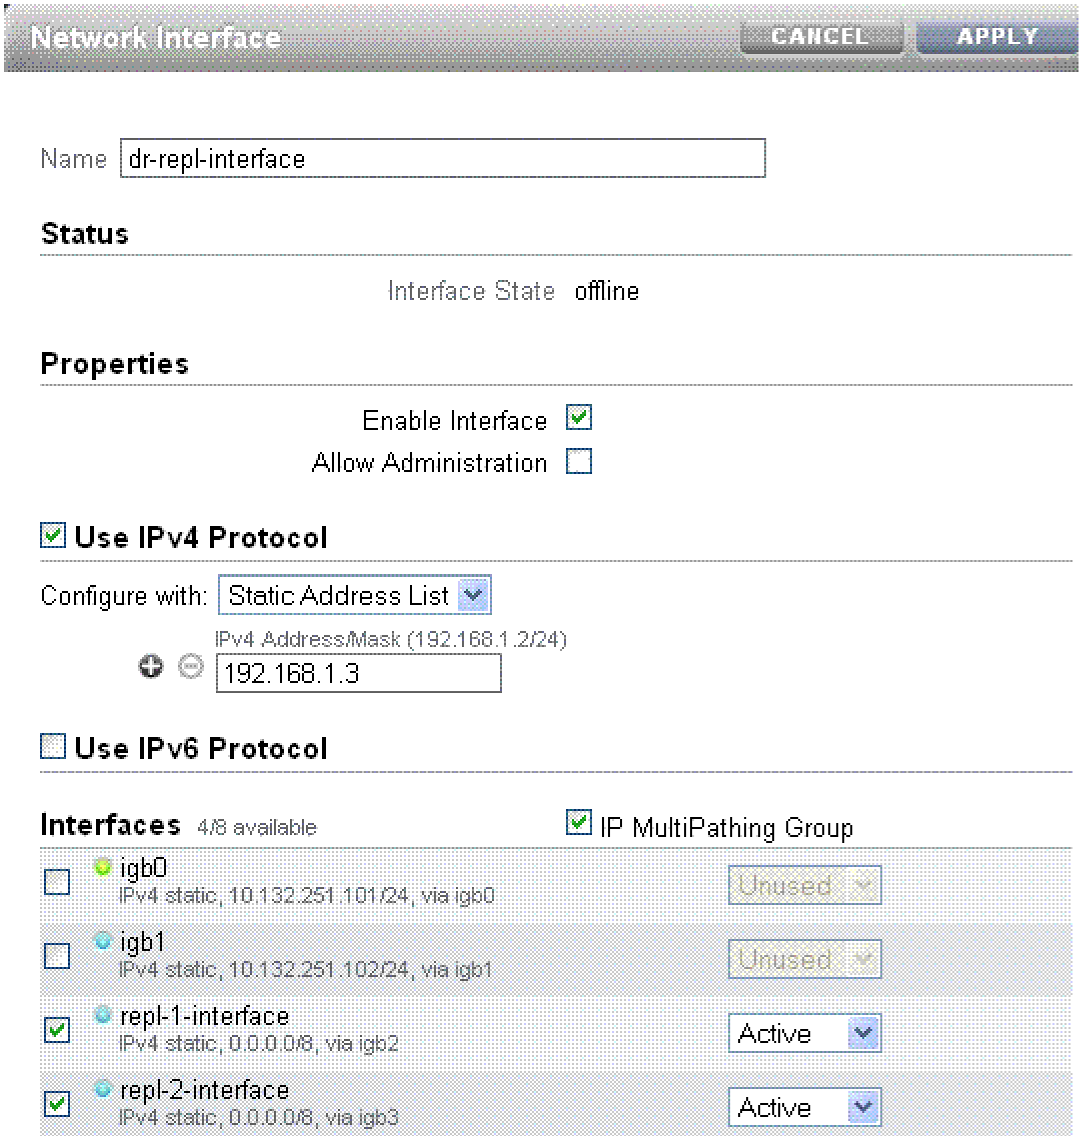

Note:

The interface names and IP addresses shown on the screens in this chapter are examples only. You must verify the interface names in your environment and use them accordingly.

6. Select the **Allow Administration** option to enable management traffic on both `igb2` and `igb3` interfaces.
7. Click **APPLY**.

8.4.4.3 Configure Option 3: ASR Support and No Disaster Recovery, But with Single Management URL

In this custom configuration, the `igb0` port on your active storage head (head 1) is used, and the management option is enabled. The `igb0` port on your stand-by storage head (head 2) is not used. The `igb1` port on your stand-by storage head (head 2) is used, and the management option is disabled. The `igb2` port uses a virtual IP, and the management option is enabled. The `igb3` port is not used.

This configuration option offers the following benefits:

- Supports Automated Service Request (ASR) for the storage appliance included in the Exalogic machine, using ports `igb0` and `igb1`
- Provides Exalogic Configuration Utility, which is used to reconfigure the Exalogic machine based on your specific requirements, with ports `igb0` and `igb1`
- Provides single management URL for both storage heads, using the port `igb2`

> **Note:**
>
> This option does not offer disaster recovery support. When you use this configuration option, you may connect the free hanging cable from `igb3` to the Cisco Management switch.

To configure this option, complete the following steps:

1. Ensure that the physical connections are correct, as shown in Figure 8-1.
2. In a web browser, enter the IP address or host name you assigned to the `NET0` port of either storage head as follows:

   ```
   https://ipaddress:215
   ```

   or

   ```
   https://hostname:215
   ```

   The login screen appears.
3. Type `root` into the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen is displayed.
4. Click the **Configuration** tab, and click **NETWORK**. The default networking configuration is displayed.
5. On the network configuration screen (Figure 8-2), click the delete symbol next to the IPMP interface, such as `dr-repl-interface` (the bonded interface of `igb2` and `igb3`). Delete this IPMP interface.
6. On the network configuration screen (Figure 8-2), click the pencil symbol next to the `igb3` interface. The Network Interface screen for `igb3` is displayed. Click the `Enable Interface` option to disable the interface, which is enabled, by default.
7. Click **APPLY**.

8.4.5 Default Storage Configuration

By default, a single storage pool is configured. Active-passive clustering for the server heads is configured. Data is mirrored, which yields a highly reliable and high-performing system.

The default storage configuration is done at the time of manufacturing, and it includes the following shares:

- Two exclusive NFS shares for each of the Exalogic compute nodes - one for crash dumps, and another for general purposes

  In this scenario, you can implement access control for these shares, based on your requirements.
- Two common NFS shares to be accessed by all compute nodes - one for patches, and another for general purposes

Table 8-1 Default Configuration of the storage appliance

Default Configuration	Name
Storage pool	`exalogic`

Projects	• Projects at the compute node level
	`NODE_1` to `NODE_N`
	where `N` represents the number of compute nodes in your Exalogic machine rack configuration.
	• Common project
	`common`

Shares	• NODE_SHARES
	`dumps`, and `general`
	These shares are at the compute node level.
	• COMMON_SHARES
	`common/patches`, `common/general`, and `common/images`

> **Note:**
>
> This table represents the default configuration of the storage appliance before the Exalogic machine rack configuration is modified at the customer's site. Oracle Exalogic Configuration Utility does not alter this configuration.

8.4.6 Custom Configuration

You can create and configure a number of projects and shares on the storage appliance to meet your specific storage requirements in the enterprise.

You can implement custom configuration, such as the following:

- Custom projects, such as `Dept_1`, `Dept_2`.
- Custom shares, such as `jmslogs`, `jtalogs`.

- Creation and administration of users.
- Access control for custom shares.

> **Note:**
>
> For information about the recommended directory structure and shares, see the *Oracle Fusion Middleware Exalogic Enterprise Deployment Guide*.

8.5 Create Custom Projects

Shares are grouped together as Projects. For example, you can create a project for `Dept_1`. `Dept_1` will contain department-level shares.

To create the `Dept_1` project, do the following:

1. In the Browser User Interface (BUI), click the **Shares** tab.

   The shares page is displayed.

2. Click the **Projects** panel.
3. Click the **+** button above the list of projects in the project panel.
4. Enter a name for the project, such as `Dept_1`. The new project `Dept_1` is listed on the Project Panel, which is on the left navigation pane.
5. Click the **General** tab on the `Dept_1` project page to set project properties. This section of the BUI controls overall settings for the project that are independent of any particular protocol and are not related to access control or snapshots. While the CLI groups all properties in a single list, this section describes the behavior of the properties in both contexts.

   The project settings page contains three sections: Space Usage (Users and Groups), Inherited Properties, and Default Settings (File systems and LUNs). Table 8-2 describes the project settings.

Table 8-2 Project Settings

Section and Setting	Description
Space Usage	Space within a storage pool is shared between all shares. File systems can grow or shrink dynamically as needed, though it is also possible to enforce space restrictions on a per-share basis. • **Quota** - Sets a maximum limit on the total amount of space consumed by all file systems and LUNs within the project. • **Reservation** - Guarantees a minimum amount of space for use across all file systems and LUNs within the project.

Table 8-2 (Cont.) Project Settings

Section and Setting	Description
Inherited Properties	Standard properties that can either be inherited by shares within the project. The behavior of these properties is identical to that at the shares level. • **Mountpoint** - The location where the file system is mounted. This property is only valid for file systems. Oracle recommends that you use specify `/export/<project_name>` as the default mountpoint. By using this consistently, you can group all shares and mount under the relevant project. It also prevents multiple shares from using the same mount points. Note that the same storage appliance is used by a multiple departments (15 in the case of Exalogic machine full rack configuration). The departments will have a similar share structure, such as `/export/dept_1/<share1>`, `/export/dept_2/share1`, and so on. • **Read only** - Controls whether the file system contents are read only. This property is only valid for file systems. • **Update access time on read** - Controls whether the access time for files is updated on read. This property is only valid for file systems. • **Non-blocking mandatory locking** - Controls whether CIFS locking semantics are enforced over POSIX semantics. This property is only valid for file systems. • **Data deduplication** - Controls whether duplicate copies of data are eliminated. • **Data compression** - Controls whether data is compressed before being written to disk. • **Checksum** - Controls the checksum used for data blocks. • **Cache device usage** - Controls whether cache devices are used for the share. • **Synchronous write bias** - Controls the behavior when servicing synchronous writes. By default, the system optimizes synchronous writes for latency, which leverages the log devices to provide fast response times. • **Database record size** - Controls the block size used by the file system. This property is only valid for file systems. By default, file systems will use a block size just large enough to hold the file, or 128K for large files. This means that any file over 128K in size will be using 128K blocks. If an application then writes to the file in small chunks, it will necessitate reading and writing out an entire 128K block, even if the amount of data being written is comparatively small. The property can be set to any power of 2 from 512 to 128K. • **Additional replication** - Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. • **Virus scan** - Controls whether this file system is scanned for viruses. This property is only valid for file systems. • **Prevent destruction** - When set, the share or project cannot be destroyed. This includes destroying a share through dependent clones, destroying a share within a project, or destroying a replication package. • **Restrict ownership change** - By default, this check box is selected and the ownership of files can only be changed by a root user. This property can be removed on a per-filesystem or per-project basis by deselecting this check box. When deselected, file ownership can be changed by the owner of the file or directory.

Table 8-2 (Cont.) Project Settings

Section and Setting	Description
Default Settings	Custom settings for file systems, to be used as default, include the following: • **User** - User that is the current owner of the directory. • **Group** - Group that is the current owner of the directory. • **Permissions** - Permissions include Read (R), Write (W), or Execute (X). Custom settings for LUNs, to be used as default, include the following: • **Volume Size** - Controls the size of the LUN. By default, LUNs reserve enough space to completely fill the volume • **Thin provisioned** - Controls whether space is reserved for the volume. This property is only valid for LUNs. By default, a LUN reserves exactly enough space to completely fill the volume. This ensures that clients will not get out-of-space errors at inopportune times. This property allows the volume size to exceed the amount of available space. When set, the LUN will consume only the space that has been written to the LUN. While this allows for thin provisioning of LUNs, most file systems do not expect to get "out of space" from underlying devices, and if the share runs out of space, it may cause instability or a corruption on clients, or both. • **Volume block size** - The native block size for LUNs. This can be any power of 2 from 512 bytes to 128K, and the default is 8K.

6. After entering your choices, click **Apply**.

8.6 Create Custom Shares

Shares are file systems and LUNs that are exported over supported data protocols to compute nodes. File systems export a file-based hierarchy and can be accessed over NFS over IPoIB in Exalogic machines.

To create a custom share, such as `domain_home` under the `Dept_1` project, do the following:

1. In the Browser User Interface (BUI), click the **Shares** tab.

   The shares page is displayed.

2. Click the **+** button next to **Filesystems** to add a file system. The Create Filesystem screen is displayed.

Figure 8-7 Create Filesystem

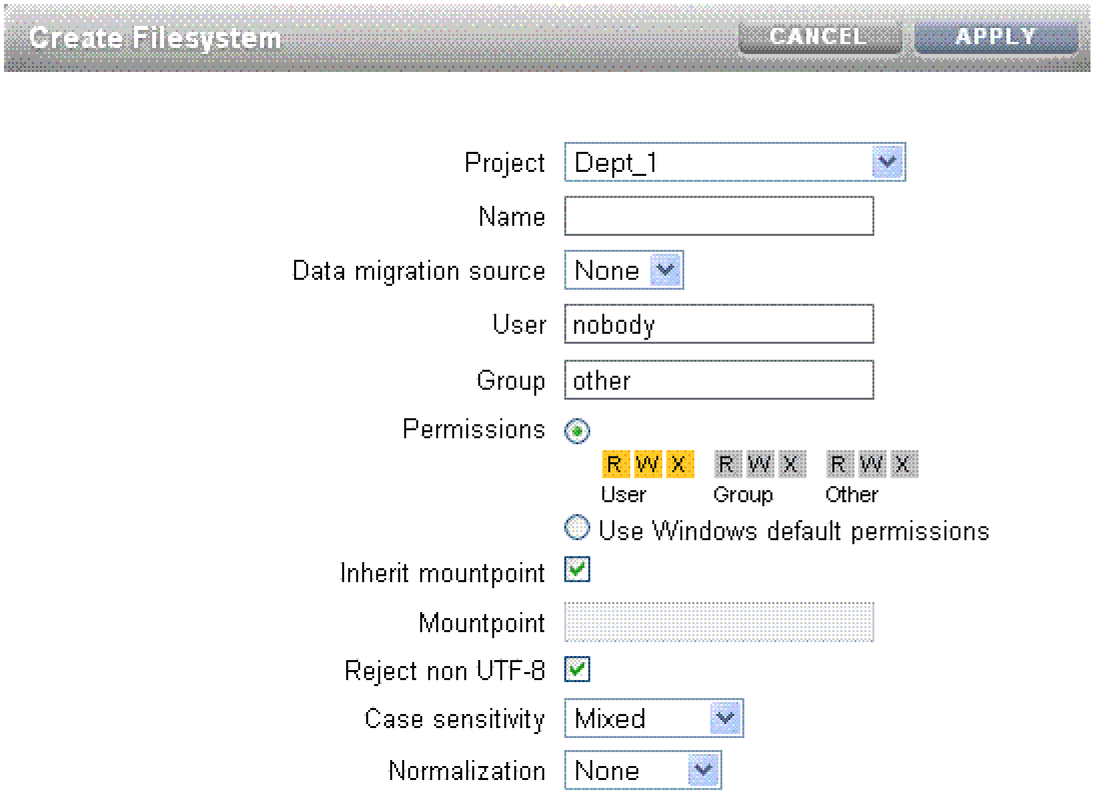

3. In the Create Filesystems screen, choose the target project from the **Project** pull-down menu. For example, choose `Dept_1`.
4. In the **Name** field, enter a name for the share. For example, enter `domain_home`.
5. From the **Data migration source** pull-down menu, choose **None**.
6. Select the **Permissions** option. Table 8-3 lists the access types and permissions.

Table 8-3 File System Access Types and Permissions

Access Type	Description	Permissions to Grant
User	User that is the current owner of the directory.	The following permissions can be granted: • **R** - Read - Permission to list the contents of the directory. • **W** - Write - Permission to create files in the directory. • **X** - Execute - Permission to look up entries in the directory. If users have execute permissions but not read permissions, they can access files explicitly by name but not list the contents of the directory.
Group	Group that is the current group of the directory.	
Other	All other accesses.	

You can use this feature to control access to the file system, based on the access types (users and groups) in `Dept_1`.

7. You can either inherit a mountpoint by selecting the **Inherit mountpoint** option or set a mountpoint.

> **Note:**
>
> The mount point must be under `/export`. The mount point for one share cannot conflict with another share. In addition, it cannot conflict with another share on cluster peer to allow for proper failover.
>
> When inheriting the mountpoint property, the current dataset name is appended to the project's mountpoint setting, joined with a slash ('/'). For example, if the `domain_home` project has the mountpoint setting `/export/domain_home`, then `domain_home/config` inherits the mountpoint `/export/domain_home/config`.

8. To enforce UTF-8 encoding for all files and directories in the file system, select the **Reject non UTF-8** option. When set, any attempts to create a file or directory with an invalid UTF-8 encoding will fail.

> **Note:**
>
> This option is selected only when you are creating the file system.

9. From the **Case sensitivity** pull-down menu, select **Mixed**, **Insensitive**, or **Sensitive** to control whether directory lookups are case-sensitive or case-insensitive.

Table 8-4 Case Sensitivity Values

BUI Value	Description
Mixed	Case sensitivity depends on the protocol being used. For NFS, FTP, and HTTP, lookups are case-sensitive. This is default, and prioritizes conformance of the various protocols over cross-protocol consistency.
Insensitive	All lookups are case-insensitive, even over protocols (such as NFS) that are traditionally case-sensitive. This setting should only be used where CIFS is the primary protocol and alternative protocols are considered second-class, where conformance to expected standards is not an issue.
Sensitive	All lookups are case-sensitive. In general, do not use this setting.

> **Note:**
>
> This option is selected only when you are creating the file system.

10. From the **Normalization** pull-down menu, select **None**, **Form C**, **Form D**, **Form KC**, or **Form KD** to control what unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encodings. Without normalization, the on-disk name stored will be different, and lookups using one of the alternative forms will fail depending on how the file was created and how it is accessed. If this property is set to anything other than **None** (the default), the **Reject non UTF-8** property must also be selected.

Table 8-5 Normalization Settings

BUI Value	Description
None	No normalization is done.
Form C	Normalization Form Canonical Composition (NFC) - Characters are decomposed and then recomposed by canonical equivalence.
Form D	Normalization Form Canonical Decomposition (NFD) - Characters are decomposed by canonical equivalence.
Form KC	Normalization Form Compatibility Composition (NFKC) - Characters are decomposed by compatibility equivalence, then recomposed by canonical equivalence.
Form KD	Normalization Form Compatibility Decomposition (NFKD) - Characters are decomposed by compatibility equivalence.

> **Note:**
>
> This option is selected only when you are creating the file system.

11. After entering the values, click **Apply**.

8.7 Use the Phone Home Service to Manage the Storage Appliance

You can use the PhoneHome service screen in the BUI to manage the appliance registration as well as the PhoneHome remote support service. Registering the storage appliance connects your appliance with the inventory portal of Oracle, through which you can manage your Sun gear. Registration is also a prerequisite for using the PhoneHome service.

The PhoneHome service communicates with Oracle support to provide:

- Fault reporting - the system reports active problems to Oracle for automated service response. Depending on the nature of the fault, a support case may be opened. Details of these events can be viewed in Problems.
- Heartbeats - daily heartbeat messages are sent to Oracle to indicate that the system is up and running. Oracle support may notify the technical contact for an account when one of the activated systems fails to send a heartbeat for too long.
- System configuration - periodic messages are sent to Oracle describing current software and hardware versions and configuration as well as storage configuration. No user data or metadata is transmitted in these messages.

> **Note:**
>
> You need a valid Oracle Single Sign-On account user name and password to use the fault reporting and heartbeat features of the Phone Home service. Go to `http://support.oracle.com` and click Register to create your account

8.7.1 Register Your Storage Appliance

To register the appliance for the first time, you must provide a Oracle Single Sign-On account and specify one of that account's inventory teams into which to register the appliance.

Using the BUI:

1. Enter your Oracle Single Sign-On user name and password. A privacy statement will be displayed for your review. It can be viewed at any time later in both the BUI and CLI.
2. The appliance will validate the credentials and allow you to choose which of your inventory teams to register with. The default team for each account is the same as the account user name, prefixed with a '$'.
3. Commit your changes.

> **Note:**
>
> You can see a log of PhoneHome events in **Maintenance**->**Logs-**>**PhoneHome**.
>
> If the phone home service is enabled before a valid Oracle Single Sign-On account has been entered, it will appear in the maintenance state. You must enter a valid Oracle Single Sign-On account to use the phone home service.

9
Configure NFS Version 4 on Exalogic

This chapter describes how to set up and configure Network File System (NFS) Version 4 (NFSv4) on Oracle Exalogic.
It contains the following topics:

- Overview
- Verify the NIS Setting on Exalogic
- Configure the Storage Appliance
- Configure an Exalogic Linux Compute Node to Use NFSv4
- Create NFSv4 Mount Points on Oracle Linux

9.1 Overview

Oracle Exalogic is a multihomed system where multiple IP addresses are configured on the InfiniBand network interface. NFSv4 addresses file locking problems on multihomed systems, such as Exalogic. The file locking problem occurs when NFSv3 is used.

9.1.1 NFSv3

NFSv3 is a stateless protocol. NFS clients do not maintain the state between requests; the clients rely on the Network Lock Manager (NLM) protocol to support file locking. There are known limitations on the NLM based locking, and NFSv3 it can cause lock ownership issues under certain multihomed system configurations, where the source IP addresses are not deterministic between lock and unlock requests.

9.1.2 NFSv4

Unlike NFSv3, NFSv4 is a stateful protocol. An NFS client obtains a client ID assigned during the initial negotiation phase and uses it for all subsequent requests. With the use of Client ID, the client IP address is not used in the file locking process. Therefore, the file locking problem on Exalogic does not occur when NFSv4 is used. In addition, with locking built into the protocol, NFSv4 provides better lock performance than NFSv3.

9.1.3 Naming Service

NFSv4 requires a naming service, such as Network Information Service (NIS) or LDAP, for ID mapping. You can continue to use the naming service on your existing network. The naming service is not related to the file locking problem.

> **Note:**
>
> It is assumed that you have a working NIS or LDAP service on your network. This chapter does not describe how to set up the NIS or LDAP service.

9.2 Verify the NIS Setting on Exalogic

To verify the NIS setting in your storage appliance, complete the following steps:

1. Log in to the Browser User Interface (BUI) of the storage appliance in your Exalogic machine.
2. Click **Configuration** > **SERVICES** > **NIS**. NIS Server configuration information is displayed, as shown in Figure 9-1.
3. Note down the domain name and the server IP address listed on the page. You require this information when setting up NIS clients on Exalogic compute nodes. If you do not see configuration information on this page, it means that you should install an NIS Server. Figure 9-1 uses an example NIS domain `nisdomain.example.com`.

Figure 9-1 NIS Service

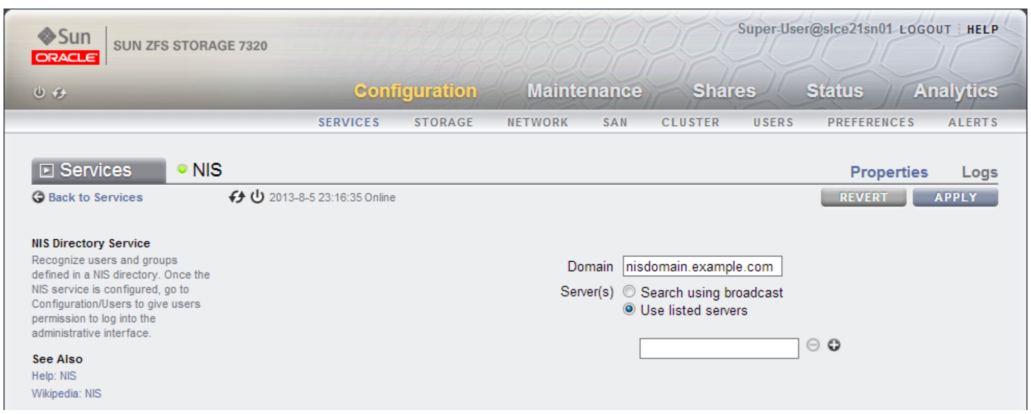

9.3 Configure the Storage Appliance

You must configure two services in the storage appliance for using NFSv4 on Exalogic:

- Configure the NFS Service
- Configure the NIS Service

9.3.1 Configure the NFS Service

Log in to the browser user interface (BUI) of the storage appliance, and configure the NFS service as follows:

1. Click **Configuration** on the home page.
2. Expand **Services** on the left navigation menu.
3. Click **NFS**. The NFS service configuration page is displayed, as shown in Figure 9-2.

Figure 9-2 NFS Service

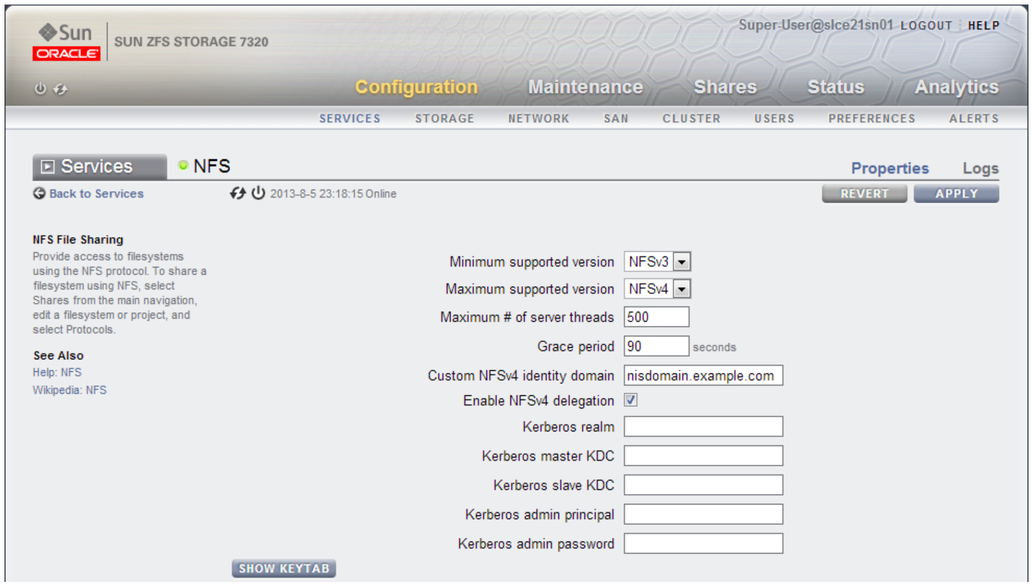

Enter information in the fields as follows:

- **Minimum supported version**: `NFSv3`
- **Maximum supported version**: `NFSv4`
- **Use DNS domain as NFSv4 identity domain**: Ensure that the option is not selected.
- **Use custom NFSv4 identity domain**: Enter the domain name of the NIS service you are using.
- **Enable NFSv4 delegation**: Select this option.
- Click **APPLY**.

9.3.2 Configure the NIS Service

Log in to the browser user interface (BUI) of the storage appliance, and configure the NIS service. For example, in Figure 9-3 the domain `nisdomain.example.com` is entered.

Figure 9-3 NIS Service

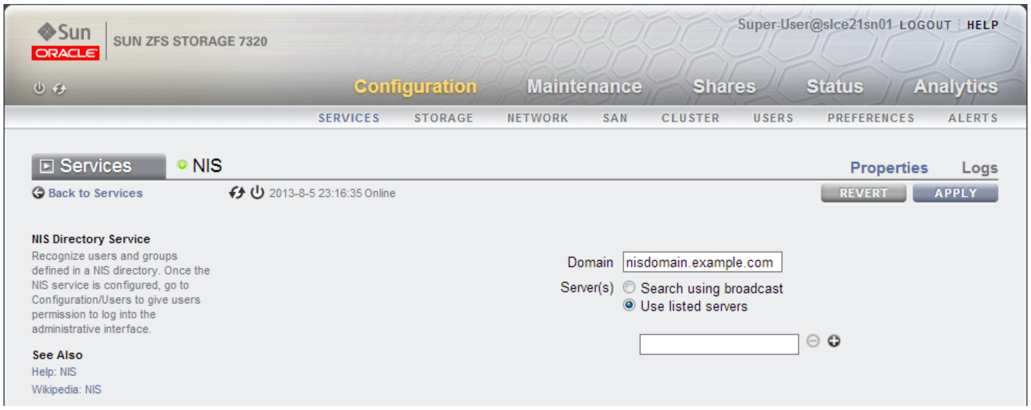

The NIS service must be up and running. For **Server**, select the **Use listed servers** option, and enter the host name or IP address of the NIS server.

9.4 Configure an Exalogic Linux Compute Node to Use NFSv4

Configure an Exalogic Linux compute node as follows:

1. Log in to the compute node as `root`.
2. Edit the `/etc/idmapd.conf` configuration file:

   ```
   vi /etc/idmapd.conf
   ```

   Set the domain value, as in the following example:

   ```
   Domain = us.myexample.com
   ```

3. Restart the `rpcidmapd` service:

   ```
   service rpcidmapd restart
   ```

4. Update the `/etc/yp.conf` configuration file, and set the correct domain value, as in the following example:

   ```
   vi /etc/yp.conf
   ```

   Add the following line:

   ```
   domain us.myexample.com server <NIS_Server_hostname_or_IP>
   ```

   Where `us.myexample.com` is the example domain and <NIS_Server_hostname_or_IP> is the host name or IP address of the NIS server. You must replace these sample values with values appropriate for your environment.

5. Set NIS domain name on the command line:

   ```
   # domainname <NIS_DOMAIN_NAME>
   ```

   For example:

   ```
   # domainname nisdomain.example.com
   ```

6. Edit the `/etc/nsswitch.conf` configuration file:

```
vi /etc/nsswitch.conf
```

   Change the following entries:

```
#passwd:     files
 passwd:     files nis
#shadow:     files
 shadow:     files nis
#group:      files
 group:      files nis
 automount:  files nis nisplus
 aliases:    files nis nisplus
```

7. Restart the `rpcidmapd` service:

```
# service rpcidmapd restart
```

8. Restart the `ypbind` service by running the following command:

```
# service ypbind restart
```

9. Check the `yp` service by running this command:

```
# ypwhich
```

10. Verify if you can access Oracle user accounts:

```
# ypcat passwd
```

11. Add `ypbind` to your boot sequence, so that it starts automatically after rebooting.

```
# chkconfig ypbind on
```

9.5 Create NFSv4 Mount Points on Oracle Linux

Create NFSv4 mount points on Oracle Linux as follows:

1. Create a new share on the storage appliance. In this example, it is assumed that the new share is accessible as `el01sn01:/export/common/patches`. This mount point is mounted on the compute node as `/u01/common/patches`.
2. On the compute node, create a new directory as follows:

```
# mkdir -p /u01/common/patches
```

3. Edit the `/etc/fstab` configuration file:

```
vi /etc/fstab
```

   Add the following line, and enter the correct values for **nfs4** and **proto=tcp**

```
el01sn01:/export/common/patches /u01/common/patches nfs4
rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp
```

4. Mount all shared volumes by running the following command:

```
# mount -a
```

10 Configure Ethernet Over InfiniBand

This chapter describes how to set up Ethernet over InfiniBand (EoIB) network configuration for your Exalogic machine.
For more information about cabling a gateway switch to an external data center 10 Gb Ethernet switch, see Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch. For information about cable and transceiver requirements, see Transceiver and Cable Requirements.

This chapter contains the following topics:

- Introduction to Virtual NICs (VNICs)
- Set Up Ethernet Over InfiniBand (EoIB) on Oracle Linux
- Set Up Ethernet Over InfiniBand (EoIB) on Oracle Solaris

10.1 Introduction to Virtual NICs (VNICs)

A virtual NIC or VNIC maps an Ethernet connector on the Sun Network QDR InfiniBand Gateway Switch to a network interface within the compute node. A connector (`0A-ETH-1` to `0A-ETH-4`, and`1A-ETH-1` to `1A-ETH-4`) hosts a 10 GbE port. In the Exalogic machine, on each Sun Network QDR InfiniBand gateway switch, you can have a maximum of eight 10 GbE uplinks. In most scenarios, the number of 10 GbE uplinks is less than 8.

One or more VNICs can be assigned to a compute node. If more than one VNIC is assigned, they can be different connectors on the same Sun Network QDR InfiniBand Gateway Switch or on different Sun Network QDR InfiniBand Gateway Switches. The default configuration is to map one connector on each of the two switches and to bond or balance the pair for increased availability.

> **Note:**
>
> In Exalogic, you create VNICs in Sun Network QDR InfiniBand Gateway Switch manual mode.

10.1.1 VNIC Resource Limit

On each Sun Network QDR InfiniBand Gateway Switch there can be a maximum of 1K VNICs per logical gateway. Logical gateway denotes either a single external Ethernet port or a defined group of external Ethernet ports. The 1K limit is shared by all VLANs on the logical gateway and each defined VLAN ID also consumes one of the 1K VNICs. For more information on VLANs see Set Up Virtual LANs .

To find out how many VNICs are currently defined, use the `showvnics` command and for VLANs, use `showvlan`.

10.2 Set Up Ethernet Over InfiniBand (EoIB) on Oracle Linux

To set up Ethernet over InfiniBand connectivity for an Exalogic compute node running Oracle Linux, complete the following steps:

1. Use an SSH client, such as PuTTY, to log in to a Sun Network QDR InfiniBand Gateway Switch. Oracle recommends that you log in as the `root` user. For example, log in to `el01gw04` as `root`.
2. At the command prompt, run the following command:

   ```
   el01gw04# listlinkup | grep Bridge
   ```

   The following is an example of the output of the `listlinkup` command:

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

   From this output, identify the uplinks. You can determine that you can use any of the following Ethernet connectors for creating a VNIC:

   - `0A-ETH-1`
   - `0A-ETH-2`
   - `0A-ETH-3`
   - `0A-ETH-4`
   - `1A-ETH-3`
   - `1A-ETH-4`

   > **Note:**
   >
   > This procedure uses `1A-ETH-3` as an example.

3. Determine GUIDs of the Exalogic compute node that requires the VNIC, as follows:

   a. On the compute node that requires the VNIC, log in as `root`, and run the `ibstat` command on the command line. For example, log in to `el01cn01` as `root`.

      Example:

      ```
      el01cn01# ibstat
      CA 'mlx4_0'
              CA type: MT26428
              Number of ports: 2
      ```

```
Firmware version: 2.7.8100
Hardware version: b0
Node GUID: 0x0021280001a0a364
System image GUID: 0x0021280001a0a367
Port 1:
        State: Active
        Physical state: LinkUp
        Rate: 40
        Base lid: 120
        LMC: 0
        SM lid: 6
        Capability mask: 0x02510868
        Port GUID: 0x0021280001a0a365
        Link layer: IB
Port 2:
        State: Active
        Physical state: LinkUp
        Rate: 40
        Base lid: 121
        LMC: 0
        SM lid: 6
        Capability mask: 0x02510868
        Port GUID: 0x0021280001a0a366
        Link layer: IB
```

In the output, information about two ports is displayed. Identify the `GUID` and `Base lid` of the port that you want to use for creating the VNIC.

For the example illustrated in this procedure, we will use the port with `GUID 0x0021280001a0a366` and `Base lid 121`.

b. On the same compute node, run the following command to view information about all the active links in the InfiniBand fabric:

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

`hostname` is the name of the compute node. You can also specify the bonded IPoIB address of the compute node.

Example:

```
el01cn01# iblinkinfo.pl -R | grep el01cn01
65   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    121   2[  ] "el01cn01 EL-
C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
64   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    120   1[  ] "el01cn01 EL-
C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
```

From the output of the `iblinkinfo` command, note the switch `lid` value (`65`, in first column) associated with the `Base lid` of the compute node port that you noted earlier (`121`, in the first line):

4. Determine the gateway switch that corresponds to the switch LID `65` by running the `ibswitches` command, as in the following example:

Example:

```
el01cn01# ibswitches
Switch  : 0x002128548042c0a0 ports 36 "SUN IB QDR GW switch el01gw03" enhanced
port 0 lid 63 lmc 0
Switch  : 0x002128547f22c0a0 ports 36 "SUN IB QDR GW switch el01gw02" enhanced
port 0 lid 6 lmc 0
Switch  : 0x00212856d0a2c0a0 ports 36 "SUN IB QDR GW switch el01gw04" enhanced
port 0 lid 65 lmc 0
```

```
Switch  : 0x00212856d162c0a0 ports 36 "SUN IB QDR GW switch el01gw05" enhanced
port 0 lid 64 lmc 0
```

`lid 65` corresponds to gateway switch `el01gw04` with GUID `0x00212856d0a2c0a0`.

5. Define a dummy MAC address in the following format:

   ```
   last3_octets_of_switchGUID : last3_octets_of_computenode_adminIP_in_hex_format
   ```

   Example:

   GUID of switch: `00:21:28:56:d0:a2:c0:a0`

   Last three octets: `a2:c0:a0`

   Administrative IP of the compute node that requires the VNIC: `192.168.1.1`

   Last three octets: `168.1.1` (in hexadecimal notation: `a8:01:01`)

   MAC address: `a2:c0:a0:a8:01:01`

   > **Note:**
   >
   > The dummy MAC address should be unique to the Exalogic network. Only even numbers are supported for the most significant byte of the MAC address (unicast). The above address is an example only.

6. As `ilom-admin`, log in to the gateway switch (`el01gw04`) that you identified in Step 4.
7. Run the following command to create a VLAN:

   ```
   hostname# createvlan connector -vlan vlan_ID -pkey default
   ```

   Example:

   ```
   e101gw04# createvlan 1A-ETH-3 -vlan 0 -pkey default
   ```

8. Run the following command to create a VNIC:

   ```
   hostname# createvnic connector -guid compute_node_port_GUID -mac
   unique_mac_address -pkey default
   ```

   Example:

   ```
   el01gw04# createvnic 1A-ETH-3 -guid 0021280001a0a366 -mac a2:c0:a0:a8:01:01 -
   pkey default
   ```

   > **Note:**
   >
   > This new resource is not tagged with any VLAN. At this time, Exalogic uses a single partition (the default partition).

   The VNIC is created.

9. To verify the VNIC, on the switch CLI, run the `showvnics` command. The following example output is displayed:

   ```
   ID  STATE FLG IOA_GUID                NODE                        IID
   MAC               VLN PKEY GW
   ```

```
--- ----- --- ---------------------- -------------------------- ----
----------------- --- ---- --------
8   UP    N   00:21:28:00:01:A0:A3:66 e101cn01 EL-C 192.168.10.29 0000
a2:c0:a0:a8:01:01 NO  ffff 1A-ETH-3
```

10. On the compute node, run the following command to display the list of VNICs available on the compute node:

    ```
    el01cn01# mlx4_vnic_info -l
    ```

    This command displays the name of the new interface, as seen on the compute node, such as `eth4`. Note this ID.

11. Create another VNIC for the same compute node, but using a connector on a different gateway switch. Note the `ethX` ID of this VNIC too.

    It is recommended that you configure the two EoIB interfaces as a bonded interface, such as `bond1`.

12. Create interface files for the VNICs on the compute node.

    To ensure correct failover behavior, the name of the VNIC interface file and the value of the `DEVICE` directive in the interface file must *not* be based on the kernel-assigned `ethX` interface name (`eth4`, `eth5`, and so on). Instead, Oracle recommends that the interface file name and value of the `DEVICE` directive in the interface file be derived from the `EPORT_ID` and `IOA_PORT` values, as follows:

    > **Note:**
    >
    > Any other unique naming scheme is also acceptable.

    a. Run the following command to find the `EPORT_ID`:

       ```
       #mlx4_vnic_info -i ethX | grep EPORT_ID
       ```

       Example:

       ```
       e101cn01#mlx4_vnic_info -i eth4 | grep EPORT_ID
       EPORT_ID     331
       ```

       Note the `EPORT_ID` that is displayed, `331` in this example.

    b. Run the following command to find the `IOA_PORT`:

       ```
       #mlx4_vnic_info -i ethX | grep IOA_PORT
       ```

       Example:

       ```
       e101cn01#mlx4_vnic_info -i eth4 | grep IOA_PORT
       IOA_PORT     mlx4_0:1
       ```

       Note the number after the colon (:) in the `IOA_PORT` value that is displayed, in this case `1`.

    c. Build the interface file name and device name by using the following convention:

       Interface file name: `ifcfg-ethA_B`

       Device name: `ethA_B`

A is the EPORT_ID, and B is the number after the colon (:) in the IOA_PORT value.

Example:

Interface file name: ifcfg-eth**331**_1

Device name: eth**331_1**

In this example, 331 is the EPORT_ID, and 1 is the value derived from the IOA_PORT.

13. Create the interface file for the first VNIC, eth4 in the example, by using a text editor such as vi.

    Save the file in the /etc/sysconfig/network-scripts directory.

    Example for Oracle Linux 6.x or previous version:

```
# more /etc/sysconfig/network-scripts/ifcfg-eth331_1
DEVICE=eth331_1
BOOTPROTO=none
ONBOOT=yes
HWADDR=a2:c0:a0:a8:01:01
MASTER=bond1
SLAVE=yes
```

    Example for Oracle Linux 7.x:

```
# more /etc/sysconfig/network-scripts/ifcfg-eth331_1
DEVICE=eth331_1
BOOTPROTO=none
ONBOOT=yes
HWADDR=a2:c0:a0:a8:01:01
MASTER=bond1
SLAVE=yes
DEVTIMEOUT=30
```

    - Make sure that the name of the interface file (ifcfg-eth331_1 in the example) is the name derived in step 12.
    - For the DEVICE directive, specify the device name (eth331_1 in the example) derived in step 12.
    - For the HWADDR directive, specify the dummy MAC address created in step 5.
    - For Oracle Linux 7.x deployments, append DEVTIMEOUT=30 to the interface file.

    > **Note:**
    >
    > The DEVTIMEOUT parameter needs to be added only in the interface configuration files, not in the corresponding bond configuration file.

14. Create an interface file for the second VNIC, say eth5. Be sure to name the interface file and specify the DEVICE directive by using a derived interface name and not the kernel-assigned name, as described earlier. In addition, be sure to specify the relevant dummy MAC address for the HWADDR directive.

15. After creating the interface files, create the ifcfg-bond1 file. If the file already exists, verify its contents.

    Example:

```
# more /etc/sysconfig/network-scripts/ifcfg-bond1
DEVICE=bond1
IPADDR=192.168.48.128
NETMASK=255.255.255.0
BOOTPROTO=none
USERCTL=no
TYPE=Ethernet
ONBOOT=yes
IPV6INIT=no
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000"
GATEWAY=192.168.48.1
```

16. Bring up the new `bond1` interface using the `ifup` command.

    You must also reboot the compute node for the changes to take effect.

10.3 Set Up Ethernet Over InfiniBand (EoIB) on Oracle Solaris

This section includes the following procedures to set up EoIB on Oracle Solaris:

- Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1
- Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2

10.3.1 Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1

You can set up Ethernet over InfiniBand connectivity for Exalogic compute nodes running Oracle Solaris 11.1 by doing the following:

1. Use an SSH client, such as PuTTY, to log in to a Sun Network QDR InfiniBand Gateway Switch as a `root`. For example, log in to `el01gw04` as `root`.
2. At the command prompt, run the following command:

   ```
   el01gw04# listlinkup | grep Bridge
   ```

   A section of the output of this command is as follows:

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

   From this example, identify the uplinks. In this example, you can use any of the following Ethernet connectors for creating a VNIC:

   - `0A-ETH-1`
   - `0A-ETH-2`
   - `0A-ETH-3`
   - `0A-ETH-4`
   - `1A-ETH-3`

- `1A-ETH-4`

> **Note:**
>
> This example procedure uses `1A-ETH-3`.

3. Determine GUIDs of an Exalogic compute node as follows:

   a. On the compute node that requires the VNIC, log in as `root`, and run the `dladm show-ib` command on the command line. For example, log in to `el01cn02` as `root`. This command displays port information, as in the following example output:

   ```
   el01cn02# dladm show-ib
   LINK     HCAGUID           PORTGUID          PORT  STATE  PKEYS
   ibp0     21280001A0A694    21280001A0A695    1     up     FFFF
   ibp1     21280001A0A694    21280001A0A696    2     up     FFFF
   ```

   In the output, information about two ports is displayed. From this output, you must determine which port GUID to use. This example procedure uses the port GUID `21280001A0A695` (port 1).

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

   ```
   el01cn02# iblinkinfo.pl -R | grep hostname
   ```

   Where `hostname` is the name of the compute node. For example, `el01cn02`.

   The following is the example output of this command:

   ```
   el01cn02# iblinkinfo.pl -R | grep el01cn02
   65   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    121   2[  ] "el01cn02 EL-
   C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
   64   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    120   1[  ] "el01cn02 EL-
   C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
   ```

   From this example output, note down the switch lid values. The switch lid of port 1 is `64` (the first column in the output). The switch lid of port 2 is `65`.

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

   a. On the compute node, run the `ibswitches` command on the command line. The example output of this command is as follows:

   ```
   el01cn02# ibswitches
   Switch  : 0x002128548042c0a0 ports 36 "SUN IB QDR GW switch el01gw03"
   enhanced port 0 lid 63 lmc 0
   Switch  : 0x002128547f22c0a0 ports 36 "SUN IB QDR GW switch el01gw02"
   enhanced port 0 lid 6 lmc 0
   Switch  : 0x00212856d0a2c0a0 ports 36 "SUN IB QDR GW switch el01gw04"
   enhanced port 0 lid 65 lmc 0
   Switch  : 0x00212856d162c0a0 ports 36 "SUN IB QDR GW switch el01gw05"
   enhanced port 0 lid 64 lmc 0
   ```

   b. In this example output, identify the switches that lid values `64` and `65` are associated with. In this example, the switch lid `64` of the gateway switch `el01gw05` with GUID `0x00212856d162c0a0` is associated with port 1 of the HCA in the compute node `el01cn02`.

> **Note:**
>
> This example procedure uses LID `64` of this gateway switch.

5. Define a dummy MAC address in the following format:

```
<last three octets from el01gw05 switch ib GUID> : <last three octets of the
administrative IP of the compute node in hexadecimal format>
```

   Example:

   GUID of switch `el01gw05`: `00:21:28:56:d1:62:c0:a0`

   Last three octets of the switch GUID: `62:c0:a0`

   Administrative IP address of compute node: `192.168.1.5`

   Last three octets of the compute node's IP address: `168.1.5`

   Last three octets in hexadecimal notation: `a8:01:05`.

   MAC address of the VNIC: `62:c0:a0:a8:01:05`

> **Note:**
>
> Each MAC address should be unique. Only even numbers are supported for the most significant byte of the MAC address (unicast). The above address is an example only.

6. As `root`, log in to `el01gw05` that you identified in Step 4. Use its IP address or host name to log in.
7. Upon login, to permit the configuration of VNICs, run the following command:

```
el01gw05# allowhostconfig
```

8. To create a VLAN, run the following command:

```
e101gw05# createvlan 1A-ETH-3 -vlan 1706 -pkey default
```

9. Note the ID of the VLAN you created by running the `showvlan` command as follows:

```
# showvlan
  Connector/LAG  VLN   PKEY
  -------------  ---   ----
  1A-ETH-3        0    ffff
  1A-ETH-3       1706 ffff
```

   In this example, the VLAN ID is `1706`.

10. Run the following command to create a VNIC on the switch:

```
el01gw05# createvnic 1A-ETH-3 -guid 00:21:28:00:01:A0:A6:95 -mac
62:c0:a0:a8:01:05 -pkey default
```

> **Note:**
>
> This new resource is not tagged with any VLAN.

A VNIC is created.

11. To verify that the VNIC was created, run the `showvnics` command. The following example output is displayed:

```
ID  STATE FLG IOA_GUID                NODE                         IID
MAC               VLN PKEY GW
--- ----- --- ----------------------- -------------------------- ----
----------------- --- ---- --------
0   UP    N   00:21:28:00:01:A0:A6:95 e101cn01 EL-C 192.168.10.29 0000
62:c0:a0:a8:01:05 NO  ffff 1A-ETH-3
```

12. On the compute node, run the following command to display the list of VNICs available on the compute node:

```
el01cn02# dladm show-phys | grep eoib
```

This command displays the name of the new interface, as seen on the compute node, such as `eoib0`. Note the corresponding link, such as `net7`. It also displays the state of the interface.

> **Note:**
>
> You may repeat the above steps to create more network- administered tagless VNICs on the same compute node as long as a unique {ETH connector, port GUID} tuple is chosen each time. When this second VNIC is configured in the same manner, the VNIC is seen on the compute node (for example, as the `eoib1` interface with the link `net8`). It is recommended that you configure these two Ethernet over InfiniBand (EoIB) interfaces in an IPMP group, such as `bond1`.
>
> To create a host-administered VNIC on a {ETH connector, port GUID} tuple with a network-administered tagless VNIC already created on it, complete the steps described in Oracle Solaris: Creating VNICs and Associating Them with VLANs.

13. Create another VNIC for the same compute node, using a connector on a different gateway switch, by following steps 1 to 12. Note the name of this interface and its corresponding link. For example, `eoib1` interface with the link `net8`.
14. Delete the following files:
    - `/etc/hostname.bond1`
    - `/etc/hostname.eoib0`
    - `/etc/hostname.eoib0`
15. Restart the compute node by running the `reboot` command.
16. Create the VNIC you created in step 8 again on the compute node by running the following command:

```
hostname# dladm create-vnic -l link_name [-v vlan_id] interface_name
```

Example:

```
el01cn02# dladm create-vnic -l net7 eoib0
el01cn02# dladm create-vnic -l net8 eoib1
```

If you are creating a VLAN tagged VNIC, use the `-v` option to add the VLAN ID as follows:

```
el01cn02# dladm create-vnic -l net7 -v 1706 eoib0
el01cn02# dladm create-vnic -l net8 -v 1706 eoib1
```

17. You can verify if the VNICs were created by using the `dladm show-vnic` command as follows:

```
hostname# dladm show-vnic
```

18. To configure `eoib0` and `eoib1` in an IPMP group for high availability purposes, do the following:

    a. Identify the data links associated with the VNICs you created on the InfiniBand switch by running the following command:

    ```
    el01cn02# dladm show-phys -m
    ```

    Identify the link names associated with the VNICs you created, such as `net7` and `net8`.

    b. Create the IPMP group by running the following command:

    ```
    hostname# ipadm create-ipmp bond_name
    ```

    Example:

    ```
    el01cn02# ipadm create-ipmp bond1
    ```

    c. Create the IP interfaces for the two links you noted in step 18.a by running the `ipadm create-ip` command as follows:

    ```
    hostname# ipadm create-ip link_name
    ```

    Example:

    ```
    el01cn02# ipadm create-ip net7
    el01cn02# ipadm create-ip net8
    ```

    d. Create interfaces for the VNICs you created in step 16 by running the following commands:

    ```
    hostname# ipadm create-ip interface_name
    ```

    Example:

    ```
    el01cn02# ipadm create-ip eoib0
    el01cn02# ipadm create-ip eoib1
    ```

    e. Set one of the interfaces as a standby for the bonded interface, by running the following command:

    ```
    hostname# ipadm set-ifprop -p standby=on -m ip interface_name
    ```

    Example:

    ```
    e101cn02# ipadm set-ifprop -p standby=on -m ip eoib1
    ```

f. Add the two interfaces to the ipmp bond you created in step 18.b, by running the following command:

```
hostname# ipadm add-ipmp -i interface_name1 -i interface_name2 bond_name
```

Example:

```
e101cn02# ipadm add-ipmp -i eoib0 -i eoib1 bond1
```

g. Set an IP address for the bonded interface you created, by running the following command:

```
hostname# ipadm create-addr –T static –a local=ipv4_address/CIDR_netmask
bond_name/v4
```

Example:

```
e101cn02# ipadm create-addr –T static –a local=10.100.44.68/22 bond1/v4
```

h. Verify that your bonded interface is up, by running the following command:

```
hostname# ipadm show-if
IFNAME     CLASS    STATE    ACTIVE OVER
lo0        loopback ok       yes    --
net0       ip       ok       yes    --
net4       ip       ok       yes    --
net8       ip       down     no     --
net9       ip       down     no     --
bond0_0    ip       ok       yes    --
bond0_1    ip       ok       no     --
bond1     ipmp     ok       yes    eoib1 eoib0
eoib1      ip       ok       no     --
eoib0      ip       ok       yes    --
```

i. Verify that your bonded interface was given an IP address by running the following command:

```
# ipadm show-addr
ADDROBJ           TYPE     STATE        ADDR
lo0/v4            static   ok           127.0.0.1/8
net0/v4           static   ok           138.3.2.87/21
net4/v4           static   ok           169.254.182.77/24
bond0/v4          static   ok           192.168.14.101/24
bond1/v4          static   ok           138.3.48.35/22
bond1/v4a         static   ok           138.3.51.1/22
lo0/v6            static   ok           ::1/128
net0/v6           addrconf ok          fe80::221:28ff:fed7:e944/10
```

10.3.2 Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2

To set up Ethernet over InfiniBand connectivity for Exalogic compute nodes running Oracle Solaris 11.2 Base Image of EECS 2.0.6.2.0 perform the following procedure:

1. Use an SSH client, such as PuTTY, to log in to a compute node as `root`. For this example log in to `el01cn16` as `root`.
2. Run the following command to verify that the image version is EECS 2.0.6.2.0 or greater, and the kernel version is SunOS 11.2.

   ```
   root@el01cn16:~# imageinfo
   ```

   A section of the output of this command is as follows:

```
Exalogic 2.0.6.2.0 (build:r240216)
Image version       : 2.0.6.2.0
. . .
Kernel version      : SunOS 11.2
. . .
```

3. Get the names of the InfiniBand (IB) datalink by running the following command:

```
root@el01cn16:~# dladm show-phys
```

The following is a section of the output of the command that displays `net4` and `net5` as the names for the IB datalink:

```
LINK              MEDIA                STATE      SPEED  DUPLEX    DEVICE
. . .
net4              Infiniband           up         32000  unknown   ibp0
net5              Infiniband           up         32000  unknown   ibp1
```

4. Open a second terminal and log in as `root` to the switch `net5` is connected to (`el01sw-ib02` for this example).

Run the `showvlan` command to verify that the VLAN 0 is associated with the default partition and that the VLAN ID is created on the correct IB partition. The following example displays the output from this command showing that the VLAN ID is `3066` and it's associated to the correct IB partition:

```
[root@el01sw-ib02 ~]# showvlan
   Connector/LAG  VLN   PKEY
   -------------  ---   ------
   0A-ETH-1        0    0xffff
   0A-ETH-1        3066 0x8206
```

5. Repeat the previous step on the switch net4 is connected to (`el01sw-ib03` for this example).

6. From the compute node session, use the `dladm` command to verify that the compute node GUIDs are included in the IB partition that the VLAN `3066` is using.

See the following extract of the command output for `net4` and net5:

```
root@el01cn16:~# dladm show-ib net4
LINK  HCAGUID     PORTGUID         PORT STATE  GWNAME      GWPORT   PKEYS
net4  2128000...  21280001EFF369  1    up     el01sw-ib02 0a-eth-1 7FFF,8206,FFFF
                                               el01sw-ib02 0a-eth-2
                                               el01sw-ib02 0a-eth-3
                                               . . .

root@el01cn16:~# dladm show-ib net5
LINK  HCAGUID     PORTGUID         PORT STATE  GWNAME      GWPORT   PKEYS
net5  2128000...  21280001EFF36A  2    up     el01sw-ib02 0a-eth-1 7FFF,8206,FFFF
                                               el01sw-ib02 0a-eth-2
                                               el01sw-ib02 0a-eth-3
                                               . . .
```

Make a note of the data that the command displays in the PORTGUID column for both IB datalinks.

7. Run the commands `iblinkinfo`, `ibswitches`, and `ibstat` to determine the mapping among the IB HCA ports, IB datalinks and IB switches. See the following section of the first command output:

```
root@el01cn16:~# iblinkinfo|grep cn16
. . .
```

```
   14    33[  ] ==(...)==>  72     2[  ] "el01cn16 EL-C  192.168.10.16 HCA-1" ( )
   15    33[  ] ==(...)==>  71     1[  ] "el01cn16 EL-C  192.168.10.16 HCA-1" ( )
```

The output of the command displays a pair of value sets:

- switch lid 14, base lid 72, port 2
- switch lid 15, base lid 71, port 1

The following is an extract of the second command:

```
root@el01cn16:~# ibswitches
. . .
Switch  : 0x0010e00b4520c0a0 ports 36 "SUN IB QDR GW switch el01sw-ib02
10.128.21.186 leaf:1" enhanced port 0 lid 14 lmc 0
. . .
Switch  : 0x0010e00b6d80c0a0 ports 36 "SUN IB QDR GW switch el01sw-ib03
10.128.21.187 leaf:2" enhanced port 0 lid 15 lmc 0
```

The lid and port data from the previous command output matches the following switches:

- switch lid 14, base lid 72 matches `el01sw-ib02`
- switch lid 15, base lid 71 matches `el01sw-ib03`

The following is an extract of the third and last command:

```
root@el01cn16:~# ibstat
. . .
        Port 1:
                . . .
                Base lid: 71
                . . .
                Port GUID: 0x0021280001eff369
                Link layer: IB
        Port 2:
                . . .
                Base lid: 72
                . . .
                Port GUID: 0x0021280001eff36a
                Link layer: IB
```

With this information you can see that `0x0021280001eff369` is the port GUID for `net4` and `0x0021280001eff36a` is the port GUID for `net5` (see 6). Now you can determine the following mappings:

- `port1 -> net4 -> el01sw-ib03`
- `port2 -> net5 -> el01sw-ib02`

8. Run the following commands to create the EoIB datalink over `net4` and display the results of the procedure:

```
root@el01cn16:~# dladm create-eoib  -l net4 -g  el01sw-ib03 -c 0A-ETH-1 eoib0

root@el01cn16:~# dladm show-eoib
LINK       GWNAME       GWPORT    GWID FLAGS  SPEED  MACADDRESS      OVER
eoib0      el01sw-ib03  0a-eth-1  506  aH---- 10000  0:0:0:0:0:0     net4
```

9. Open another terminal and log in as `root` to the switch `net4` is connected to (`el01sw-ib03` for this example).

   Run the following commands to create a VNIC with no VLAN tag and display the results of the procedure. The following is an example of the commands' output:

```
[root@el01sw-ib03 ~]# createvnic 0A-ETH-1 -guid 0021280001EFF369 -mac
80:C0:A0:09:16:01
vNIC created

[root@el01sw-ib03 ~]# showvnics |grep cn16

105 WAIT-IOA    N 0021280001EFF369        el01cn16 EL-C  192.168.10.16  0000
80:C0:A0:09:16:01 NO  0xffff 0A-ETH-1
```

10. From the compute node session create a host-based VNIC with a VLAN tag. Run the following commands:

```
root@el01cn16:~# dladm create-vnic -l eoib0 -v 3066 vnic3066_0

root@el01cn16:~# dladm show-vnic
```

The following is an example of the `show-vnic` command output:

```
LINK                OVER              SPEED  MACADDRESS        MACADDRTYPE VIDS
vnic3066_0          eoib0             10000  2:8:20:42:a1:f1   random      3066
```

11. Run the following commands to create the EoIB datalink over `net5`:

```
root@el01cn16:~# dladm create-eoib  -l net5 -g  el01sw-ib02 -c 0A-ETH-1 eoib1

root@el01cn16:~# dladm show-eoib
```

The following is an example of the `show-eoib` command output:

```
LINK        GWNAME      GWPORT   GWID FLAGS  SPEED  MACADDRESS        OVER
eoib0       el01sw-ib03 0a-eth-1 506  aHnU-- 10000  80:c0:a0:9:16:1   net4
eoib1       el01sw-ib02 0a-eth-1 286  aH---- 10000  0:0:0:0:0:0       net5
```

12. Log into the switch that `net5` is connected to and run the following command to create a VNIC with no VLAN tag:

```
[root@el01sw-ib02 ~]# createvnic 0A-ETH-1 -guid 0021280001eff36a -mac
00:14:4F:09:16:02
vNIC created
```

Run the following command to display the result of the creation of the VNIC:

```
[root@el01sw-ib02 ~]# showvnics|grep cn16
108 WAIT-IOA    N 0021280001EFF36A        el01cn16 EL-C  192.168.10.16  0000
00:14:4F:09:16:02 NO  0xffff 0A-ETH-1
```

13. From the compute node session run the following command to create a VNIC with a VLAN tag:

```
root@el01cn16:~# dladm create-vnic -l eoib1 -v 3066 vnic3066_1
00:14:4F:09:16:02
vNIC created
```

Run the following command to display the result of the creation of the VNIC:

```
root@el01cn16:~# dladm show-vnic
LINK                OVER              SPEED  MACADDRESS        MACADDRTYPE VIDS
vnic3066_0          eoib0             10000  2:8:20:42:a1:f1   random      3066
vnic3066_1          eoib1             10000  2:8:20:10:7f:d3   random      3066
```

14. Run the following commands to create the IPMP `bond1` group:

```
root@el01cn16:~# ipadm create-ip vnic3066_0

root@el01cn16:~# ipadm create-ip vnic3066_1
```

```
root@el01cn16:~# ipadm delete-ipmp bond1

root@el01cn16:~# ipadm create-ipmp -i vnic3066_0 -i vnic3066_1 bond1

root@el01cn16:~# ipadm create-addr -T static -a 192.168.100.16/24 bond1/v4

root@el01cn16:~# ipadm set-ifprop -p standby=on -m ip vnic3066_1

root@el01cn16:~# ipmpstat -i
INTERFACE   ACTIVE  GROUP       FLAGS     LINK      PROBE     STATE
vnic3066_1  no      bond1       is-----   up        disabled  ok
vnic3066_0  yes     bond1       --mbM--   up        disabled  ok
bond0_1     no      bond0       is-----   up        disabled  ok
bond0_0     yes     bond0       --mbM--   up        disabled  ok
```

15. From the session to switch `el01sw-ib03` run the following command to verify that the active VNIC is up on the switch `el01sw-ib03`:

```
[root@el01sw-ib03 ~]# showvnics |grep cn16
105 UP         N 0021280001EFF369       el01cn16 EL-C  192.168.10.16  31744
80:C0:A0:09:16:01 NO  0xffff 0A-ETH-1
106 UP         H 0021280001EFF369       el01cn16 EL-C  192.168.10.16  64513
02:08:20:42:A1:F1 3066 0x8206 0A-ETH-1
```

16. On the switch `el01sw-ib02`, the passive VNIC is not expected to appear until the IPMP group failover process runs. From the session to the switch `el01sw-ib02` run the following command to verify the state of the VNIC:

```
[root@el01sw-ib02 ~]# showvnics|grep cn16
108 UP         N 0021280001EFF36A       el01cn16 EL-C  192.168.10.16  31744
00:14:4F:09:16:02 NO  0xffff 0A-ETH-1
```

11
Set Up Virtual LANs

This chapter describes how to set up a VLAN on the Ethernet connector of a Sun Network QDR InfiniBand Gateway Switch, and it describes how to assign a virtual NIC (VNIC) on a compute node to use that VLAN.
The tasks described in this chapter are optional.

It contains the following topics:

- Introduction to VLAN
- Example Scenario
- Tag the Ethernet Connectors With a VLAN Identifier
- Oracle Linux: Creating VNICs and Associating Them with VLANs
- Oracle Solaris: Creating VNICs and Associating Them with VLANs

11.1 Introduction to VLAN

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

> **Note:**
>
> You can create more than one VLAN per Ethernet connection.
>
> For a general introduction to VNICs, see Introduction to Virtual NICs (VNICs).

11.2 Example Scenario

To understand the use of VLANs in an Exalogic environment, consider the following example scenario.

You want to combine Production, Test, and Development environments in the same Exalogic machine. However, you do not want these systems to communicate with each other directly. The production systems require dedicated Ethernet interfaces. You wish to share resources, such as Ethernet connectors, between Test and Development systems.

For the Production systems, you may dedicate a few of the external 10 GbE connectors on the gateway switch (for example, `0A-ETH-1` to `0A-ETH-4`). Production systems will be on one VLAN using these four dedicated 10 GbE external uplinks.

For the Development systems, you may use one Ethernet connector on the gateway switch and a VLAN of their own. For example, `1A-ETH-3` associated with VLAN ID `10`. In this VLAN, the resources using `1A-ETH-3` are dedicated to Development systems.

The Test systems require two Ethernet interfaces, and they can use Development's Ethernet connector `1A-ETH-3`, but on a different VLAN. For example, you can create two Ethernet interfaces using `1A-ETH-3` associated with a VLAN ID `11` for use by the Test systems. In this manner, Development systems on their VLAN get their resources while sharing the Ethernet connector or uplink with Test systems. Since the two VLANs exist on the same 10 GbE Ethernet link on the gateway switch, any traffic between the two VLANs should travel through an external router if they are required to be seen by each other.

> **Note:**
>
> On a single Ethernet connection using the connector on the gateway switch, you can create up to 4094 VLANs.

11.3 Tag the Ethernet Connectors With a VLAN Identifier

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

In this process, you are mapping the following:

- Ethernet connector on the Sun Network QDR InfiniBand Gateway Switch (`0A-ETH-1` to `0A-ETH-4`, and `1A-ETH-1` to `1A-ETH-4`)
- VLAN ID (2 to 4094)
- InfiniBand partition key (`0xFFFF`)

> **Note:**
>
> Exalogic uses the default partition, and the partition key is `0xFFFF`. You can associate multiple VLANs to a single Ethernet connector.

For example, you can associate VLAN identifiers `10` and `11` to the same Ethernet connector `1A-ETH-3`.

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

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

   ```
   # createvlan 1A-ETH-3 -VLAN 10 -PKEY default
   ```

   Where `1A-ETH-3` is the Ethernet connector on the gateway switch, `10` is the VLAN identifier, and `default` is the partition key used in Exalogic.

```
# createvlan 1A-ETH-3 -VLAN 11 -PKEY default
```

Where `1A-ETH-3` is the Ethernet connector, `11` is the VLAN identifier, and `default` is the partition key used in Exalogic.

If you are using Oracle Solaris compute nodes, you should also enable the connector for untagged traffic, by running the following command on the gateway switch:

```
# createvlan 1A-ETH-3 -VLAN -1 -PKEY default
```

2. To verify, run the following command:

```
# showvlan
```

The following information is displayed:

```
Connector/LAG    VLN   PKEY
--------------   ---   -----
1A-ETH-3           0   ffff
1A-ETH-3          10   ffff
1A-ETH-3          11   ffff
```

> **Tip:**
>
> See the Example Scenario for more information.

11.4 Oracle Linux: Creating VNICs and Associating Them with VLANs

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

To create a VNIC and associate with a VLAN, use the following example procedure:

1. On the gateway switch CLI, as `root`, complete the steps **1** through **6**, as described in Set Up Ethernet Over InfiniBand (EoIB) on Oracle Linux.
2. Run the following command to create a VNIC and associate it with a VLAN (for example, with VLAN `10`, as shown in Tag the Ethernet Connectors With a VLAN Identifier):

```
# createvnic 1A-ETH-3 -GUID 00:21:28:56:d0:a2:c0:a0 -mac a2:c0:a0:a8:1:1 -vlan
10 -pkey default
```

Where `1A-ETH-3` is the Ethernet connector, `00:21:28:56:d0:a2:c0:a0` is the GUID, `a2:c0:a0:a8:1:1` is the dummy MAC address, `10` is the VLAN identifier, and `default` is the partition key used in Exalogic.

This example creates a VNIC, such as `eth4`.

3. To create a second VNIC using the same Ethernet connector and GUID, run the following command to tag the VNIC with a different VLAN identifier (`11`):

```
# createvnic 1A-ETH-3 -GUID 00:21:28:56:d0:a2:c0:a0 -mac a2:c0:a0:a8:1:a -vlan
11 -pkey default
```

Where `1A-ETH-3` is the Ethernet connector, `00:21:28:56:d0:a2:c0:a0` is the GUID, `a2:c0:a0:a8:1:a` is the dummy MAC address defined for this second interface, `11` is the VLAN identifier, and `default` is the partition key used in Exalogic.

This example creates a VNIC, such as `eth5`.

4. Run the following command to verify the VNICs:

```
# showvnics
```

The following message is displayed:

```
ID  STATE    FLG  IOA_GUID          NODE      IID     MAC     VLN  PKEY  GW
--- --------  ---  ---------------------- ---------- ---- -------------
  8 UP    N 00:21:28:00:01:A0:A3:65 computenode1 EL-C 192.168.10.29 0000
a2:c0:a0:a8:1:1 10  ffff   1A-ETH-3
  9 UP    N 00:21:28:00:01:A0:A3:65 computenode1 EL-C 192.168.10.29 0001
a2:c0:a0:a8:1:a 11  ffff   1A-ETH-3
```

> **Tip:**
>
> After creating the interfaces, you can run the `ifconfig` command with the `-a` option to verify the MAC address on the compute node. For example, to verify the new interface and its MAC address, run the following command on the Oracle Linux compute node for which the VNIC was created:
>
> ```
> # ifconfig -a eth4
> ```
>
> The output of this command shows the `HWADDR`, which is the MAC address you defined for the VNIC in Set Up Ethernet Over InfiniBand (EoIB) on Oracle Linux.

5. If you want your VNIC configuration to persist across reboots, you should save VNIC configuration to a file. For information about doing this on Oracle Linux, seeSet Up Ethernet Over InfiniBand (EoIB) on Oracle Linuxfor more information. Be sure to create a bonded interface comprising two VNICs on Oracle Linux, for high availability purposes.

11.5 Oracle Solaris: Creating VNICs and Associating Them with VLANs

If you wish to associate a VNIC with a VLAN, you should provide a VLAN identifier when mapping the MAC address.

To create a VNIC and associate with a VLAN, use the following example procedure:

1. On the gateway switch CLI, as `root`, complete one of the following procedures depending on the OS version:
   - steps 1 to 15 as described in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1
   - steps 1 to 13 as described in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2

   This procedure creates a VNIC, such as `eoib0`.

2. On Oracle Solaris, VLAN-tagged VNICs may only be created from the Oracle Solaris compute node. Run one of the following procedures depending on the OS version:
   - step 16 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1
   - steps 10 and 13 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2
3. You can verify that the VNICs were created by running one of the following procedures depending on the OS version:
   - step 17 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1.
   - second command in step 10 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2
4. Run the following command on the gateway switch CLI, as root, to verify the VNICs:

   ```
   # showvnics
   ```

   The following message is displayed:

   ```
   ID  STATE    FLG  IOA_GUID              NODE     IID    MAC    VLN    PKEY    GW
   --- -------- --- ----------------------- ---------- ---- -------------
   0   UP       N    00:21:28:00:01:A0:A6:95  computenode1  EL-C  192.168.10.29
   0000  62:C0:A0:A8:01:05  NO  ffff   1A-ETH-3
   1   UP       H    00:21:28:00:01:A0:A6:95  computenode1  EL-C  192.168.10.29
   8001  62:C0:A0:A8:01:02  10  ffff   1A-ETH-3
   2   UP       H    00:21:28:00:01:A0:A6:95  computenode1  EL-C  192.168.10.29
   8002  62:C0:A0:A8:01:03  11  ffff   1A-ETH-3
   ```

   Note that the VNIC with ID `0` corresponds to an already created network-administered VNIC (created using the steps described in Set Up Ethernet Over InfiniBand (EoIB) on Oracle Solaris). The two new host-administered interfaces are the ones with IDs `1` and `2` (with VLANs `10` and `11`, respectively).

   > **Tip:**
   >
   > After creating the interfaces, you can run the `dladm` command to verify the MAC address on the compute node. For example, to verify the new interface and its MAC address, run one of the following commands on the Oracle Solaris compute node for which the VNIC was created:
   >
   > ```
   > # dladm show-vnic eoib0_v10
   > ```
   >
   > The output of this command shows the MAC address, which is the MAC address you defined for the VNIC.

5. If you want to configure the VNIC interfaces in an IPMP group on Oracle Solaris, for high availability purposes, perform on of the following procedures depending on the OS version:
   - step 18 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.1.
   - step 14 in Set Up Ethernet Over InfiniBand on Oracle Solaris 11.2

12 Use the InfiniBand Gateway Switches

This chapter describes how to use the Sun Network QDR InfiniBand Gateway switches in your Exalogic machine. The number of gateway switches depends on your purchased Exalogic machine rack configuration.
It also describes how to manage the InfiniBand network using Subnet Manager.

This chapter contains the following sections:

- Using Sun Network QDR InfiniBand Gateway Switches
- What Next?

12.1 Using Sun Network QDR InfiniBand Gateway Switches

This section contains the following topics:

- Physical Specifications
- Access the Command-Line Interface (CLI) of a Gateway Switch
- Verify the Status of a Gateway Switch
- Start the Subnet Manager Manually
- Check Link Status
- Verify the InfiniBand Fabric
- Monitor a Gateway Switch Using Web Interface

12.1.1 Physical Specifications

This section introduces Sun Network QDR InfiniBand Gateway Switches, which are also referred to as leaf switches in this guide.

Table 12-1 provides the physical specifications of the Sun Network QDR InfiniBand Gateway Switch.

Table 12-1 NM2-GW Specifications

Dimension	Measurements
Width	17.52 in. (445.0 mm)
Depth	24 in. (609.6 mm)
Height	1.75 in. (44.5 mm)
Weight	23.0 lbs (11.4 kg)

12.1.2 Access the Command-Line Interface (CLI) of a Gateway Switch

With power applied, you can access the CLI of a gateway switch in your Exalogic machine.

The number of gateway switches in your Exalogic machine depends on your purchased Exalogic machine rack configuration. You must access the command-line interfaces of these gateway switches individually.

For example, to access the CLI of a gateway switch, complete the following steps:

1. If you are using a network management port, begin network communication with the CLI using the **ssh** command and the host name configured for the gateway switch:

   ```
   % ssh -l root gateway-name
   root@gateway-name's password: password
   #
   ```

   where `gateway-name` is the host name configured for the gateway switch.

   If you do not see this output or prompt, there is a problem with the network communication, host name, or CLI.

2. If you are using a USB management port, begin serial communication with the CLI as follows:

   a. Connect a serial terminal, terminal server, or workstation with a TIP connection to the USB-to-serial adapter. Configure the terminal or terminal emulator with these settings:

      115200 baud, 8 bits, No parity, 1 Stop bit, and No handshaking

   b. Press the Return or Enter key on the serial device several times to synchronize the connection. You might see text similar to the following:

      ```
      …
      CentOS release 5.2 (Final)
      Kernel 2.6.27.13-nm2 on an i686

      gateway-name login: root
      Password: password
      #
      ```

      where `gateway-name` is the host name assigned to the gateway switch.

      If you do not see this output or prompt, there is a problem with the network communication, host name, or command-line interface (CLI).

3. **Note:**

   Repeat these steps to access the CLI for the other gateway switches in your Exalogic machine.

12.1.3 Verify the Status of a Gateway Switch

For each gateway switch, you can check the status of the CLI, power supplies, fans, and switch chip. Verify that the voltage and temperature values of the gateway switch are within specification:

```
# showunhealthy
# env_test
```

An unfavorable output from these commands indicates a hardware fault with that particular component. A voltage or temperature deviating more than 10% from the provided specification means a problem with the respective component.

For example, on the CLI of one of the gateway switches, enter the following command to check its status:

```
# env_test
```

This command performs a set of checks and displays the overall status of the gateway switch, as in the following example:

```
Environment test started:
Starting Voltage test:
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.37 V
Measured 12V = 12.06 V
Measured 5V = 5.03 V
Measured VBAT = 3.25 V
Measured 1.0V = 1.01 V
Measured I4 1.2V = 1.22 V
Measured 2.5V = 2.52 V
Measured V1P2 DIG = 1.17 V
Measured V1P2 AND = 1.16 V
Measured 1.2V BridgeX = 1.21 V
Measured 1.8V = 1.80 V
Measured 1.2V Standby = 1.20 V
Voltage test returned OK
Starting PSU test:
PSU 0 present
PSU 1 present
PSU test returned OK
Starting Temperature test:
Back temperature 23.00
Front temperature 32.62
SP temperature 26.12
Switch temperature 45, maxtemperature 45
Bridge-0 temperature 41, maxtemperature 42
Bridge-1 temperature 43, maxtemperature 44
Temperature test returned OK
Starting FAN test:
Fan 0 not present
Fan 1 running at rpm 11212
Fan 2 running at rpm 11313
Fan 3 running at rpm 11521
Fan 4 not present
FAN test returned OK
Starting Connector test:
Connector test returned OK
```

```
Starting onboard ibdevice test:
Switch OK
Bridge-0 OK
Bridge-1 OK
All Internal ibdevices OK
Onboard ibdevice test returned OK
Environment test PASSED
```

When the status is operational, you can start the Subnet Manager (SM).

> **Note:**
>
> Repeat these steps to verify the status of the other gateway switches in your Exalogic machine.

12.1.4 Start the Subnet Manager Manually

The Subnet Manager (SM) is enabled on the gateway switches in a single Exalogic rack configuration, by default.

However, if the SM is not running on the InfiniBand switches, you can start and activate the SM as follows:

1. On the CLI of a switch, start the SM by running the following command:

   ```
   # enablesm
   ```

2. Set the SM priority within the command-line interface (CLI) as follows:

   ```
   # setsmpriority priority
   ```

   > **Note:**
   >
   > For information about the switches on which the SM should run in various rack configurations and the SM priorities for the switches, see Subnet Manager Operation in Different Rack Configurations.

   For example, to set the SM on a gateway switch to priority **5**, run the following command:

   ```
   # setsmpriority 5
   ```

   The following output is displayed:

   ```
   -------------------------------------------------
   OpenSM 3.2.6_20090717
     Reading Cached Option File: /etc/opensm/opensm.conf
     Loading Cached Option:routing_engine = ftree
     Loading Cached Option:sminfo_polling_timeout = 1000
     Loading Cached Option:polling_retry_number = 3
   Command Line Arguments:
     Priority = 5
     Creating config file template '/tmp/osm.conf'.
     Log File: /var/log/opensm.log
   -------------------------------------------------
   ```

For the changes to take effect, restart the SM as follows:

```
# disablesm
```

```
# enablesm
```

12.1.5 Check Link Status

After starting the SM, you can verify that the Link LEDs for cabled links are green. If the Link LED is dark, the link is down. If the Link LED flashes, there are symbol errors.

To check the link status of the cables:

```
# listlinkup
```

If the link for a connector is reported as not present, the link at either end of the cable is down. If a port is down, use the `enableswitchport 0 portnumber` command to bring the port up. Alternatively, use the `ibdevreset` command to reset the switch chip.

See the *Sun Network QDR InfiniBand Gateway Switch Administration Guide*, "Enable a Switch Chip Port" and "Reset the Switch Chip".

After making sure that the link is up, you can verify the InfiniBand fabric.

The following is an output example of the `listlinkup` command:

```
# listlinkup
Connector  0A Present <-> Switch Port 20 up (Enabled)
Connector  1A Present <-> Switch Port 22 up (Enabled)
Connector  2A Present <-> Switch Port 24 up (Enabled)
.
.
.
Connector 15A Not present
Connector 0A-ETH Present
 Bridge-0-1 Port 0A-ETH-1 up (Enabled)
 Bridge-0-1 Port 0A-ETH-2 up (Enabled)
 Bridge-0-0 Port 0A-ETH-3 up (Enabled)
 Bridge-0-0 Port 0A-ETH-4 up (Enabled)
Connector 1A-ETH Present
 Bridge-1-1 Port 1A-ETH-1 up (Enabled)
 Bridge-1-1 Port 1A-ETH-2 up (Enabled)
 Bridge-1-0 Port 1A-ETH-3 up (Enabled)
 Bridge-1-0 Port 1A-ETH-4 up (Enabled)
Connector 0B Present <-> Switch Port 19 up (Enabled)
Connector 1B Present <-> Switch Port 21 up (Enabled)
.
.
.
Connector 15B Not present
#
```

12.1.6 Verify the InfiniBand Fabric

Use the following commands on the command-line interface (CLI) to verify that the InfiniBand fabric is operational:

1. `ibnetdiscover`

   Discovers and displays the InfiniBand fabric topology and connections. See Discover the InfiniBand Network Topology.

2. `ibdiagnet`

   Performs diagnostics upon the InfiniBand fabric and reports status. See Perform Diagnostics on the InfiniBand Fabric.

3. `ibcheckerrors`

   Checks the entire InfiniBand fabric for errors. See Validate and Check Errors in the InfiniBand Fabric.

12.1.6.1 Discover the InfiniBand Network Topology

To discover the InfiniBand network topology and build a topology file which is used by the OpenSM Subnet Manager, run the following command on the command-line interface (CLI) of a gateway switch:

```
# ibnetdiscover
```

The output is displayed, as in the following example:

The topology file is used by InfiniBand commands to scan the InfiniBand fabric and validate the connectivity as described in the topology file, and to report errors as indicated by the port counters.

```
# Topology file: generated on Sat Apr 13 22:28:55 2002
#
# Max of 1 hops discovered
# Initiated from node 0021283a8389a0a0 port 0021283a8389a0a0
vendid=0x2c9
devid=0xbd36
sysimgguid=0x21283a8389a0a3
switchguid=0x21283a8389a0a0(21283a8389a0a0)
Switch   36 "S-0021283a8389a0a0" # "Sun DCS 36 QDR switch localhost" enhanced port 0
lid 15 lmc 0
[23]     "H-0003ba000100e388"[2](3ba000100e38a) # "nsn33-43 HCA-1" lid 14 4xQDR
vendid=0x2c9
devid=0x673c
sysimgguid=0x3ba000100e38b
caguid=0x3ba000100e388
Ca   2 "H-0003ba000100e388" # "nsn33-43 HCA-1"
[2](3ba000100e38a)   "S-0021283a8389a0a0"[23] # lid 14 lmc 0 "Sun DCS 36 QDR switch
localhost" lid 15 4xQDR
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

12.1.6.2 Perform Diagnostics on the InfiniBand Fabric

To perform a collection of tests on the InfiniBand fabric and generate several files that contain parameters and aspects of the InfiniBand fabric, run the following command on the command-line interface (CLI) on a gateway switch:

```
# ibdiagnet
```

In the following example, the `ibdiagnet` command is minimized to determine which links are utilized:

```
# ibdiagnet -lw 4x -ls 10 -skip all

Loading IBDIAGNET from: /usr/lib/ibdiagnet1.2
-W- Topology file is not specified.
 Reports regarding cluster links will use direct routes.
Loading IBDM from: /usr/lib/ibdm1.2
-I- Using port 0 as the local port.
-I- Discovering ... 2 nodes (1 Switches & 1 CA-s) discovered.
.
.
.
-I- Links With links width != 4x (as set by -lw option)
-I---------------------------------------------------
-I- No unmatched Links (with width != 4x) were found
-I---------------------------------------------------
-I- Links With links speed != 10 (as set by -ls option)
-I---------------------------------------------------
-I- No unmatched Links (with speed != 10) were found
.
.
.
-I- Stages Status Report:
 STAGE               Errors Warnings
 Bad GUIDs/LIDs Check               0   0
 Link State Active Check               0   0
 Performance Counters Report               0   0
 Specific Link Width Check               0   0
 Specific Link Speed Check               0   0
 Partitions Check               0   0
 IPoIB Subnets Check               0   0
Please see /tmp/ibdiagnet.log for complete log
----------------------------------------------------------------
-I- Done. Run time was 1 seconds.
```

Note:

The actual output for your InfiniBand fabric will differ from that in the example.

12.1.6.3 Validate and Check Errors in the InfiniBand Fabric

Use the `ibcheckerrors` command that uses the topology file to scan the InfiniBand fabric and validate the connectivity as described in the topology file, and to report errors as indicated by the port counters.

On the command-line interface (CLI), enter the following command:

```
# ibcheckerrors

## Summary: 4 nodes checked, 0 bad nodes found
##      34 ports checked, 0 ports have errors beyond threshold
```

Note:

The actual output for your InfiniBand fabric will differ from that in the example.

12.1.7 Monitor a Gateway Switch Using Web Interface

1. Open a web browser and go to the following URL:

   ```
   http://gateway-IP
   ```

   where `gateway-IP` is the IP address of a gateway switch.
2. Log in to the interface as the `root` user.
3. Click the **Switch/Fabric Monitoring Tools** tab.
4. Click **Launch Sun DCS GW Monitor**.

   The Fabric Monitor is displayed.

12.2 What Next?

The Sun Network QDR InfiniBand Gateway Switch is installed with the default vNICs (vnic0 and vnic1) configured on separate Sun Network QDR InfiniBand Gateway Switches for the Ethernet over InfiniBand (EoIB) BOND1 interface.

Optionally, you can create VLANs and vNICs using the InfiniBand gateway switches.

13
Manage the InfiniBand Network Using Subnet Manager

This chapter describes how to manage the InfiniBand Network Using Subnet Manager in your Exalogic machine.
This chapter contains the following sections:

- Understand Administrative Commands
- Manage InfiniBand Network Using Subnet Manager
- Work with the Default Rack-Level InfiniBand Partition

13.1 Understand Administrative Commands

The following topics provide an overview of administrative tasks and the command sets to perform those tasks. Administering the gateway requires accessing the command-line interface (CLI), which is also referred to as the **management controller**.

This section contains the following topics:

- Hardware Command Overview
- InfiniBand Command Overview

13.1.1 Hardware Command Overview

The CLI (management controller) uses a simplified Linux operating system and file system. From the **#** prompt on the CLI, you can type hardware commands to perform some administrative and management tasks. Hardware commands are user-friendly and can perform some testing upon the switch chip, enabling greater control of a gateway switch and its operation.

After you log in to the `root` account, the shell prompt (**#**) appears, and you can enter shell commands. Enter the hardware commands in the following format:

```
# command [arguments][arguments]...
```

13.1.2 InfiniBand Command Overview

The InfiniBand commands are a means of monitoring and controlling aspects of the InfiniBand fabric. These commands are also installed on and run from the CLI, which is also the host of the Subnet Manager. Use of these commands requires thorough knowledge of InfiniBand architecture and technology.

After you log in to the `root` account, the shell prompt (**#**) appears, and you can enter shell commands. Enter the InfiniBand commands in the following format:

```
# command [option][option] ...
```

13.2 Manage InfiniBand Network Using Subnet Manager

This section contains the following topics:

- Overview of Subnet Manager
- Subnet Manager Operation in Different Rack Configurations
- Monitor the Subnet Manager
- Control the Subnet Manager

13.2.1 Overview of Subnet Manager

The subnet manager (SM) manages all operational characteristics of the InfiniBand network, such as the following:

- Discovering the network topology
- Assigning a local identifier (LID) to all ports connected to the network
- Calculating and programming switch forwarding tables
- Programming Partition Key (PKEY) tables at HCAs and switches
- Programming QoS tables (Service Level to Virtual Lane mapping tables, and Virtual Lane arbitration tables)
- Monitoring changes in the fabric

The InfiniBand network typically has more than one SM, but only one SM is active at a time. The active SM is Master SM, others are Standby SMs. If the master SM shuts down or fails, a standby SM will automatically become the master SM.

Note:

In the Exalogic machine, the InfiniBand switches (both leaf and spine) are automatically configured to separate the IP over InfiniBand (IPoIB) traffic and the Ethernet over InfiniBand (EoIB) traffic.

13.2.2 Subnet Manager Operation in Different Rack Configurations

Table 13-1 provides information about the switches on which the subnet manager should run in different rack configurations.

Table 13-1 Running the Subnet Manager in Different Rack Configurations

Rack Configuration	SM Should Run On...	SM Priority
Single Exalogic machine	All leaf switches	All leaf switches: **5**
Two half- or full-rack Exalogic machines	Spine switches	Spine switch: **8**
Two quarter-rack Exalogic machines	All leaf switches	All leaf switches: **5**
Three or more Exalogic machines	Spine switches	Spine switch: **8**

Table 13-1 (Cont.) Running the Subnet Manager in Different Rack Configurations

Rack Configuration	SM Should Run On...	SM Priority
Half- or full-rack Exalogic machine connected to a half- or full-rack Exadata machine. **See also**: Run the SM in Configurations with Varying Switch Firmware Versions.	Spine switches	Spine switch: **8**
Quarter-rack Exalogic machine connected to a quarter-rack Exadata machine. **See also**: Run the SM in Configurations with Varying Switch Firmware Versions.	All leaf switches	All leaf switches: **5**
Two or more Exalogic machines connected to two or more Exadata machines. **See also**: Run the SM in Configurations with Varying Switch Firmware Versions.	Spine switches	Spine switch: **8**

13.2.2.1 Run the SM in Configurations with Varying Switch Firmware Versions

In a multirack configuration consisting of both Exalogic and Exadata machines, if firmware upgrades result in switches with varying firmware versions across the configuration, the SM should run on only the switches with the latest firmware version. This is necessary to benefit from the features of the latest firmware.

Note that the SM should run on at least two switches in the fabric.

Consider a configuration that consists of three or more spine switches—for example, two Exalogic machines connected to two Exadata machines—but with varying firmware versions.

- If two or more of the available spine switches, across the configuration, have the highest firmware version, the SM should run on those spine switches, with the priority set to **8**.
- If only *one* of the spine switches in the entire configuration has the highest firmware version:
  - The SM should run on that spine switch. The SM priority should be set to **8**.
  - *In addition*, the SM should run on one or more leaf switches having the latest firmware version. The SM priority of the leaf switches should be set to **5**.

  In this case, running the SM on one or more leaf switches, besides running it on the spine switch, is necessary to fulfill the requirement that at least two SMs should be running in the fabric.

For more information about running the subnet manager, see the following topics:

- Start the Subnet Manager Manually
- Identify the Location of Master Subnet Manager
- Relocate the Master Subnet Manager
- Enable Subnet Manager on a Switch
- Disable Subnet Manager on a Switch

13.2.3 Monitor the Subnet Manager

This section contains the following topics:

- Display the Subnet Manager Status
- Display Recent Subnet Manager Activity

13.2.3.1 Display the Subnet Manager Status

If you want to quickly determine your Subnet Manager's priority and state, you can use the `sminfo` command.

On the command-line interface (CLI), run the following command:

```
# sminfo
```

The output is displayed, as in the following example:

```
sminfo: sm lid 15 sm guid 0x21283a8389a0a0, activity count 32046 priority 8 state3
SMINFO_MASTER
```

In the example output, the Subnet Manager's hosting HCA has LID 15 and GUID 0x21283a8620b0f0. The Subnet Manager has a priority of 8 (high) and its state is 3 (master).

13.2.3.2 Display Recent Subnet Manager Activity

On the command-line interface (CLI), run the following command:

```
# getmaster -l
```

The output is displayed, as in the following example:

```
# getmaster -l
Last ring buffer history listed:
whereismaster-daemon is running
20091204 15:00:53 whereismaster started
20091204 15:00:55 No OpenSM Master seen in the system
20091204 15:06:19 OpenSM Master on Switch : 0x0002c9000100d050 ports 36 Sun DCS
36 QDR switch o4nm2-36p-2.norway.test.com enhanced port 0 lid 7 lmc 0
```

13.2.4 Control the Subnet Manager

This section contains the following topics:

- Identify the Location of Master Subnet Manager
- Relocate the Master Subnet Manager
- Enable Subnet Manager on a Switch
- Disable Subnet Manager on a Switch

13.2.4.1 Identify the Location of Master Subnet Manager

From any InfiniBand switch in the network (leaf switch or spine switch), log in as `root` and run the `getmaster` command to obtain the location of the master SM as follows:

```
# getmaster
```

This command displays the host name or IP address and the IP address of the switch where the master SM is running.

13.2.4.2 Relocate the Master Subnet Manager

You are required to relocate the master SM from a leaf switch (Sun Network QDR InfiniBand Gateway Switch) to the spine switch (Sun Datacenter InfiniBand Switch 36) when you are connecting more than one Exalogic machine. This step is also necessary when you are connecting an Exalogic machine to an Oracle Exadata Database Machine.

Relocating the master SM does not affect the availability of the InfiniBand network. You can perform this task while normal workload is running.

To relocate the master SM from a leaf switch (Sun Network QDR InfiniBand Gateway Switch) to the spine switch (Sun Datacenter InfiniBand Switch 36):

1. Identify the location of the master SM, as described in Identify the Location of Master Subnet Manager.
2. If the master SM is not running on a spine switch, log in as a `root` user to the leaf switch where the master SM is located.
3. Disable SM on the switch, as described in Disable Subnet Manager on a Switch. This step relocates the master SM to another switch in the network.
4. Perform the above steps until the master SM relocates to the spine switch (Sun Datacenter InfiniBand Switch 36).
5. Enable SM on the leaf switches where SM was disabled during this procedure. For information about enabling SM on a switch, see Enable Subnet Manager on a Switch.

13.2.4.3 Enable Subnet Manager on a Switch

To enable SM on a switch:

1. Log in as a `root` user.
2. At the command prompt, run the following command:

   ```
   # enablesm
   ```

13.2.4.4 Disable Subnet Manager on a Switch

To disable SM on a switch:

1. Log in as a `root` user.
2. At the command prompt, run the following command:

   ```
   # disablesm
   ```

13.3 Work with the Default Rack-Level InfiniBand Partition

This section contains the following topics:

- Partition in Exalogic Machine

- Verify the Default Partition

13.3.1 Partition in Exalogic Machine

By default, the Exalogic machine includes a single partition at the rack level. All Exalogic compute nodes and the storage appliance are full members of this default partition.

> **Note:**
>
> Oracle recommends that you create IP subnets over the default IP over InfiniBand (IPoIB) link to implement isolate application deployments in the Exalogic environments. Each IP subnet will have a single multicast domain. When you create IP subnets, ensure that each of the interfaces per Exalogic compute node for these additional IP subnets above the default IPoIB subnet is bonded, for high availability (HA) purposes.
>
> For more information, see the "Application Isolation by Subnetting over IPoIB" topic in the *Oracle Exalogic Enterprise Deployment Guide*.

13.3.2 Verify the Default Partition

You can verify the default partition and the partition key by running the `smpartition list` command on the command-line interface (CLI) for one of the gateway switches.

14
Use the Sun Datacenter InfiniBand Switch 36 in Multirack Configurations

This chapter describes how to set up and configure Sun Datacenter InfiniBand Switch 36, which is used as the spine switch in multirack configurations (an Exalogic machine to another Exalogic machine, or an Exalogic machine to an Oracle Exadata Database Machine) only. This spine switch is not included in Exalogic machine configurations and must be purchased separately.
By using this spine switch, you can connect multiple Exalogic machines or a combination of Exalogic machines and Oracle Exadata Database Machines together on the same InfiniBand fabric.

This chapter contains the following topics:

- Physical Specifications
- Access the CLI of a Sun Datacenter InfiniBand Switch 36
- Verify the Switch Status
- Start the Subnet Manager in Multirack Configuration Scenarios
- Check Link Status
- Verify the InfiniBand Fabric in a Multirack Configuration
- Monitor the Spine Switch Using Web Interface
- What Next?

14.1 Physical Specifications

Table 14-1 provides the physical specifications of the Sun Datacenter InfiniBand Switch 36.

Table 14-1 Sun Datacenter InfiniBand Switch 36 Specifications

Dimension	Measurements
Width	17.52 in. (445.0 mm)
Depth	24 in. (609.6 mm)
Height	1.75 in. (44.5 mm)
Weight	23.0 lbs (11.4 kg)

14.2 Access the CLI of a Sun Datacenter InfiniBand Switch 36

The Sun Datacenter InfiniBand Switch 36 is connected and used in the Exalogic machine in multirack configuration scenarios only. Therefore, you can access the CLI of this switch after connecting the switch in a multirack configuration scenario.

After connecting this switch and applying power, you can access its command-line interface (CLI).

To access the command-line interface (CLI):

1. If you are using a network management port, begin network communication with the command-line interface (CLI) using the **ssh** command and the host name configured with the DHCP server.:

   ```
   % ssh -l root switch-name
   root@switch-name's password: password
   #
   ```

   where `switch-name` is the host name assigned to the Sun Datacenter InfiniBand Switch 36.

   If you do not see this output or prompt, there is a problem with the network communication or cabling of the switch.

2. If you are using a USB management port, begin serial communication with the command-line interface (CLI) as follows:

   a. Connect a serial terminal, terminal server, or workstation with a TIP connection to the USB-to-serial adapter. Configure the terminal or terminal emulator with these settings:

      115200 baud, 8 bits, No parity, 1 Stop bit, and No handshaking

   b. Press the Return or Enter key on the serial device several times to synchronize the connection. You might see text similar to the following:

      ```
      …
      CentOS release 5.2 (Final)
      Kernel 2.6.27.13-nm2 on an i686

      switch-name login: root
      Password: password
      #
      ```

      where `switch-name` is the host name assigned to the Sun Datacenter InfiniBand Switch 36.

      If you do not see this output or prompt, there is a problem with the network communication or the cabling of the switch.

14.3 Verify the Switch Status

For the Sun Datacenter InfiniBand Switch 36, you can check the status of the command-line interface (CLI), power supplies, fans, and switch chip. Verify that the voltage and temperature values of the switch are within specification:

```
# showunhealthy
# env_test
```

An unfavorable output from these commands indicates a hardware fault with that particular component. A voltage or temperature deviating more than 10% from the provided specification means a problem with the respective component.

For example, on the CLI of the switch, enter the following command to check its status:

```
# env_test
```

This command performs a set of checks and displays the overall status of switch, as in the following example:

```
Environment test started:
Starting Voltage test:
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.37 V
Measured 12V = 12.06 V
Measured 5V = 5.03 V
Measured VBAT = 3.25 V
Measured 1.0V = 1.01 V
Measured I4 1.2V = 1.22 V
Measured 2.5V = 2.52 V
Measured V1P2 DIG = 1.17 V
Measured V1P2 AND = 1.16 V
Measured 1.2V BridgeX = 1.21 V
Measured 1.8V = 1.80 V
Measured 1.2V Standby = 1.20 V
Voltage test returned OK
Starting PSU test:
PSU 0 present
PSU 1 present
PSU test returned OK
Starting Temperature test:
Back temperature 23.00
Front temperature 32.62
SP temperature 26.12
Switch temperature 45, maxtemperature 45
Bridge-0 temperature 41, maxtemperature 42
Bridge-1 temperature 43, maxtemperature 44
Temperature test returned OK
Starting FAN test:
Fan 0 not present
Fan 1 running at rpm 11212
Fan 2 running at rpm 11313
Fan 3 running at rpm 11521
Fan 4 not present
FAN test returned OK
Starting Connector test:
Connector test returned OK
Starting onboard ibdevice test:
Switch OK
Bridge-0 OK
Bridge-1 OK
All Internal ibdevices OK
Onboard ibdevice test returned OK
Environment test PASSED
```

When the switch status is operational, you can start the Subnet Manager (SM).

14.4 Start the Subnet Manager in Multirack Configuration Scenarios

The Sun Datacenter InfiniBand Switch 36, which is referred to as the spine switch, is connected and configured in multirack configuration scenarios. You must start the Subnet Manager manually on the switch as follows:

1. On the CLI of the spine switch, run the following command:

   ```
   # enablesm
   ```

2. On the CLI of the spine switch, set the Subnet Manager priority within the command-line interface (CLI) as follows:

   ```
   # setsmpriority priority
   ```

   > **Note:**
   >
   > For information about the switches on which the SM should run in various rack configurations and the SM priorities for the switches, see Subnet Manager Operation in Different Rack Configurations

   For example, to set the Subnet Manager on the spine switch to priority **8**, run the following command on the CLI of the spine switch:

   ```
   # setsmpriority 8
   ```

   The following output is displayed:

   ```
   -------------------------------------------------
   OpenSM 3.2.6_20090717
     Reading Cached Option File: /etc/opensm/opensm.conf
     Loading Cached Option:routing_engine = ftree
     Loading Cached Option:sminfo_polling_timeout = 1000
     Loading Cached Option:polling_retry_number = 3
   Command Line Arguments:
     Priority = 8
     Creating config file template '/tmp/osm.conf'.
     Log File: /var/log/opensm.log
   -------------------------------------------------
   ```

   For the changes to take effect, restart the Subnet Manager as follows:

   ```
   # disablesm
   ```

   ```
   # enablesm
   ```

3. After assigning the SM priority on the spine switch correctly, on the CLI of the gateway switches (Sun Network QDR InfiniBand Gateway Switches referred to as leaf switches in this guide), run the following command to disable Subnet Manager individually on the gateway switches:

   ```
   # disablesm
   ```

14.5 Check Link Status

After starting the Subnet Manager, you can verify that the Link LEDs for cabled links are green. If the Link LED is dark, the link is down. If the Link LED flashes, there are symbol errors.

To check the link status of the cables:

```
# listlinkup
```

If the link for a connector is reported as not present, the link at either end of the cable is down. If a port is down, use the `enableswitchport 0 portnumber` command to bring the port up. Alternatively, use the `ibdevreset` command to reset the switch chip.

See the *Sun Datacenter InfiniBand Switch 36 User's Guide*, "Enable a Switch Chip Port" and "Reset the Switch Chip".

After making sure that the link is up, you can verify the InfiniBand fabric.

The following is example output of the `listlinkup` command:

```
# listlinkup
Connector  0A Present <-> Switch Port 20 up (Enabled)
Connector  1A Present <-> Switch Port 22 up (Enabled)
Connector  2A Present <-> Switch Port 24 up (Enabled)
.
.
.
Connector 15A Not present
Connector 0A-ETH Present
 Bridge-0-1 Port 0A-ETH-1 up (Enabled)
 Bridge-0-1 Port 0A-ETH-2 up (Enabled)
 Bridge-0-0 Port 0A-ETH-3 up (Enabled)
 Bridge-0-0 Port 0A-ETH-4 up (Enabled)
Connector 1A-ETH Present
 Bridge-1-1 Port 1A-ETH-1 up (Enabled)
 Bridge-1-1 Port 1A-ETH-2 up (Enabled)
 Bridge-1-0 Port 1A-ETH-3 up (Enabled)
 Bridge-1-0 Port 1A-ETH-4 up (Enabled)
Connector 0B Present <-> Switch Port 19 up (Enabled)
Connector 1B Present <-> Switch Port 21 up (Enabled)
.
.
.
Connector 15B Not present
#
```

14.6 Verify the InfiniBand Fabric in a Multirack Configuration

Use the following commands on the command-line interface (CLI) of the spine switch to verify that the InfiniBand fabric in your multirack configuration is operational:

1. `ibnetdiscover`

   Discovers and displays the InfiniBand fabric topology and connections. See Discover the InfiniBand Network Topology in a Multirack Configuration.

2. `ibdiagnet`

Performs diagnostics upon the InfiniBand fabric and reports status. See Perform Diagnostics on the InfiniBand Fabric in a Multirack Configuration.

3. `ibcheckerrors`

   Checks the entire InfiniBand fabric for errors. See Check for Errors in the InfiniBand Fabric in a Multirack Configuration.

14.6.1 Discover the InfiniBand Network Topology in a Multirack Configuration

To discover the InfiniBand network topology and build a topology file which is used by the OpenSM Subnet Manager, run the following command on the command-line interface (CLI) of the spine switch:

```
# ibnetdiscover
```

The output is displayed, as in the following example:

The topology file is used by InfiniBand commands to scan the InfiniBand fabric and validate the connectivity as described in the topology file, and to report errors as indicated by the port counters.

```
# Topology file: generated on Sat Apr 13 22:28:55 2002
#
# Max of 1 hops discovered
# Initiated from node 0021283a8389a0a0 port 0021283a8389a0a0
vendid=0x2c9
devid=0xbd36
sysimgguid=0x21283a8389a0a3
switchguid=0x21283a8389a0a0(21283a8389a0a0)
Switch   36 "S-0021283a8389a0a0" # "Sun DCS 36 QDR switch localhost" enhanced port 0
lid 15 lmc 0
[23]    "H-0003ba000100e388"[2](3ba000100e38a) # "nsn33-43 HCA-1" lid 14 4xQDR
vendid=0x2c9
devid=0x673c
sysimgguid=0x3ba000100e38b
caguid=0x3ba000100e388
Ca   2 "H-0003ba000100e388" # "nsn33-43 HCA-1"
[2](3ba000100e38a)   "S-0021283a8389a0a0"[23] # lid 14 lmc 0 "Sun DCS 36 QDR switch
localhost" lid 15 4xQDR
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

14.6.2 Perform Diagnostics on the InfiniBand Fabric in a Multirack Configuration

To perform a collection of tests on the InfiniBand fabric and generate several files that contain parameters and aspects of the InfiniBand fabric, run the following command on the command-line interface (CLI) of the spine switch:

```
# ibdiagnet
```

In the following example, the `ibdiagnet` command is minimized to determine which links are underperforming:

```
# ibdiagnet -lw 4x -ls 10 -skip all

Loading IBDIAGNET from: /usr/lib/ibdiagnet1.2
-W- Topology file is not specified.
 Reports regarding cluster links will use direct routes.
Loading IBDM from: /usr/lib/ibdm1.2
-I- Using port 0 as the local port.
-I- Discovering ... 2 nodes (1 Switches & 1 CA-s) discovered.
.
.
.
-I- Links With links width != 4x (as set by -lw option)
-I---------------------------------------------------
-I- No unmatched Links (with width != 4x) were found
-I---------------------------------------------------
-I- Links With links speed != 10 (as set by -ls option)
-I---------------------------------------------------
-I- No unmatched Links (with speed != 10) were found
.
.
.
-I- Stages Status Report:
 STAGE                Errors Warnings
 Bad GUIDs/LIDs Check               0   0
 Link State Active Check               0   0
 Performance Counters Report               0   0
 Specific Link Width Check               0   0
 Specific Link Speed Check               0   0
 Partitions Check               0   0
 IPoIB Subnets Check               0   0
Please see /tmp/ibdiagnet.log for complete log
----------------------------------------------------------------
-I- Done. Run time was 1 seconds.
```

Note:

The actual output for your InfiniBand fabric will differ from that in the example.

14.6.3 Check for Errors in the InfiniBand Fabric in a Multirack Configuration

Use the `ibcheckerrors` command that uses the topology file to scan the InfiniBand fabric and validate the connectivity as described in the topology file, and to report errors as indicated by the port counters.

On the command-line interface (CLI) of the spine switch, enter the following command:

```
# ibcheckerrors

## Summary: 4 nodes checked, 0 bad nodes found
##      34 ports checked, 0 ports have errors beyond threshold
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

14.7 Monitor the Spine Switch Using Web Interface

1. Open a web browser and go to the following URL:

   ```
   http://switch-IP
   ```

   where `switch-IP` is the IP address of the spine switch.
2. Log in to the interface as the `root` user.
3. Click the **Switch/Fabric Monitoring Tools** tab.
4. Click **Launch Sun DCS GW Monitor**.

   The Fabric Monitor is displayed.

14.8 What Next?

After setting up the Sun Datacenter InfiniBand Switch 36 in a multirack configuration scenario, you can proceed to monitor and control the InfiniBand fabric.

15 Monitor and Control the InfiniBand Fabric

This chapter describes how to monitor and control the InfiniBand fabric.
It contains the following topics:

- Monitor the InfiniBand Fabric
- Control the InfiniBand Fabric
- For More Information

15.1 Monitor the InfiniBand Fabric

This section contains the following topics:

- Identify All Switches in the Fabric
- Identify All HCAs in the Fabric
- Display the InfiniBand Fabric Topology
- Display a Route Through the Fabric
- Display the Link Status of a Node
- Display Counters for a Node
- Display Data Counters for a Node
- Display Low-Level Detailed Information for a Node
- Display Low-Level Detailed Information for a Port
- Map LIDs to GUIDs
- Perform Comprehensive Diagnostics for the Entire Fabric
- Perform Comprehensive Diagnostics for a Route
- Determine Changes to the InfiniBand Topology
- Determine Which Links Are Experiencing Significant Errors
- Check All Ports

15.1.1 Identify All Switches in the Fabric

You can use the `ibswitches` command to identify the Sun Network QDR InfiniBand Gateway Switches in the InfiniBand fabric in your Exalogic machine. This command displays the Global Unique Identifier (GUID), name, Local Identifier (LID), and LID mask control (LMC) for each switch. The output of the command is a mapping of GUID to LID for switches in the fabric.

On any command-line interface (CLI), run the following command:

```
# ibswitches
```

The output is displayed, as in the following example:

```
Switch : 0x0021283a8389a0a0 ports 36 "Sun DCS 36 QDR switch localhost" enhancedport 0
 lid 15 lmc 0
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.2 Identify All HCAs in the Fabric

You can use the `ibhosts` command to display identity information about the host channel adapters (HCAs) in the InfiniBand fabric in a subnet. This command displays the GUID and name for each HCA.

On the command-line interface (CLI), run the following command:

```
# ibhosts
```

The output is displayed, as in the following example:

```
Ca : 0x0003ba000100e388 ports 2 "nsn33-43 HCA-1"
Ca : 0x5080020000911310 ports 1 "nsn32-20 HCA-1"
Ca : 0x50800200008e532c ports 1 "ib-71 HCA-1"
Ca : 0x50800200008e5328 ports 1 "ib-70 HCA-1"
Ca : 0x50800200008296a4 ports 2 "ib-90 HCA-1"
.
.
.
#
```

> **Note:**
>
> The output in the example is just a portion of the full output and varies for each InfiniBand topology.

15.1.3 Display the InfiniBand Fabric Topology

To understand the routing that happens within your InfiniBand fabric, the `ibnetdiscover` command displays the node-to-node connectivity. The output of the command is dependent upon the size of your fabric. You can also use this command to display the LIDs of HCAs.

On the command-line interface (CLI), enter the following command:

```
# ibnetdiscover
```

The output is displayed, as in the following example:

```
# Topology file: generated on Sat Apr 13 22:28:55 2002
#
# Max of 1 hops discovered
# Initiated from node 0021283a8389a0a0 port 0021283a8389a0a0
vendid=0x2c9
devid=0xbd36
```

```
sysimgguid=0x21283a8389a0a3
switchguid=0x21283a8389a0a0(21283a8389a0a0)
Switch   36 "S-0021283a8389a0a0" # "Sun DCS 36 QDR switch localhost" enhanced port 0
lid 15 lmc 0
[23]    "H-0003ba000100e388"[2](3ba000100e38a) # "nsn33-43 HCA-1" lid 14 4xQDR
vendid=0x2c9
devid=0x673c
sysimgguid=0x3ba000100e38b
caguid=0x3ba000100e388
Ca   2 "H-0003ba000100e388" # "nsn33-43 HCA-1"
[2](3ba000100e38a)   "S-0021283a8389a0a0"[23] # lid 14 lmc 0 "Sun DCS 36 QDR switch
localhost" lid 15 4xQDR
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.4 Display a Route Through the Fabric

You sometimes need to know the route between two nodes in the InfiniBand fabric. The `ibtracert` command can provide that information by displaying the GUIDs, ports, and LIDs of the nodes.On the command-line interface (CLI), run the following command:

```
# ibtracert slid dlid
```

where `slid` is the LID of the source node and `dlid` is the LID of the destination node in the fabric.

The output is displayed, as in the following example:

```
# ibtracert 15 14
#
From switch {0x0021283a8389a0a0} portnum 0 lid 15-15 "Sun DCS 36 QDR switch localhost
"
[23] -> ca port {0x0003ba000100e38a}[2] lid 14-14 "nsn33-43 HCA-1"
To ca {0x0003ba000100e388} portnum 2 lid 14-14 "nsn33-43 HCA-1"
#
```

For this example:

The route starts at switch with GUID `0x0021283a8389a0a0` and is using port `0`. The switch is `LID 15` and in the description, the switch host's name is `Sun DCS 36 QDR switch localhost`. The route enters at port `23` of the HCA with GUID `0x0003ba000100e38a` and exits at port `2`. The HCA is `LID 14`.

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.5 Display the Link Status of a Node

If you want to know the link status of a node in the InfiniBand fabric, run the `ibportstate` command to display the state, width, and speed of that node:

On the command-line interface (CLI), run the following command:

```
# ibportstate lid port
```

where `lid` is the LID of the node in the fabric, `port` is the port of the node.

The output is displayed, as in the following example:

```
# ibportstate 15 23

PortInfo:
# Port info: Lid 15 port 23
LinkState:.......................Active
PhysLinkState:...................LinkUp
LinkWidthSupported:..............1X or 4X
LinkWidthEnabled:................1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.................10.0 Gbps
Peer PortInfo:
# Port info: Lid 15 DR path slid 15; dlid 65535; 0,23
LinkState:.......................Active
PhysLinkState:...................LinkUp
LinkWidthSupported:..............1X or 4X
LinkWidthEnabled:................1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.................10.0 Gbps
#
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.6 Display Counters for a Node

To help ascertain the health of a node in the fabric, use the `perfquery` command to display the performance, error, and data counters for that node:

On the command-line interface (CLI), enter the following command:

```
# perfquery lid port
```

where `lid` is the LID of the node in the fabric, and `port` is the port of the node.

> **Note:**
>
> If a port value of 255 is specified for a switch node, the counters are the total for all switch ports.

For example:

```
# perfquery 15 23
#
# Port counters: Lid 15 port 23
PortSelect:......................23
CounterSelect:...................0x1b01
SymbolErrors:....................0
.
.
.
VL15Dropped:.....................0
XmtData:.........................20232
RcvData:.........................20232
XmtPkts:.........................281
RcvPkts:.........................281
```

> **Note:**
>
> The output in the example is just a portion of the full output.

15.1.7 Display Data Counters for a Node

To list the data counters for a node in the fabric, use the `ibdatacounts` command.

On the command-line interface (CLI), enter the following command:

```
# ibdatacounts lid port
```

where `lid` is the LID of the node in the fabric, and `port` is the port of the node.

For example:

```
# ibdatacounts 15 23
#
XmtData:.........................6048
RcvData:.........................6048
XmtPkts:.........................84
RcvPkts:.........................84
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.8 Display Low-Level Detailed Information for a Node

If intensive troubleshooting is necessary to resolve a problem, you can use the `smpquery` command to display very detailed information about a node in the fabric.

On the command-line interface (CLI), enter the following command:

```
# smpquery switchinfo lid
```

where `lid` is the LID of the node in the fabric.

For example:

```
# smpquery switchinfo 15
#
# Switch info: Lid 15
LinearFdbCap:....................49152
RandomFdbCap:....................0
McastFdbCap:.....................4096
LinearFdbTop:....................16
DefPort:.........................0
DefMcastPrimPort:................255
DefMcastNotPrimPort:.............255
LifeTime:........................18
StateChange:.....................0
LidsPerPort:.....................0
PartEnforceCap:..................32
InboundPartEnf:..................1
OutboundPartEnf:.................1
FilterRawInbound:................1
FilterRawOutbound:...............1
EnhancedPort0:...................1
#
# smpquery portinfo lid port
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example.

15.1.9 Display Low-Level Detailed Information for a Port

If intensive troubleshooting is necessary to resolve a problem, you can use the `smpquery` command to display very detailed information about a port.

On the command-line interface (CLI), enter the following command:

```
# smpquery portinfo lid port
```

where `lid` is the LID of the node in the fabric.

For example:

```
# smpquery portinfo 15 23
#
Mkey:............................0x0000000000000000
GidPrefix:.......................0x0000000000000000
```

```
Lid:.............................0x0000
SMLid:...........................0x0000
CapMask:.........................0x0
DiagCode:........................0x0000
MkeyLeasePeriod:.................0
LocalPort:.......................0
LinkWidthEnabled:................1X or 4X
LinkWidthSupported:..............1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:.......................Active
PhysLinkState:...................LinkUp
LinkDownDefState:................Polling
ProtectBits:.....................0
LMC:.............................0
.
.
.
SubnetTimeout:...................0
RespTimeVal:.....................0
LocalPhysErr:....................8
OverrunErr:......................8
MaxCreditHint:...................85
RoundTrip:.......................16777215
#
```

> **Note:**
>
> The actual output for your InfiniBand fabric will differ from that in the example, and it is just a portion of the full output.

15.1.10 Map LIDs to GUIDs

In the InfiniBand fabric in Exalogic machines, as a Subnet Manager and Subnet administrator, you may want to assign subnet-specific LIDs to nodes in the fabric. Often in the use of the InfiniBand commands, you must provide an LID to issue a command to a particular InfiniBand device.

Alternatively, the output of a command might identify InfiniBand devices by their LID. You can create a file that is a mapping of node LIDs to node GUIDs, which can help with administrating your InfiniBand fabric.

> **Note:**
>
> Creation of the mapping file is not a requirement for InfiniBand administration.

The following procedure creates a file that lists the LID in hexadecimal, the GUID in hexadecimal, and the node description:

1. Create an inventory file:

   ```
   # osmtest -f c -i inventory.txt
   ```

The `inventory.txt` file can be used for other purposes too, besides this procedure.

2. Create a mapping file:

```
# cat inventory.txt |grep -e '^lid' -e 'port_guid' -e 'desc' |sed 's/^lid/
\nlid/'> mapping.txt
```

3. Edit the latter half of the `mapping.txt` file to remove the nonessential information. The content of the `mapping.txt` file looks similar to the following:

```
lid 0x14
port_guid 0x0021283a8620b0a0
# node_desc Sun DCS 72 QDR switch 1.2(LC)
lid 0x15
port_guid 0x0021283a8620b0b0
# node_desc Sun DCS 72 QDR switch 1.2(LC)
lid 0x16
port_guid 0x0021283a8620b0c0
# node_desc Sun DCS 72 QDR switch 1.2(LC)
```

Note:

The output in the example is just a portion of the entire file.

15.1.11 Perform Comprehensive Diagnostics for the Entire Fabric

If you require a full testing of your InfiniBand fabric, you can use the `ibdiagnet` command to perform many tests with verbose results. The command is a useful tool to determine the general overall health of the InfiniBand fabric.

On the command-line interface (CLI), run the following command:

```
# ibdiagnet -v -r
```

The `ibdiagnet.log` file contains the log of the testing.

15.1.12 Perform Comprehensive Diagnostics for a Route

You can use the `ibdiagpath` command to perform some of the same comprehensive tests for a particular route.

On the command-line interface (CLI), run the following command:

```
# ibdiagpath -v -l slid dlid
```

where `slid` is the LID of the source node in the fabric, and `dlid` is the LID of the destination node.

The `ibdiagpath.log` file contains the log of the testing.

15.1.13 Determine Changes to the InfiniBand Topology

If your fabric has a number of nodes that are suspect, the `osmtest` command enables you to take a snapshot (inventory file) of your fabric and at a later time compare that file to the present conditions.

> **Note:**
>
> Although this procedure is most useful after initializing the Subnet Manager, it can be performed at any time.

Complete the following steps:

1. Ensure that Subnet Manager is initiated.
2. On the command-line interface (CLI), run the following command to take a snapshot of the topology:

   ```
   # osmtest -f c
   ```

   For example:

   ```
   # osmtest -f c
   Command Line Arguments
   Done with args
   Flow = Create Inventory
   Aug 13 19:44:53 601222 [B7D466C0] 0x7f -> Setting log level to: 0x03
   Aug 13 19:44:53 601969 [B7D466C0] 0x02 -> osm_vendor_init: 1000 pending
   umadsspecified
   using default guid 0x21283a8620b0f0
   Aug 13 19:44:53 612312 [B7D466C0] 0x02 -> osm_vendor_bind: Binding to
   port0x21283a8620b0f0
   Aug 13 19:44:53 636876 [B7D466C0] 0x02 -> osmtest_validate_sa_class_port_info:
   -----------------------------
   SA Class Port Info:
   base_ver:1
   class_ver:2
   cap_mask:0x2602
   cap_mask2:0x0
   resp_time_val:0x10
   -----------------------------
   OSMTEST: TEST "Create Inventory" PASS
   #
   ```

3. After an event, compare the present topology to that saved in the inventory file, as in the following example:

   ```
   # osmtest -f v
   Command Line Arguments
   Done with args
   Flow = Validate Inventory
   Aug 13 19:45:02 342143 [B7EF96C0] 0x7f -> Setting log level to: 0x03
   Aug 13 19:45:02 342857 [B7EF96C0] 0x02 -> osm_vendor_init: 1000 pending
   umadsspecified
   using default guid 0x21283a8620b0f0
   Aug 13 19:45:02 351555 [B7EF96C0] 0x02 -> osm_vendor_bind: Binding to
   port0x21283a8620b0f0
   Aug 13 19:45:02 375997 [B7EF96C0] 0x02 -> osmtest_validate_sa_class_port_info:
   -----------------------------
   SA Class Port Info:
   base_ver:1
   class_ver:2
   cap_mask:0x2602
   cap_mask2:0x0
   resp_time_val:0x10
   -----------------------------
   ```

```
Aug 13 19:45:02 378991 [B7EF96C0] 0x01 -> osmtest_validate_node_data:
Checkingnode 0x0021283a8620b0a0, LID 0x14
Aug 13 19:45:02 379172 [B7EF96C0] 0x01 -> osmtest_validate_node_data:
Checkingnode 0x0021283a8620b0b0, LID 0x15
.
.
.
Aug 13 19:45:02 480201 [B7EF96C0] 0x01 -
>osmtest_validate_single_path_rec_guid_pair:
Checking src 0x0021283a8620b0f0 to dest 0x0021283a8620b0f0
Aug 13 19:45:02 480588 [B7EF96C0] 0x01 -> osmtest_validate_path_data:
Checkingpath SLID 0x19 to DLID 0x19
Aug 13 19:45:02 480989 [B7EF96C0] 0x02 -> osmtest_run:
***************** ALL TESTS PASS *****************
OSMTEST: TEST "Validate Inventory" PASS
#
```

> **Note:**
>
> Depending on the size of your InfiniBand fabric, the output from the `osmtest` command could be tens of thousands of lines long.

15.1.14 Determine Which Links Are Experiencing Significant Errors

You can use the `ibdiagnet` command to determine which links are experiencing symbol errors and recovery errors by injecting packets.

On the command-line interface (CLI), run the following command:

```
# ibdiagnet -c 100 -P all=1
```

In this instance of the `ibdiagnet` command, 100 test packets are injected into each link and the `-P all=1` option returns all counters that increment during the test.

In the output of the `ibdiagnet` command, search for the `symbol_error_counter` string. That line contains the symbol error count in hexadecimal. The preceding lines identify the node and port with the errors. Symbol errors are minor errors, and if there are relatively few during the diagnostic, they can be monitored.

> **Note:**
>
> According to the InfiniBand specification 10E-12 BER, the maximum allowable symbol error rate is 120 errors per hour.

In addition, in the output of the `ibdiagnet` command, search for the `link_error_recovery_counter` string.

That line contains the recovery error count in hexadecimal. The preceding lines identify the node and port with the errors. Recovery errors are major errors and the respective links must be investigated for the cause of the rapid symbol error propagation.

Additionally, the `ibdiagnet.log` file contains the log of the testing.

15.1.15 Check All Ports

To perform a quick check of all ports of all nodes in your InfiniBand fabric, you can use the `ibcheckstate` command.

On the command-line interface (CLI), run the following command:

```
# ibcheckstate -v
```

The output is displayed, as in the following example:

```
# Checking Switch: nodeguid 0x0021283a8389a0a0
Node check lid 15: OK
Port check lid 15 port 23: OK
Port check lid 15 port 19: OK
.
.
.
# Checking Ca: nodeguid 0x0003ba000100e388
Node check lid 14: OK
Port check lid 14 port 2: OK
## Summary: 5 nodes checked, 0 bad nodes found
## 10 ports checked, 0 ports with bad state found
#
```

> **Note:**
>
> The `ibcheckstate` command requires time to complete, depending upon the size of your InfiniBand fabric. Without the `-v` option, the output contains only failed ports. The output in the example is only a small portion of the actual output.

15.2 Control the InfiniBand Fabric

This section contains the following topics:

- Clear Error Counters
- Clear Data Counters
- Reset a Port
- Set Port Speed
- Disable a Port
- Enable a Port

15.2.1 Clear Error Counters

If you are troubleshooting a port, the `perfquery` command provides counters of errors occurring at that port. To determine if the problem has been resolved, you can reset all of the error counters to 0 with the `ibclearerrors` command.

On the command-line interface (CLI), run the following command:

```
# ibclearerrors
```

The output is displayed, as in the following example:

```
## Summary: 5 nodes cleared 0 errors
#
```

15.2.2 Clear Data Counters

When you are optimizing the InfiniBand fabric for performance, you might want to know how the throughput increases or decreases according to changes you are making to the fabric and Subnet Manager. The `ibclearcounters` command enables you to reset the data counters for all ports to 0.

On the command-line interface (CLI), run the following command:

```
# ibclearcounters
```

The output is displayed, as in the following example:

```
## Summary: 5 nodes cleared 0 errors
#
```

15.2.3 Reset a Port

You might need to reset a port to determine its functionality.

On the command-line interface (CLI), run the following command:

```
# ibportstate lid port reset
```

where `lid` is the LID of the node in the fabric, and `port` is the port of the node.

For example:

```
# ibportstate 15 23 reset
Initial PortInfo:
# Port info: Lid 15 port 23
LinkState:.......................Down
PhysLinkState:...................Disabled
LinkWidthSupported:..............1X or 4X
LinkWidthEnabled:................1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.................2.5 Gbps
After PortInfo set:
# Port info: Lid 15 port 23
LinkState:.......................Down
PhysLinkState:...................Disabled
After PortInfo set:
# Port info: Lid 15 port 23

LinkState:.......................Down
PhysLinkState:...................PortConfigurationTraining
#
```

15.2.4 Set Port Speed

You can manually set the speed of a single port to help determine symbol error generation. The `ibportstate` command can set the speed to 2.5, 5.0, or 10.0 GB/sec.

On the command-line interface (CLI), run the following command:

```
# ibportstate lid port speed <value>
```

where `lid` is the LID of the node in the fabric, `port` is the port of the node, and `<value>` is the speed you want to set.

> **Note:**
>
> Adding speed values enables either speed. For example, speed 7 is 2.5, 5.0, and 10.0 GB/sec.

For example:

```
# ibportstate 15 23 speed 1
Initial PortInfo:
# Port info: Lid 15 port 23
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
After PortInfo set:
# Port info: Lid 15 port 23
LinkSpeedEnabled:................2.5 Gbps
# ibportstate 15 23 speed 7
Initial PortInfo:
# Port info: Lid 15 port 23
LinkSpeedEnabled:................2.5 Gbps
After PortInfo set:
# Port info: Lid 15 port 23
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
#
```

15.2.5 Disable a Port

If a port is found to be problematic due to a bad cable connection or a physical damage to the connectors, you can disable the port.

On the command-line interface (CLI), run the following command:

```
# disableswitchport [--reason=reason] connector|ibdev port
```

where `reason` is the reason for disabling the port, `Blacklist` or `Partition`. `connector` is the number of the QSFP connector (0A–15B). `ibdev` is the InfiniBand device name (Switch, Bridge-0-0, Bridge-0-1, Bridge-1-0, Bridge-1-1). `port` is the number of the port (1–36).

This hardware command disables a QSFP connector and port on the switch chip or a port on the BridgeX chips. The command addresses either the connector or the port on the switch chip or the BridgeX port.

The `--reason` option enables you to use a passphrase to lock the state of the port:

- `Blacklist` – A connector and port pair are identified as being inaccessible because of unreliable operation.
- `Partition` – A connector and port pair are identified as being isolated from the InfiniBand fabric.

Both the `Blacklist` and `Partition` passphrases survive reboot. You unlock these passphrases using the `enableswitchport` command with the `--reason` option.

> **Note:**
>
> State changes made with the `ibportstate` command are not recognized by the `disableswitchport`, `enableswitchport`, or `listlinkup` commands.

The following example shows how to disable and blacklist connector 14A with the `disableswitchport` command.:

```
# disableswitchport --reason=Blacklist 14A
Disable Switch port 7 reason: Blacklist
Initial PortInfo:
# Port info: DR path slid 65535; dlid 65535; 0 port 7
LinkState:.......................Down
PhysLinkState:...................Polling
LinkWidthSupported:..............1X or 4X
LinkWidthEnabled:................1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.................2.5 Gbps
After PortInfo set:
# Port info: DR path slid 65535; dlid 65535; 0 port 7
LinkState:.......................Down
PhysLinkState:...................Disabled
#
```

> **Note:**
>
> After fixing the cable connection or any connector problems, you should enable the port.

15.2.6 Enable a Port

After fixing any connection- or connector-related problem related to a port, you should enable the port with the `enableswitchport` command.

On the command-line interface (CLI), run the following command:

```
enableswitchport [--reason=reason] connector|ibdev port
```

where `reason` is the reason for disabling the port, `connector` is the number of the QSFP connector (0A–15B), `ibdev` is the InfiniBand device name (Switch, Bridge-0-0, Bridge-0-1, Bridge-1-0, Bridge-1-1), and `port` is the number of the port (1–36).

For example:

```
# enableswitchport --reason=Blacklist 14A
Enable Switch port 7
Initial PortInfo:
# Port info: DR path slid 65535; dlid 65535; 0 port 7
```

```
LinkState:.......................Down
PhysLinkState:...................Disabled
LinkWidthSupported:..............1X or 4X
LinkWidthEnabled:................1X or 4X
LinkWidthActive:.................4X
LinkSpeedSupported:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.................2.5 Gbps
After PortInfo set:
# Port info: DR path slid 65535; dlid 65535; 0 port 7
LinkState:.......................Down
PhysLinkState:...................Polling
#
```

15.3 For More Information

For more information about Sun Network QDR InfiniBand Gateway Switches, see the product documentation at the following URL:

http://download.oracle.com/docs/cd/E19671-01/index.html

16 Use InfiniBand Partitions in Exalogic Physical Environments

This chapter describes how to use InfiniBand partitions for network isolation on Exalogic's InfiniBand fabric in the Exalogic physical environment.

> **Note:**
>
> If you are connecting your Exalogic machine to Oracle Exadata Database Machine on the same InfiniBand fabric, you must use the default partition for data traffic between Exalogic machine and Oracle Exadata Database Machine. In this scenario, if you wish to implement network isolation, you can configure IP subnets on the default IPoIB network interface.
>
> In addition, see Important Notes for Combined Exalogic-Exadata Fabric Users for more information about using partitions in this scenario.

This chapter contains the following sections:

16.1 Overview of Partitioning

An InfiniBand partition defines a group of InfiniBand nodes that are allowed to communicate with one another. You can use InfiniBand partitions to increase security by implementing network isolation on Exalogic machine's InfiniBand fabric. In addition, you can associate InfiniBand nodes with specific VLANs.

An InfiniBand node can be a member of multiple partitions. When a packet arrives at a compute node, the partition key (*pkey*) of the packet is matched with the Subnet Manager configuration. This validation prevents a compute node from communicating with another compute node outside its partition.

Based on your requirements, you can create additional partitions as follows:

- Create a unique partition for Exalogic's private InfiniBand fabric by setting nondefault partition keys.

  This scenario applies to both single rack and multiple Exalogic racks.
- Create Virtual LANs (VLANs) on the client access network for EoIB configuration by specifying nondefault partition keys.

  VLAN tagging for a virtual network interface (VNIC) on the EoIB network is optional.

16.2 Understand Partition Keys

A partition key (`pkey`) is a unique ID assigned to an InfiniBand partition. The pkey of the default partition is `0x7fff`. When a pkey is created, it is a 15-bit number. After the membership type is set, the `pkey` value becomes a 16-bit number. The Most Significant Bit (MSB) of the 16-bit `pkey` value denotes the membership type. A limited member has a value of 0, and a full member has a value of 1.

A full member can communicate with both full and limited members of the partition. However, a limited member can only communicate with a full member.

When assigning a `pkey` value for a unique, nondefault partition, you should select a 15-bit value. For example, `0x1234` with values from `0x0001` to `0x7fff`. A total of 32767 pkeys are available. Do not assign `pkey`s that differ only in the MSB of their 16-bit numbers (for example, `0x8005` and `0x0005`).

16.2.1 Guidelines for Managing pkey Allocation in a Hybrid Rack

The term *hybrid rack* denotes an Exalogic machine on which half the compute nodes are in a physical configuration and the other half constitutes a virtualized data center. For more information about hybrid racks, see the *Exalogic Elastic Cloud Release Notes*.

On a hybrid Exalogic rack, Exalogic Control ensures that a unique pkey is assigned to each partition in the virtual environment. However, in the physical half of the rack, pkeys continue to be assigned manually, typically by the network administrator. The following guidelines will help ensure that the pkeys assigned manually to the partitions created in the physical half of a hybrid rack are different from those that Exalogic Control assigns automatically to partitions created in the virtual half of the rack.

1. Make a list (say, **L1**) of all the pkeys assigned to partitions that were created before the rack was converted to a hybrid configuration.

   This set includes the pkey for the IPoIB-default partition (0x7fff) and pkeys for any nondefault partitions that were created in the physical configuration. Note that these pkeys are not guaranteed to be sequential, because they are assigned manually by administrators, who may be using different conventions for assigning pkeys to partitions. For example, for EoIB partitions, some administrators may follow the convention of assigning pkey values that match the VLAN IDs used for the EoIB networks.
2. Identify a list (say, **L2**) of pkeys to be assigned to partitions created in the physical half of the hybrid rack.

Select a list that is preferably near the upper end of the 0x0001–0x7ffe range. For example, if you identify 0x7000 to 0x7ffe as you range, you can create up to 4096 partitions. We recommend the upper end, because Exalogic Control assigns pkeys starting from lower end—that is, 0x0001.

3. In this list (**L2**), mark or remove the pkeys that were assigned before the rack was converted to a hybrid configuration—that is, the **L1** list you created earlier.
4. As you create partitions, select pkeys from only the predetermined list (**L2**) and keep track of the pkeys that you are assigning.

   Such an approach provides a reasonable guarantee that pkeys assigned in the physical environment are different from the pkeys that Exalogic Control assigns in the virtual half of the rack.

After the Exalogic machine is converted to a hybrid rack, for every network (either IPoIB or EoIB) that you create in the virtual half of the hybrid rack, Exalogic Control automatically assigns a unique pkey, starting from 0x0001. While selecting an unused pkey for a new partition, Exalogic Control will skip any pkeys (both **L1** and **L2**) that are used for partitions in the physical half of the configuration. This way, every partition on the hybrid rack—regardless of whether it is on the physical or virtual part—will have a unique pkey.

16.3 Before You Begin

Before you can start creating unique InfiniBand partitions, you must complete the following tasks:

1. Verify the switch firmware version
2. Gather the port GUIDs of compute nodes and BridgeX ports of gateway switches
3. Identify the InfiniBand switches in your Exalogic machine's InfiniBand fabric and note down their IP addresses
4. Determine which InfiniBand switch is running the master Subnet Manager (SM)
5. Log in to the InfiniBand switch that is running the master Subnet Manager (SM)

16.3.1 Verify the Firmware of InfiniBand Switch

Ensure that the InfiniBand switches in your Exalogic machine are installed with firmware versions 2.0.4 or above. This requirement is mandatory.

16.3.2 Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches

Before creating an InfiniBand partition, you must identify the port GUIDs of Exalogic compute nodes that will be added to the partition. In addition, you must identify the BridgeX ports of the gateway switches that are connected to those Exalogic compute nodes.

Identifying Port GUIDs on Compute Nodes

To identify the port GUIDs on an Exalogic compute node, run the following command on the command line:

```
# ibstat
```

This command displays output, as in the following example:

```
CA 'mlx4_0'
     CA type: MT26428
     Number of ports: 2
     Firmware version: 2.9.1000
     Hardware version: b0
     Node GUID: 0x0021280001cef972
     System image GUID: 0x0021280001cef975
     Port 1:
         State: Active
         Physical state: LinkUp
         Rate: 40
         Base lid: 36
         LMC: 0
         SM lid: 5
         Capability mask: 0x02510868
         Port GUID: 0x0021280001cef973
         Link layer: IB
     Port 2:
         State: Active
         Physical state: LinkUp
         Rate: 40
         Base lid: 37
         LMC: 0
         SM lid: 5
         Capability mask: 0x02510868
         Port GUID: 0x0021280001cef974
         Link layer: IB
```

In the above example, Port GUID values are highlighted in a rectangle for illustration purposes only. The actual command does not highlight Port GUID. You must see the command output and note down the values for both InfiniBand ports on each compute node.

Alternatively, you can run the following command to display only GUIDs:

```
# ibstat | grep 'Port GUID:'
```

16.3.2.1 Identify BridgeX Ports on Gateway Switches

To identify the BridgeX ports on the gateway switches that are connected to your compute nodes, run the following command at the command prompt on each gateway switch that your compute node is connected to:

```
# showgwports
```

This command displays the BridgeX ports. Note down the values in the INTERNAL PORTS section of the output, as in the following example:

```
INTERNAL PORTS:
---------------
Device Port Portname PeerPort PortGUID LID IBState GWState
--------------------------------------------------------------------
Bridge-0 1 Bridge-0-1 4 0x002128548062c001 0x0015 Active Up
Bridge-0 2 Bridge-0-2 3 0x002128548062c002 0x000d Active Up
Bridge-1 1 Bridge-1-1 2 0x002128548062c041 0x000f Active Up
Bridge-1 2 Bridge-1-2 1 0x002128548062c042 0x0010 Active Up
```

> **Tip:**
>
> In an Exalogic machine full rack, compute nodes 1 to 15 (start from the bottom of the rack) connect their InfiniBand port 1 to gateway switch1 and their InfiniBand port2 to gateway switch 2. Similarly, compute nodes 16 to 30 are connected to gateway switches 3 and 4.

16.3.3 Identify All InfiniBand Switches in the Fabric

To identify all InfiniBand switches (Sun Network InfiniBand Gateway Switch or Sun Datacenter InfiniBand Switch 36) running master or standby instances of Subnet Manager (SM) on the fabric, run the following command on any of the InfiniBand switches:

```
# ibswitches
```

This command displays the GUID, name, LID, and LMC for each switch. The output of the command is a mapping of GUID to LID for switches in the fabric.

16.3.4 Determine the SM Priority on an InfiniBand Switch

After identifying the InfiniBand switches and their IP addresses, you must log in to each of the switches and run the following command to identify the InfiniBand switch where the master Subnet Manager (SM) is running:

```
# getmaster
```

This command displays output, as shown in the following example:

```
Local SM enabled and running
20111122 08:45:02 Master SubnetManager on sm lid 11 sm guid 0x21283bad45c0a0 : SUN
IB QDR GW switch el01gw04 10.10.10.10
```

16.3.5 Log In to the InfiniBand Switch That Runs Master SM

After identifying the InfiniBand switch where master SM is running, log in to the ILOM shell for the InfiniBand switch as the ILOM administrator (`ilom-admin`). After logging in, run the `show /SYS/Fabric_Mgmt` command to log in to the restricted Linux shell. To view a list of available commands, you can run the `help all` command.

16.4 Move from a Default Partition to a Custom Partition

Moving from a configuration that does not use InfiniBand partitions (that is, uses the default partition only) to a configuration with partitions involves the following steps:

- Making all Exalogic compute nodes limited members of the default partition

> **Note:**
>
> By default, all Exalogic compute nodes are full members of the default partition.

- Disabling IPoIB on the default partition

> **Note:**
>
> Do not complete this step if your Exalogic machine is connected to Oracle Exadata Database Machine on the same InfiniBand fabric.

See the following example:

1. Run the following command to start the process:

   ```
   # smpartition start
   ```

2. Run the following command:

   ```
   # smpartition list modified
   ```

3. Run the following command to make Exalogic compute nodes limited members of the default partition and to disable IPoIB on the default partition:

   ```
   # smpartition modify -n Default -port ALL_CAS -m limited -flag
   ```

16.5 Create an IPoIB Partition and Adding Ports

In this example procedure, you are creating a unique, non-default partition named `myIPoIB` for network isolation on Exalogic's private InfiniBand fabric by configuring a non-default partition key (`pkey`) value `0x005`.

To do so, complete the following steps:

1. Log in to the InfiniBand switch where master SM is running. For more information, see Before You Begin.
2. To start the configuration process, run the following command:

   ```
   # smpartition start
   ```

3. Create the `myIPoIB` partition with the `pkey 0x005` with full membership by running the following command:

   ```
   # smpartition create -n myIPoIB -pkey 0x8005 -m full -flag ipoib
   ```

4. Run the following command to add the compute node port GUIDs, which you noted down in Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches, to the `myIPoIB` partition.

   ```
   # smpartition add -n myIPoIB -port portGUID1 portGUID2
   ```

   In this example, `portGUID1` and `portGUID2` are the ports that you want to add to the partition. This command example shows a few port entries only. You can add as many ports as necessary. An example port value is `0021280001cef8e3`.

5. If you intend to use the partition for creating vNICs, run the following command to add the gateway switch's BridgeX ports to the `myIPoIB` partition; otherwise, proceed to the next step. The gateway switch's BridgeX ports were noted down in Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches.

   ```
   # smpartition add -n myIPoIB -port BridgeXPort1 BridgeXPort2 BridgeXPort3
   BridgeXPort4
   ```

   In this example, `BridgeXPort1`, `BridgeXPort2`, `BridgeXPort3`, and `BridgeXPort4` are the BridgeX ports that you want to add to the partition.

6. Follow these steps to add the ibp0 and the ibp1 network device port GUIDs to the partition:

   a. SSH to the storage appliance and run the following commands to determine the port GUID of the ibp0 network device:

      ```
      :> configuration
      :configuration> net
      :configuration net> devices
      :configuration net devices> select ibp0
      :configuration net devices ibp0> show
      Properties:
                               speed = 32000 Mbit/s
                                  up = true
                              active = false
                               media = Infiniband
                         factory_mac = not available
                                port = 1
                                guid = 0x212800013e8fbf
      configuration net devices ibp0> done
      ```

   b. Repeat the previous step with `select ibp1` to determine the port GUID of the ibp1 device.

   c. Run the following command to add storage appliance GUIDs to the myIPoIB partition:

      ```
      # smpartition add -n myIPoIB -port ibp0GUID ibp1GUID
      ```

7. Run the following command to view the changed partition configuration:

   ```
   # smpartiiton list modified
   ```

   This command displays the new partition with its `pkey`, ports added to the partition, and membership type.

8. Run the following command to confirm the partition configuration:

   ```
   # smpartition commit
   ```

9. Create interfaces for the ibp0 and the ibp1 network devices and bond them by running these steps:

   a. Log in to the Browser User Interface (BUI) of the storage appliance in your Exalogic machine.

   b. Under the **Configuration** tab, select **Network**.

   c. Create a new datalink with the following properties by dragging **ibp0** under Devices to the Datalinks column to:

      i. In the **Name** field, enter ibp0.8005, where 8005 is the partition key.

      ii. In the **Partition Key** field, enter the partition key specified for the partition.

iii. Set the **Link Mode** as Connected Mode.

d. Create a datalink for **ibp1** called ibp1.8005 by performing Step Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches.

e. Create an interface for the ibp0 network device with the following properties by dragging the ibp0.8005 datalink to the interfaces column:

   i. In the **Name** field, enter ib0.8005, where 8005 is the partition key.

   ii. In the **Configure with** field, select **Static Address List**.

   iii. Enter the **IPv4 Address/Mask** as 0.0.0.0/8.

f. Create an interface for **ibp1** called ib1.8005 by performing Step Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches.

g. Bond the interfaces by adding another interface. Click the plus button next to the Interfaces column and enter the following information:

   i. In the **Name** field, enter a name that denotes the IB partition, such as IB_IF_8005.

   ii. In the **Configure with** field, select **Static Address List**.

   iii. Enter the **IPv4 Address/Mask** to configure as the Infiniband device. You can use any unused IP address. In this example, we will use 192.168.33.15/24.

   iv. Check the **IP MultiPathing Group** box.

10. The network must be configured on the compute node and the share mounted by following these steps:

   a. SSH to a compute node that is a member of the partition as the root user.

   b. Determine the active infiniband device by running the following commands on the compute node:

```
# ifconfig ib0
# ifconfig ib1
```

   The active device will have non-zero Rx/Tx bytes.

   c. Configure the network for the active device by running the following commands:

```
# echo 0x8005 > /sys/class/net/ib0/create_child
# ifconfig -a | grep -A6 8005
ib0.8005  Link encap:InfiniBand  HWaddr 80:50:05:4C:FE:
80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
          BROADCAST MULTICAST  MTU:2044  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:256
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

# ifconfig ib0.8005 192.168.33.15
```

   d. Unmount any mounted shares using the `umount` command.

   e. Mount the share from the remote device. Use the IP specified in the configuration step. In our example we used `192.168.33.15`.

```
# mount 192.168.33.15:/export/share_name /mnt/share_name
```

f. Repeat the previous steps for all compute nodes that are members of the partition.

16.6 Delete a Partition

You can delete a non-default partition by running the following command:

```
# smpartition delete -n myIPoIB
```

This command deletes the `myIPoIB` partition.

> **Note:**
>
> Do not attempt to delete the default partition.

16.7 Create a Partition for EoIB and Associating the pkey with a VNIC and VLAN

You can create a partition for EoIB (both inbound and outbound) and associate the partition's `pkey` with a VLAN and VNIC on the edge network.

> **Note:**
>
> The port GUID values, MAC addresses, VLAN IDs, compute node names, gateway switch names, and Ethernet connector names used in this procedure are examples only.

1. At the command prompt on one of the gateway switches, run the following command:

   ```
   el01gw04# listlinkup | grep Bridge
   ```

   The following is an example of the output of the `lislinkup` command:

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

   From this example, identify the uplinks. You can determine that you can use any of the following Ethernet connectors for creating a VNIC:

   - `0A-ETH-1`
   - `0A-ETH-2`

- `0A-ETH-3`
- `0A-ETH-4`
- `1A-ETH-3`
- `1A-ETH-4`

> **Note:**
>
> This procedure uses `1A-ETH-3` as an example.

2. Determine GUIDs of the Exalogic compute node the requires the VNIC as follows:

   a. On the compute node that requires the VNIC, log in as `root`, and run the `ibstat` command on the command line. For example, log in to `el01cn01` as `root`.

   Example:

```
el01cn01# ibstat
CA 'mlx4_0'
        CA type: MT26428
        Number of ports: 2
        Firmware version: 2.7.8100
        Hardware version: b0
        Node GUID: 0x0021280001a0a364
        System image GUID: 0x0021280001a0a367
        Port 1:
                State: Active
                Physical state: LinkUp
                Rate: 40
                Base lid: 120
                LMC: 0
                SM lid: 6
                Capability mask: 0x02510868
                Port GUID: 0x0021280001a0a365
                Link layer: IB
        Port 2:
                State: Active
                Physical state: LinkUp
                Rate: 40
                Base lid: 121
                LMC: 0
                SM lid: 6
                Capability mask: 0x02510868
                Port GUID: 0x0021280001a0a366
                Link layer: IB
```

   In the output, information about two ports is displayed. Identify the `GUID` and `Base lid` of the port that you want to use for creating the VNIC.

   For the example illustrated in this procedure, we will use the port with `GUID 0x0021280001a0a366` and `Base lid 121`.

   b. On the same compute node, run the following command to view information about all the active links in the InfiniBand fabric:

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

`hostname` is the name of the compute node. You can also specify the bonded IPoIB address of the compute node.

Example:

```
el01cn01# iblinkinfo.pl -R | grep el01cn01
65   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    121   2[  ] "el01cn01 EL-
C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
64   15[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>    120   1[  ] "el01cn01 EL-
C 192.168.10.29 HCA-1" (Could be 5.0 Gbps)
```

From the output of the `iblinkinfo` command, note the switch `lid` value (`65`, in first column) associated with the `Base lid` of the compute node port that you noted earlier (`121`, in the first line):

3. Determine the gateway switch that corresponds to the switch lid `65` by running the `ibswitches` command, as in the following example:

   Example:

   ```
   el01cn01# ibswitches
   Switch  : 0x002128548042c0a0 ports 36 "SUN IB QDR GW switch el01gw03" enhanced
   port 0 lid 63 lmc 0
   Switch  : 0x002128547f22c0a0 ports 36 "SUN IB QDR GW switch el01gw02" enhanced
   port 0 lid 6 lmc 0
   Switch  : 0x00212856d0a2c0a0 ports 36 "SUN IB QDR GW switch el01gw04" enhanced
   port 0 lid 65 lmc 0
   Switch  : 0x00212856d162c0a0 ports 36 "SUN IB QDR GW switch el01gw05" enhanced
   port 0 lid 64 lmc 0
   ```

   `lid 65` corresponds to gateway switch `el01gw04` with GUID `0x00212856d0a2c0a0`.

4. Define a dummy MAC address in the following format:

   ```
   last3_octets_of_switchGUID : last3_octets_of_computenode_adminIP_in_hex_format
   ```

   Example:

   GUID of switch: `00:21:28:56:d0:a2:c0:a0`

   Last three octets: `a2:c0:a0`

   Administrative IP of the compute node that requires the VNIC: `192.168.1.1`

   Last three octets: `168.1.1` (in hexadecimal notation: `a8:01:01`)

   MAC address: `a2:c0:a0:a8:01:01`

   > **Note:**
   >
   > The dummy MAC address should be unique to the Exalogic network. Only even numbers are supported for the most significant byte of the MAC address (unicast). The above address is an example only.

5. Ensure that you have noted down all port GUIDs and BridgeX ports.

6. Log in to the InfiniBand switch where master SM is running. For more information, see Before You Begin.

7. Run the following command to start the configuration process:

   ```
   # smpartition start
   ```

8. Run the following command to create a `myEoIB` partition with the `pkey 0x005` with a full membership:

```
# smpartition create -n myEoIB -pkey 0x005 -m full
```

9. Run the following command to add port GUIDs and BridgeX ports, which you noted down in Gather Port GUIDs of Compute Nodes and BridgeX Ports of Gateway Switches, to the `myEoIB` partition:

```
# smpartition add -n myEoIB -port port_guid1 port_guid2 bridgex_port1
bridgex_port2
```

   Where `port_guid1`, `port_guid2`, `bridgex_port1`, and `bridgex_port2` are the ports that you want to add to the partition. This command example shows a few port entries only. You can add as many ports as necessary. An example port value is `0021280001cef8e3`.

10. Run the following command to view the changed partition configuration:

```
# smpartition list modified
```

   This command displays the new partition with its `pkey`, ports added to the partition, and membership type.

11. Run the following command to confirm the partition configuration:

```
# smpartition commit
```

   The `myEoIB` partition with `0x005 pkey` is created.

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

```
# createvlan 1A-ETH-3 -vlan 10 -pkey 0x005
```

   Where `1A-ETH-3` is the Ethernet connector on the gateway switch, `10` is the VLAN identifier, and `0x005` is the partition key that you created earlier.

13. To verify, run the following command:

```
# showvlan
```

   The following information is displayed:

```
Connector/LAG    VLN    PKEY
--------------   ---    -----
1A-ETH-3          10    0x005
0A-ETH-1          11     ffff
```

14. As `root`, log in to `el01gw04` that you identified in Step 4. Use its IP address or host name to log in.

15. Upon login, run the following command to create a VNIC:

```
# createvnic 1A-ETH-3 -GUID 00212856d0a2c0a0 -mac a2:c0:a0:a8:01:01 -vlan 10 -
pkey 0x005
```

   Where `1A-ETH-3` is the Ethernet connector, `00:21:28:56:d0:a2:c0:a0` is the GUID, `a2:c0:a0:a8:01:01` is the dummy MAC address defined in Step 4, `10` is the VLAN identifier, and `0x005` is the partition key that you created earlier.

   This example creates a VNIC, such as `eth4` (on Oracle Linux) or `eoib0` (on Oracle Solaris) associated with VLAN `10` associated with a partition with `0x005` as the `pkey`.

16. Run the following command to verify the VNICs:

```
# showvnics
```

The following message is displayed:

```
ID  STATE FLG IOA_GUID                NODE                           IID
MAC              VLN PKEY  GW
--- ----- --- ----------------------- --------------------------- ----
----------------- --- ----- --------
8   UP    N   00:21:28:00:01:A0:A3:66 e101cn01 EL-C 192.168.10.29 0000
a2:c0:a0:a8:01:01 10  0x005 1A-ETH-3
```

> **Tip:**
>
> After creating the interfaces, you can run the `ifconfig` command with the `-a` option to verify the MAC address on the compute node. For example, to verify the new interface and its MAC address, run the following command on the Oracle Linux compute node for which the VNIC was created:
>
> ```
> # ifconfig -a eth4
> ```
>
> The output of this command shows the `HWADDR`, which is the MAC address you defined for the VNIC in Step 5.

17. On the compute node, run the following command to display the list of VNICs available on the compute node:

    ```
    el01cn01# mlx4_vnic_info -l
    ```

    This command displays the name of the new interface, as seen on the compute node, such as `eth4`. Note this ID.

18. Create another VNIC for the same compute node, but using a connector on a different gateway switch. Note the `ethX` ID of this VNIC too.

    It is recommended that you configure the two EoIB interfaces as a bonded interface, such as `bond1`.

19. Create interface files for the VNICs on the compute node.

    To ensure correct failover behavior, the name of the VNIC interface file and the value of the `DEVICE` directive in the interface file must *not* be based on the kernel-assigned `ethX` interface name (`eth4`, `eth5`, and so on). Instead, Oracle recommends that the interface file name and value of the `DEVICE` directive in the interface file be derived from the `EPORT_ID` and `IOA_PORT` values, as follows:

    > **Note:**
    >
    > Any other unique naming scheme is also acceptable.

    a. Run the following command to find the `EPORT_ID`:

       ```
       #mlx4_vnic_info -i ethX | grep EPORT_ID
       ```

       Example:

       ```
       e101cn01#mlx4_vnic_info -i eth4 | grep EPORT_ID
       EPORT_ID     331
       ```

Note the EPORT_ID that is displayed, 331 in this example.

b. Run the following command to find the IOA_PORT:

```
#mlx4_vnic_info -i ethX | grep IOA_PORT
```

Example:

```
e101cn01#mlx4_vnic_info -i eth4 | grep IOA_PORT
IOA_PORT     mlx4_0:1
```

Note the number after the colon (:) in the IOA_PORT value that is displayed, in this case 1.

c. Build the interface file name and device name by using the following convention:

Interface file name: ifcfg-eth**A_B**

Device name: eth**A_B**

A is the EPORT_ID, and B is the number after the colon (:) in the IOA_PORT value.

Example:

Interface file name: ifcfg-eth**331**_1

Device name: eth**331_1**

In this example, 331 is the EPORT_ID, and 1 is the value derived from the IOA_PORT.

20. Create the interface file for the first VNIC, eth4 in the example, by using a text editor such as vi.

Save the file in the /etc/sysconfig/network-scripts directory.

Example:

```
# more /etc/sysconfig/network-scripts/ifcfg-eth331_1
DEVICE=eth331_1
BOOTPROTO=none
ONBOOT=yes
HWADDR=a2:c0:a0:a8:01:01
MASTER=bond1
SLAVE=yes
```

- Make sure that the name of the interface file (ifcfg-eth331_1 in the example) is the name derived in step 19.
- For the DEVICE directive, specify the device name (eth331_1 in the example) derived in step 19.
- For the HWADDR directive, specify the dummy MAC address created in step 4.

21. Create an interface file for the second VNIC, say eth5. Be sure to name the interface file and specify the DEVICE directive by using a derived interface name and not the kernel-assigned name, as described earlier. In addition, be sure to specify the relevant dummy MAC address for the HWADDR directive.

22. After creating the interface files, create the ifcfg-bond1 file. If the file already exists, verify its contents.

Example:

```
# more /etc/sysconfig/network-scripts/ifcfg-bond1
DEVICE=bond1
IPADDR=192.168.48.128
NETMASK=255.255.255.0
BOOTPROTO=none
USERCTL=no
TYPE=Ethernet
ONBOOT=yes
IPV6INIT=no
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000"
GATEWAY=192.168.48.1
```

23. Restart the network services by running the following command:

```
# service network restart
```

24. Bring up the new `bond1` interface using the `ifup` command.

    You must also reboot the compute node for the changes to take effect.

16.8 Perform the Post-Configuration Steps

After creating a partition on the InfiniBand switch, you must create a child interface for the IPoIB interface on your Exalogic compute node.

For example, on the InfiniBand switch, if you defined a partition with `pkey 0x33`, with IPoIB enabled, you must complete the following steps on a compute node with port 1 that is either full or limited member of that partition:

> **Note:**
>
> Even though the example uses port 1, you can create child interfaces for both `ib0` and `ib1` and bond them together on the partitioned network.

1. Log in as a `root` user.
2. Run the following commands on the command line:

```
# cd /sys/class/net/ib0
```

```
echo 0x8033 > create_child
```

3. Run the following command to verify that the child interface was created:

```
# ifconfig ib0.8033
```

4. Specify your setup for the child interface in an `ibcfg-ib0.8033` file in the `/etc/sysconfig/networks-scripts` directory. Note that it is `.8033` even if it might be limited member.

16.9 Important Notes for Combined Exalogic-Exadata Fabric Users

Read the following notes if you are using partitions in a scenario where your Exalogic machine is connected to an Oracle Exadata Database Machine on the same InfiniBand fabric:

- Oracle Exadata Database Machine currently uses the default InfiniBand partition only. Therefore, Oracle Exadata Database Machine nodes are full members of the default partition.
- If your Exalogic machine is connected to the Oracle Exadata Database Machine on the same InfiniBand fabric, ensure that all Exalogic compute nodes are limited members of the default partition. By default, all Exalogic compute nodes are full members of the default partition. To make an Exalogic compute node a limited member of the default partition, add the port GUIDs of the compute node as limited members of the default partition. In addition, ensure that IPoIB is enabled on the default partition.
- Exalogic nodes as limited members of the default partition will not be able to communicate with any other Exalogic node in the default partition. However, client access to Oracle Exadata Database Machine is provided via IPoIB in the default partition.
- You must disable Subnet Manager (SM) on all InfiniBand switches that are not using firmware 2.0.4 or above. Exalogic's InfiniBand switches use firmware versions 2.0.4 or above for partitioning support, and SM should run on one of Exalogic's InfiniBand switches.

16.10 Partitioning Limitations

Consider the following limitations when creating non-default partitions:

- Once a new partition configuration is successfully committed using the `smpartition` command on the current master SM, the configuration is kept highly available among the defined set of SM instances. However, all Sun Network QDR InfiniBand Gateway Switches defined to have SM enabled (that is, defined by the `smnodes` command on each gateway switch) must be operational and able to communicate with the other smnodes gateway switches in order for any change in the partition configuration to take place.
- The limitation for number of partitions per end-port is a constraint defined by the various end-port implementations. For ConnectX2 and BridgeX, this limit is 128, which includes the default partition. Hence the maximum number of other partitions is 127. The CLI interface of the gateway switch does not verify this explicitly. However, if you specify partitions more than the maximum limit for any port (GUID), the SM only handles the maximum number of partitions and then logs a message.

17 Monitoring the Exalogic Machine Using Oracle Enterprise Manager Ops Center

This chapter describes how to manage assets of an Exalogic machine in a physical configuration using a standalone Oracle Enterprise Manager 11g Ops Center installation that is external to Exalogic.

The information in this chapter is intended to be a reference model, to help you understand how an external, standalone installation of an Oracle product like Enterprise Manger Ops Center 11g or Enterprise Manager Cloud Control can be used to manage the assets in an Exalogic machine in a physical configuration.

- For more information about Enterprise Manager Ops Center 12c, the preferred product for hardware-level management, see the *Enterprise Manager Ops Center 12c documentation* at `http://docs.oracle.com/cd/E27363_01/index.htm`.
- For more information about Enterprise Manager Cloud Control, see the *Enterprise Manager Cloud Control 12c documentation* at `http://docs.oracle.com/cd/E24628_01/index.htm`.

> **Note:**
>
> Exalogic virtual configurations include Exalogic Control, which provides the management and monitoring interface for Exalogic. The information in this chapter is relevant only for Exalogic physical configurations.

This chapter contains the following sections:

- About Problem Management
- Using Monitoring Profiles and Rules
- Using Reports in Oracle Enterprise Manager Ops Center
- Using Oracle Services in Oracle Enterprise Manager Ops Center

17.1 Overview

Oracle Enterprise Manager 11*g* Ops Center can optionally be used to monitor the following components in the Exalogic machine infrastructure:

- Exalogic compute nodes
- Storage appliance
- Sun Network QDR InfiniBand Gateway Switches
- Sun Datacenter InfiniBand Switch 36

Although Oracle Enterprise Manager is optional in the Exalogic machine environment, Oracle recommends that you use Oracle Enterprise Manager Ops Center to monitor the hardware components of the Exalogic machine.

17.2 Key Features

Oracle Enterprise Manager 11*g* Ops Center supports the following key features in the Exalogic machine environment:

- Hardware lifecycle management
- InfiniBand fabric views and Ethernet network view
- Console access to launch the browser user interface for managing the storage appliance and the InfiniBand switches
- Serial console access to the service processors of compute nodes, switches, and storage appliance
- Integration with Oracle Services
- Problem management for the storage appliance

> **Note:**
>
> The storage appliance has the ability to create service requests when it detects a problem condition. When the appliance is running in the Oracle Enterprise Manager Ops Center environment, Ops Center also detects the problem condition and creates an alert. You can use this feature of Ops Center to report all problems.

In addition, Oracle Enterprise Manager Ops Center supports bare metal provisioning, discovery of hardware assets, patch automation, import of firmware images, creation policies, and firmware upgrade for the hardware components of an Exalogic machine.

17.3 Prerequisites

The following are the prerequisites for using Oracle Enterprise Manager Ops Center:

- Installing and configuring Oracle Enterprise Manager Ops Center outside of the Exalogic machine

  You must connect the system that you installed Oracle Enterprise Manager Ops Center on to the Exalogic machine's Ethernet network either directly to the switch in the rack or to a datacenter switch carrying that network.

  For information about installing Oracle Enterprise Manager Ops Center, see the following URL:

  `http://download.oracle.com/docs/cd/E11857_01/nav/management.htm`

  > **Note:**
  >
  > After installing Oracle Enterprise Manager Ops Center, you can deploy the enterprise controller and proxy controllers on the same machine.
  >
  > For more information about site preparation and installation, see the "Site Preparation" and "Installation" sections in the Oracle Enterprise Manager Ops Center documentation.

- Verifying the IP addresses assigned to each of the management and data interfaces on the hardware components of the Exalogic machine

17.4 Accessing Oracle Enterprise Manager Ops Center Documentation

You can access Oracle Enterprise Manager Ops Center documentation at the following URL:

`http://docs.oracle.com/cd/E11857_01/nav/management.htm`

17.5 Launching Oracle Enterprise Manager Ops Center

When you launch Oracle Enterprise Manager Ops Center in a web browser, the login page is displayed.

After successful login, the home page is displayed, as in Figure 17-1.

Figure 17-1 Oracle Enterprise Manager Home Page

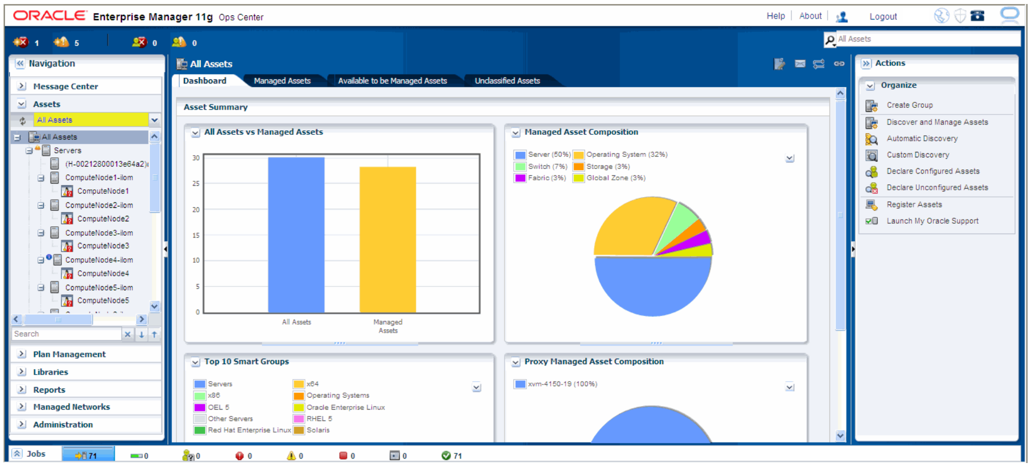

17.6 Understanding the Workflow

Figure 17-2 illustrates the workflow for the management of Exalogic machine hardware components using Oracle Enterprise Manager Ops Center.

Figure 17-2 Management of Exalogic Machine Hardware Using Oracle Enterprise Manager Ops Center

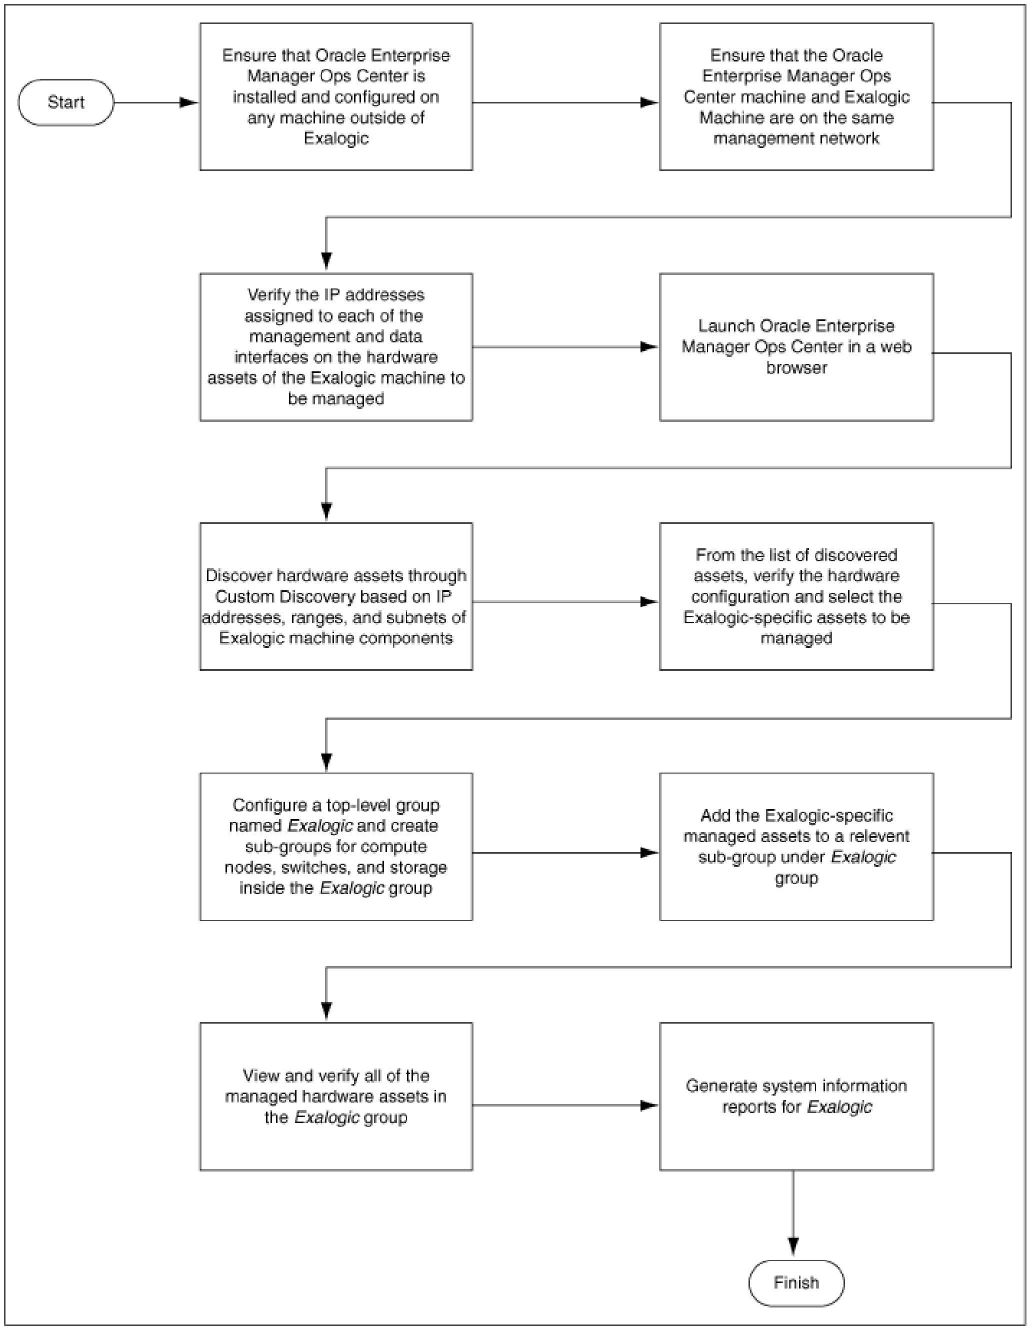

17.7 Managing Users and Roles

You can add users to Oracle Enterprise Manager Ops Center from the local authentication subsystem of the Enterprise Controller's operating system. Each user can be given a different role which grants or denies access to the different functions of Oracle Enterprise Manager Ops Center. You can view the existing Users from the

Users tab of the **Administration** section.You can view the Roles of existing users from the **Roles** tab of the **Administration** section.

> **Note:**
>
> For more information about User and Role management in Enterprise Manager Ops Center, see the "User and Role Management" topic in the Administration section in the Oracle Enterprise Manager Ops Center 11g documentation. This topic describes the following tasks:
>
> - Adding a user
> - Deleting a user
> - Roles and authorizations
> - Assigning a role
> - Configuring a notification profile
> - Deleting a notification profile

17.8 Discovering and Managing Exalogic Machine Hardware

This section describes how to discover and manage the hardware components of the Exalogic machine. For a list of the hardware components managed by Oracle Enterprise Manager Ops Center, see Overview.

To discover and manage Exalogic hardware:

1. On the Oracle Enterprise Manager Ops Center home page, click **Custom Discovery** on the **Actions** panel.

   The following screen is displayed.

Figure 17-3 Custom Asset Discovery Screen

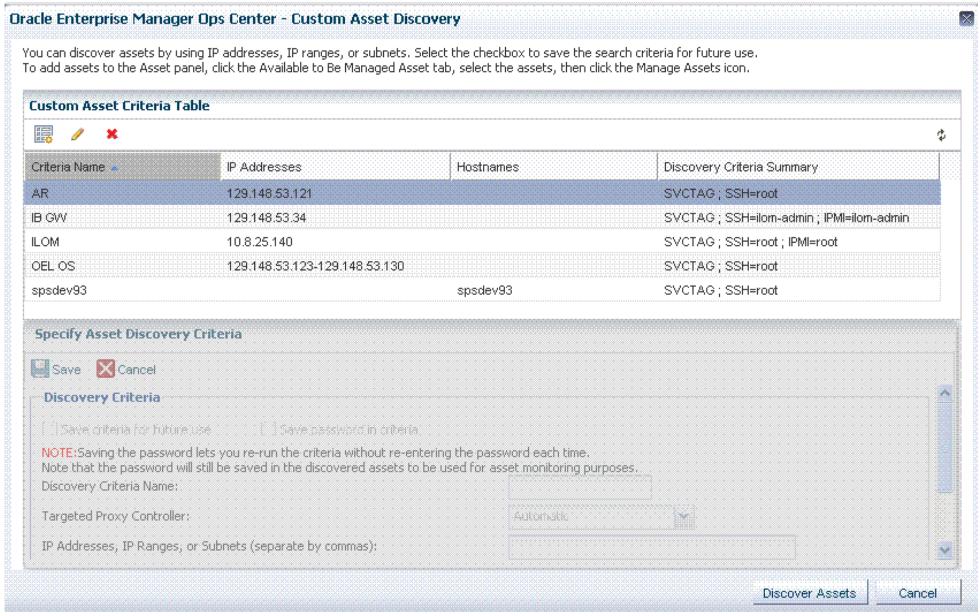

2. In the Custom Asset Criteria Table section, click the **New Criteria** icon to define new discovery criteria. The Specify Asset Discovery Criteria section in Figure 17-3 is enabled. You must define discovery criteria for the following assets in your Exalogic machine individually:
   - Compute Nodes
   - Storage appliance
   - Sun Network QDR InfiniBand Gateway Switches
   - Sun Datacenter InfiniBand Switch 36

   > **Note:**
   >
   > This switch is used in multirack configuration scenarios only. It is not connected in a Exalogic machine single rack configuration.

3. For each of the above asset categories, in the Specify Asset Discovery Criteria section, enter the following:
   - A name for the criteria

     > **Note:**
     >
     > For example, you can enter `ComputeNode` for the compute nodes' discovery criteria, `Storage` for the storage appliance's discover criteria, `Gateway` for the gateway switches' discovery criteria.

   - IP addresses, ranges, or subnets

> **Note:**
>
> In this field, enter the IP addresses or ranges assigned to your Exalogic machine components.

- Host names

> **Note:**
>
> In this field, enter the host names, if any, assigned to your Exalogic machine components.

4. After setting the discovery criteria for each of the Exalogic machine asset categories, name and save the criteria. You can check **Save Criteria for Future Use** to save the criteria for future use.

> **Note:**
>
> You can launch the discoveries all at once by multiple selection of the saved criteria.

5. Click **Discover Assets**.

   When the assets are discovered, they are listed under the **Available to be Managed Assets** tab on the home page.

6. To add assets to the Assets panel, click the **Available to be Managed Assets** tab, select the assets, and click the **Managed Assets** icon. The selected asset is now listed under the **Managed Assets** tab on the home page.

> **Note:**
>
> In addition to discovering hardware assets, you may discover operating systems running on the host compute nodes.

17.9 Grouping Exalogic Machine Hardware Assets

This section describes how to group the Exalogic-specific managed assets to reflect their physical containment hierarchy.

It contains the following topics:

- Prerequisites
- Creating the Exalogic Top-Level Group
- Creating a Sub-Group for Exalogic Compute Nodes
- Creating a Sub-Group for the Storage Appliance

- Creating a Sub-Group for InfiniBand Switches
- Adding Assets to a Group

17.9.1 Prerequisites

Before grouping Exalogic-specific managed assets, ensure that all Exalogic-specific hardware assets, such as compute nodes, storage appliance, and network switches are discovered and managed. However, you can extend or modify the groups at any time. For more information, see Discovering and Managing Exalogic Machine Hardware.

In addition, verify that all Exalogic-specific assets are listed under **Managed Assets**, as shown in the example (Figure 17-9).

17.9.2 Creating the Exalogic Top-Level Group

To create a top-level group for Exalogic machine assets, complete the following steps:

1. On the right navigation pane named **Actions**, click **Create Group** under **Organize**. The Configure/Modify Group screen is displayed.
2. In the Configure Group screen, in the **Group Name** field, enter a name for the new Exalogic group. For example, enter `Exalogic`. Enter a description in the **Description** field. Select the **Top Level (Root)** option as **Location**. See Figure 17-4.

Figure 17-4 Configure Group

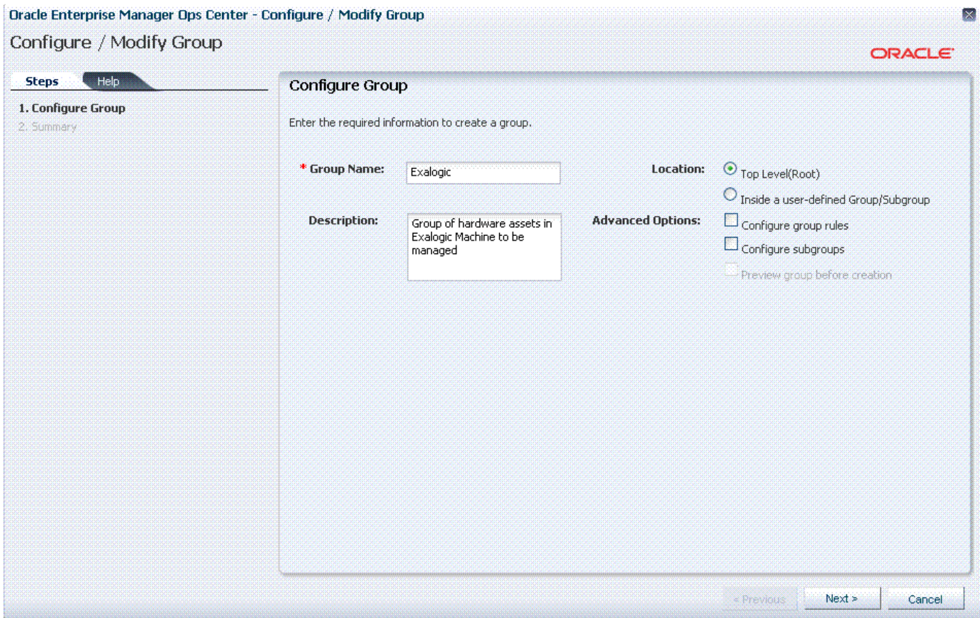

Note:

When you are naming your group, you can also specify the host name of an Exalogic rack. For example, you can enter `Exalogic el01` as a group name.

3. After entering the group information and selecting the location, click **Next**. The Summary screen is displayed, as in Figure 17-5.

Figure 17-5 Group Summary Screen

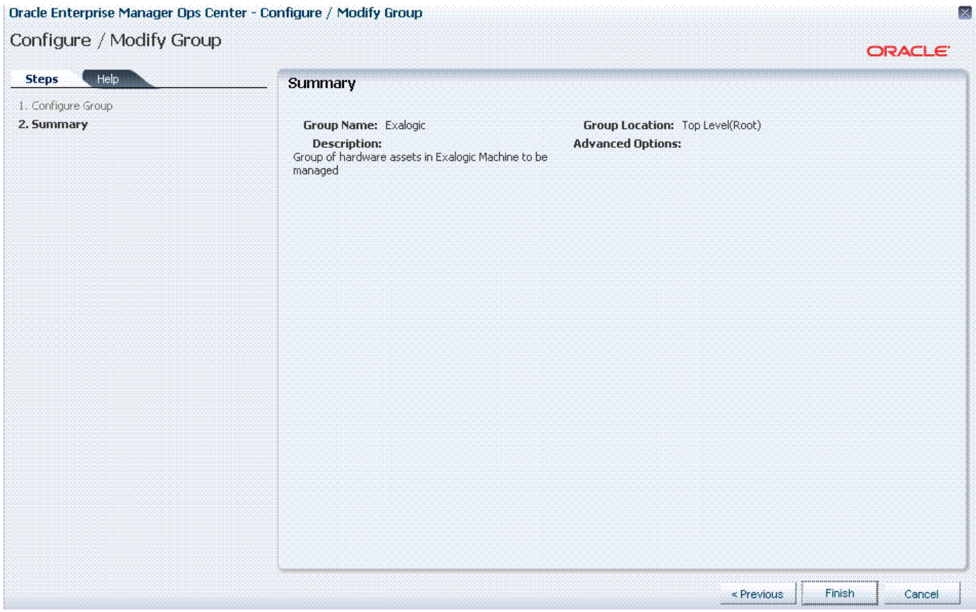

4. On the Summary screen, click **Finish**. The `Exalogic` group is created at the top-level (root).

17.9.3 Creating a Sub-Group for Exalogic Compute Nodes

After creating a top-level group for the Exalogic machine, you can create a sub-group for Exalogic compute nodes as follows:

1. On the right navigation pane named **Actions**, click **Create Group** under **Organize**. The Configure/Modify Group screen is displayed.
2. In the Configure Group screen, in the **Group Name** field, enter a name for the new Exalogic group. For example, enter `ComputeNodes`. Enter a description in the **Description** field. Select the **Inside a user-defined Group/Subgroup** option as **Location**. In the Group Selection Panel, select the `Exalogic` top-level group that you created in Creating the Exalogic Top-Level Group. The Configure Group screen should be similar to Figure 17-6.

Figure 17-6 Configure Sub-Group for Compute Nodes

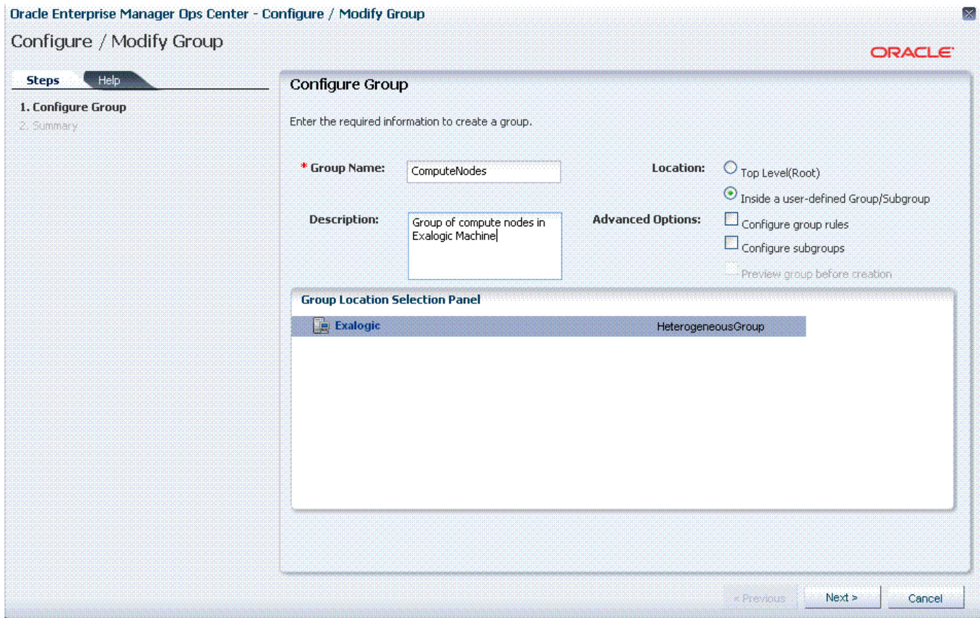

3. After entering the sub-group information and selecting the group location, click **Next**. The Summary screen is displayed.
4. On the Summary screen, click **Finish**. The `ComputeNodes` sub-group is created under the `Exalogic` group.

17.9.4 Creating a Sub-Group for the Storage Appliance

After creating a top-level group for the Exalogic machine, you can create a sub-group for the storage appliance, which is included in the Exalogic machine.

You can create a sub-group for this storage appliance as follows:

1. On the right navigation pane named **Actions**, click **Create Group** under **Organize**. The Configure/Modify Group screen is displayed.
2. In the Configure Group screen, in the **Group Name** field, enter a name for the new Exalogic group. For example, enter `Storage`. Enter a description in the **Description** field. Select the **Inside a user-defined Group/Subgroup** option as **Location**. In the Group Selection Panel, select the `Exalogic` top-level group that you created in Creating the Exalogic Top-Level Group. The Configure Group screen should be similar to Figure 17-7.

Figure 17-7 Configure Sub-Group for Storage Appliance

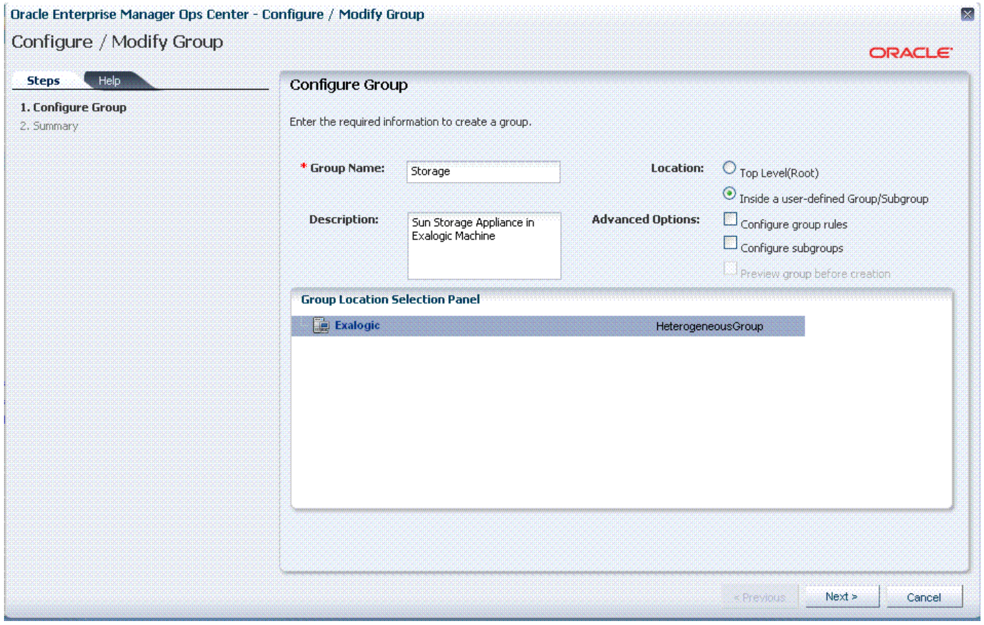

3. After entering the sub-group information and selecting the group location, click **Next**. The Summary screen is displayed.

> **Note:**
>
> Ensure that you do not select the **Configure group rules** option.

4. On the Summary screen, click **Finish**. The `Storage` sub-group is created under the `Exalogic` group.

17.9.5 Creating a Sub-Group for InfiniBand Switches

After creating a top-level group for the Exalogic machine, you can create a sub-group for the InfiniBand switches and gateways (leaf switches and spine switches), which are included in the Exalogic machine.

You can create a sub-group for these InfiniBand switches as follows:

1. On the right navigation pane named **Actions**, click **Create Group** under **Organize**. The Configure/Modify Group screen is displayed.
2. In the Configure Group screen, in the **Group Name** field, enter a name for the new Exalogic group. For example, enter `IBswitches`. Enter a description in the **Description** field. Select the **Inside a user-defined Group/Subgroup** option as **Location**. In the Group Selection Panel, select the `Exalogic` top-level group that you created in Creating the Exalogic Top-Level Group. The Configure Group screen should be similar to Figure 17-8.

Figure 17-8 Configure Sub-Group for InfiniBand Switches

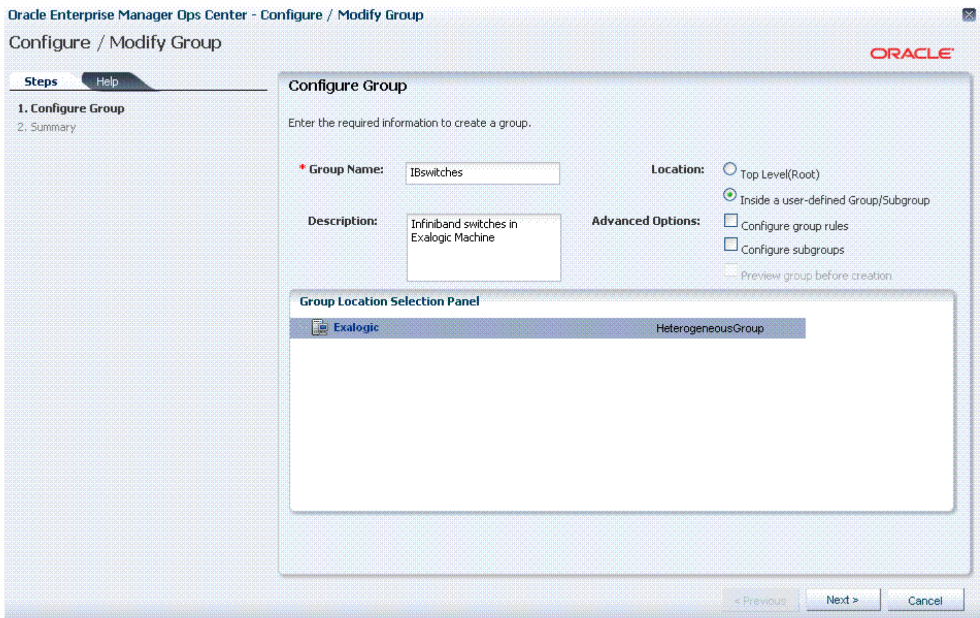

3. After entering the sub-group information and selecting the group location, click **Next**. The Summary screen is displayed.
4. On the Summary screen, click **Finish**. The `IBswitches` sub-group is created under the `Exalogic` group.

17.9.6 Adding Assets to a Group

After creating the top-level Exalogic group and its sub-groups, you must add Exalogic-specific managed assets to a group.

To add an Exalogic-specific managed asset to a group, do the following:

1. On the All Assets page, click the **Managed Assets** tab. Exalogic-specific managed assets are listed under **All Managed Assets**.
2. Select a managed asset, and click the **Add Asset to Group** icon.

   Note that you can add more than one asset to a group or subgroup simultaneously by selecting the assets. Hold the **Ctrl** key on your keyboard while selecting the assets.

   The Add Assets to Group/SubGroup screen is displayed.
3. Select the relevant sub-group (`ComputeNodes`, `Storage`, or `IBswitches`) under the `Exalogic` group. Click the **Add Assets to Group** button. The selected asset is added to the selected group/sub-group.
4. Repeat these steps for each of the Exalogic compute nodes, the storage appliance, and the InfiniBand switches listed under Managed Assets.

17.10 Viewing Exalogic Compute Nodes

To view Exalogic compute nodes, do the following:

1. On the left navigation pane, under **Assets**, from the drop-down list, select the **All User Defined Groups** option.
2. Select the `Exalogic` top-level group. The Exalogic group page is displayed, as in Figure 17-9.

Figure 17-9 Exalogic Group View

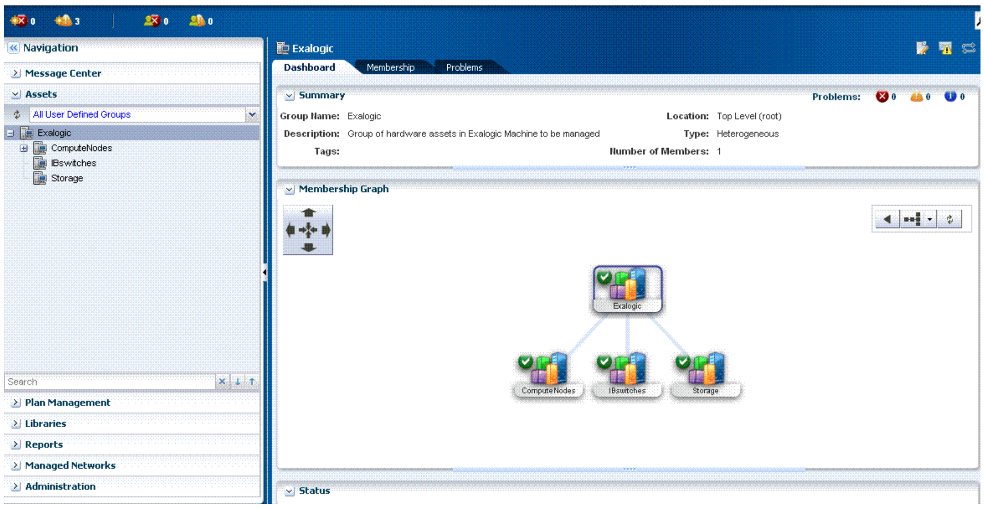

3. On the left navigation pane, select a particular compute node under the `ComputeNodes` sub-group to view information about that compute node.
4. Click the **Hardware** tab to view detailed information about that compute node. This information includes CPU, memory, network adapter, disk, power supply, and fan tray (fan). For example, see Figure 17-10.

Figure 17-10 Compute Node Hardware Information

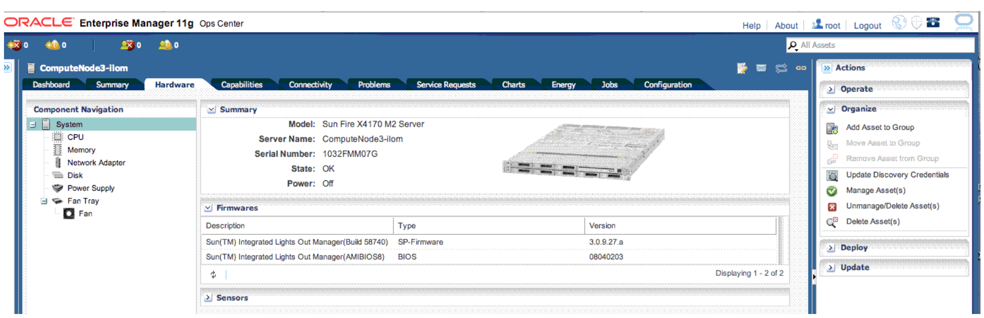

5. Review this information to verify that the configuration corresponds to the compute node configuration in your Exalogic machine.
6. If you wish to launch the Service Processor (SP) console from within this screen, click the **Console** tab. If your console connection is not enabled, enable it by clicking the **Enable the Console Connection** icon.
7. Click the **Connect to the Console** icon. The SP console for the selected compute node is launched, as shown in Figure 17-11.

Figure 17-11   Launching Service Processor Console

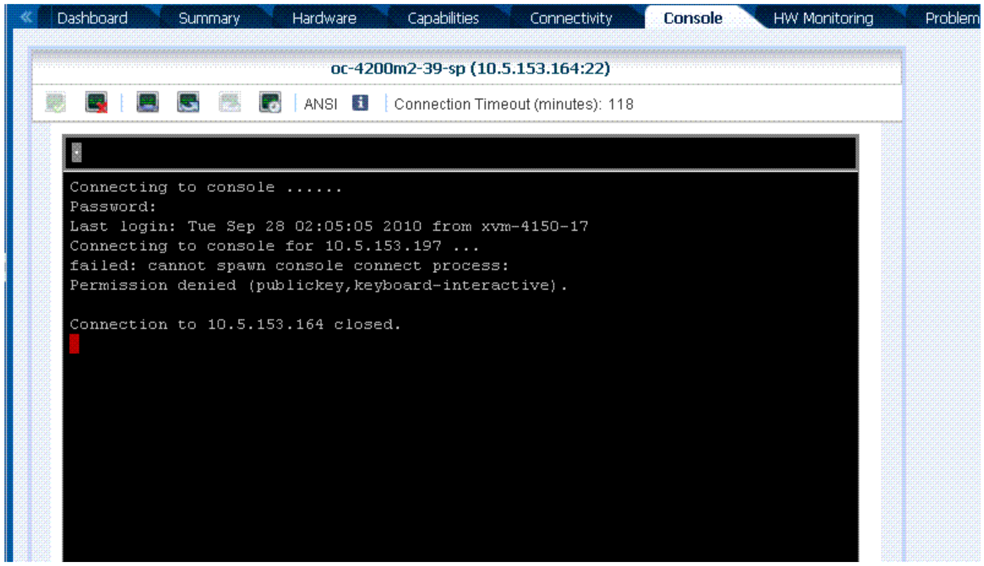

8. On the right navigation pane named **Actions**, click the **Launch LOM Controller** option to verify that the Sun Integrated Lights Out Manager (ILOM) web interface for the compute node can be launched.

17.11 Viewing InfiniBand Switches

To view Exalogic-specific InfiniBand switches, do the following:

1. On the left navigation pane, under **Assets**, from the drop-down list, select the **All User Defined Groups** option.
2. Select the `Exalogic` top-level group. The Exalogic group page is displayed, as in Figure 17-9.
3. On the left navigation pane, select a particular switch under the `IBswitches` sub-group to view information about that switch. You should see a page similar to the compute node page shown in Figure 17-12.

Figure 17-12   Switch View

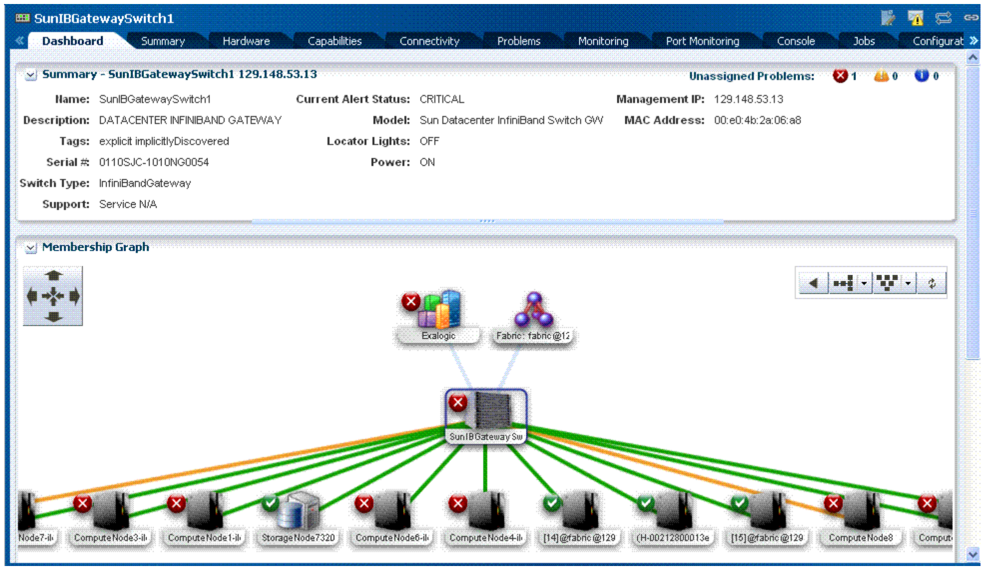

4. Click the **Hardware** tab to view detailed information about that switch. This information includes ports, network adapters connected to the switch, and so on. The hardware information about the switch is displayed, as in Figure 17-13.

Figure 17-13   Switch Hardware Information

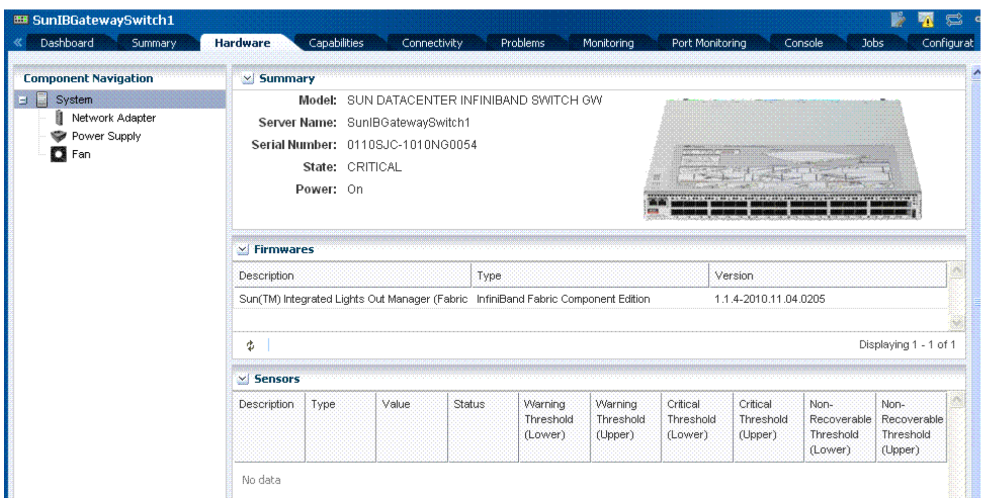

5. Review this information to verify that the configuration corresponds to the corresponding switch specifications in your Exalogic machine.

   In addition, verify that the information matches with the configuration of the network managed by Oracle Enterprise Manager Ops Center. To view information about managed networks, on the left navigation page, click **Managed Networks** to view network information.

6. If you wish to view connectivity information about service processor and server port, click the **Connectivity** tab.

7. If you wish to view information about switch ports, click the **Port Monitoring** tab.

8. If you wish to launch the Service Processor (SP) console from within this interface, click the **Console** tab. If your console connection is not enabled, enable it by clicking the **Enable the Console Connection** icon.
9. Click the **Connect to the Console** icon. The SP console for the selected switch is launched, as shown in Figure 17-11.
10. On the right navigation pane named **Actions**, click the **Launch LOM Controller** option to verify that the Sun Integrated Lights Out Manager (ILOM) web interface for the switch can be launched.

17.12 Viewing the Storage Appliance

To view the storage appliance included in the Exalogic machine, do the following:

1. On the left navigation pane, under **Assets**, from the drop-down list, select the **All User Defined Groups** option.
2. Select the `Exalogic` top-level group. The Exalogic group page is displayed, as in Figure 17-9.
3. On the left navigation pane, select the storage appliance under the `Storage` sub-group to view information about the storage appliance, as shown in Figure 17-14.

Figure 17-14 Storage Appliance View

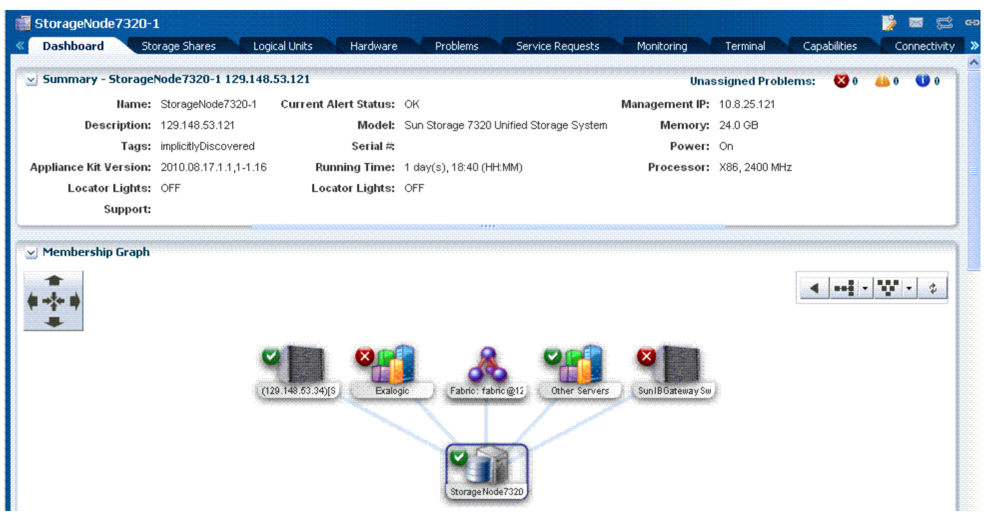

4. Click the **Hardware** tab to view detailed information about the storage appliance. This information includes CPU, memory, network adapters, disks, and so on.

Figure 17-15 Storage appliance Hardware View

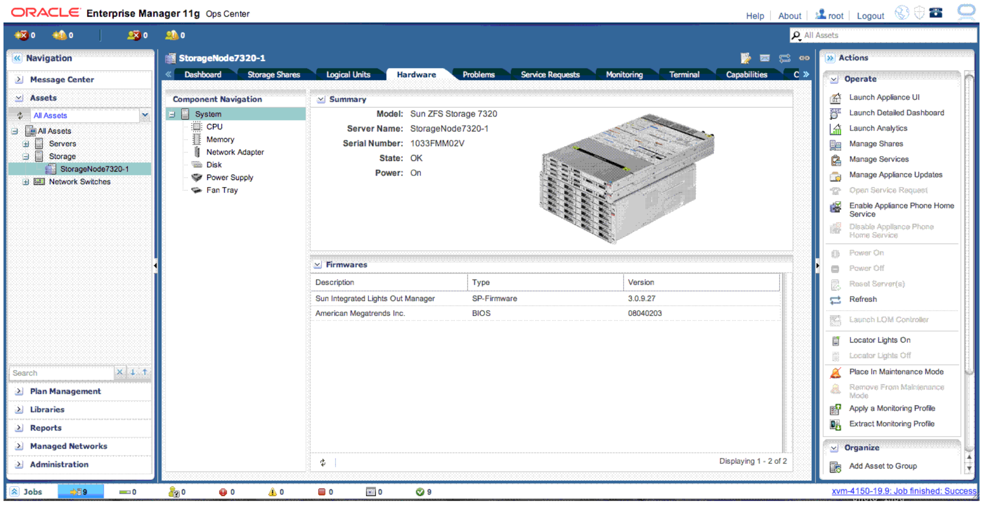

5. Review this information to verify that the configuration corresponds to the corresponding storage appliance specifications in your Exalogic machine.
6. Click the **Storage Shares** tab to view the shares (exported file systems) configured on the storage appliance, as shown in Figure 17-16.

Figure 17-16 Viewing Shares Configured on the Storage Appliance

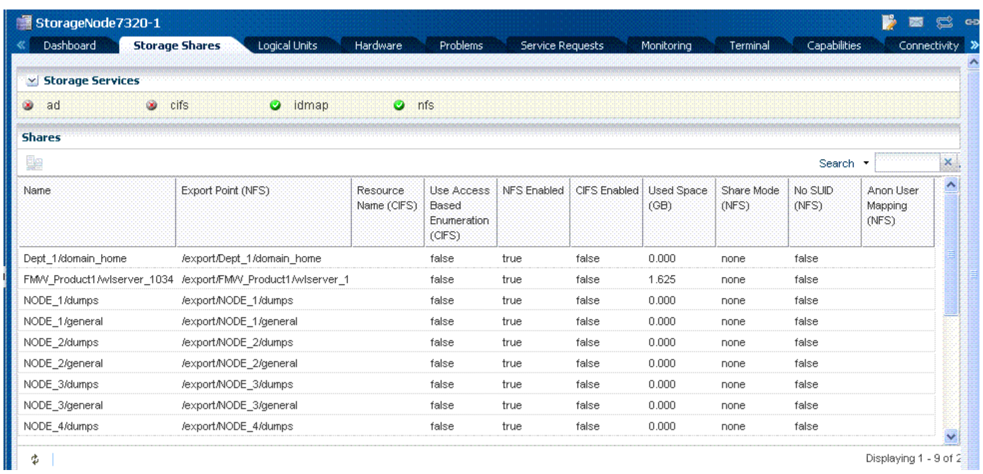

7. If you wish to launch the administration console for the storage appliance from within the Oracle Enterprise Manager Ops Center interface, click the **Terminal** tab.

> **Note:**
>
> You can also launch the storage appliance UI, detailed dashboard, analytics, and manage shares and services by selecting an appropriate action from the **Operate** actions pane, as shown in Figure 17-15.

17.13 Viewing the InfiniBand Fabric and Its Nodes

You can view the InfiniBand fabric and its nodes in Enterprise Manager Ops Center.

To view the InfiniBand fabric, do the following:

1. On the left navigation pane, under **Managed Assets**, from the drop-down list, select the option that starts with the `Fabric:` string. This string is of the format:

   `Fabric: fabric@<IP address>`

   Figure 17-17 Fabric View

   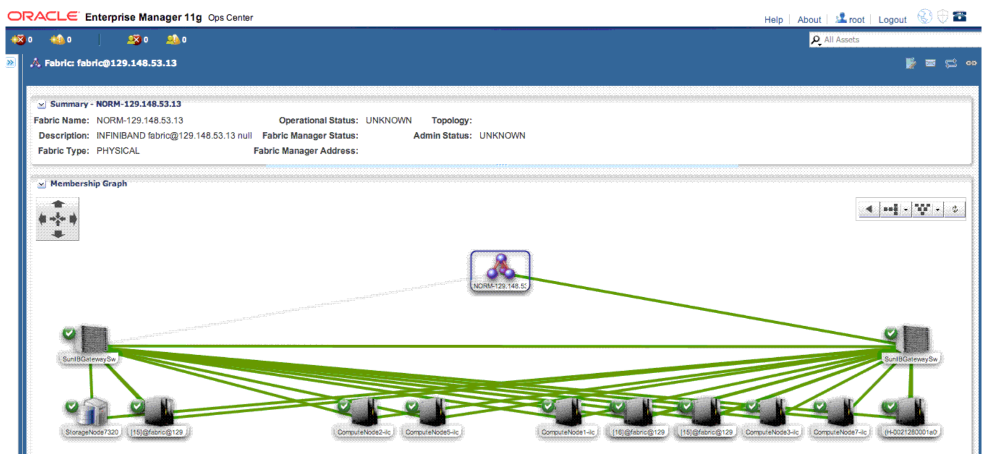

   When you collapse this fabric string, you will see compute nodes, gateway switches, and the storage appliance. You can select each of them to view information about a particular fabric node.

2. In the navigation pane, click a compute node under the `ComputeNodes` sub-group under the fabric. Information about that particular fabric node.
3. From this view, you can view detailed information about the fabric node, such as Summary, Hardware, Capabilities, Connectivity, Problems, Monitoring, Port Monitoring, Console, Jobs, Configuration, Charts, and Service Requests.
4. Similarly, on the left navigation pane, under the fabric drop-down list, click a gateway switch. Information about the switch as a fabric node is shown, as shown in Figure 17-18.

Figure 17-18   Viewing a Switch in the Fabric

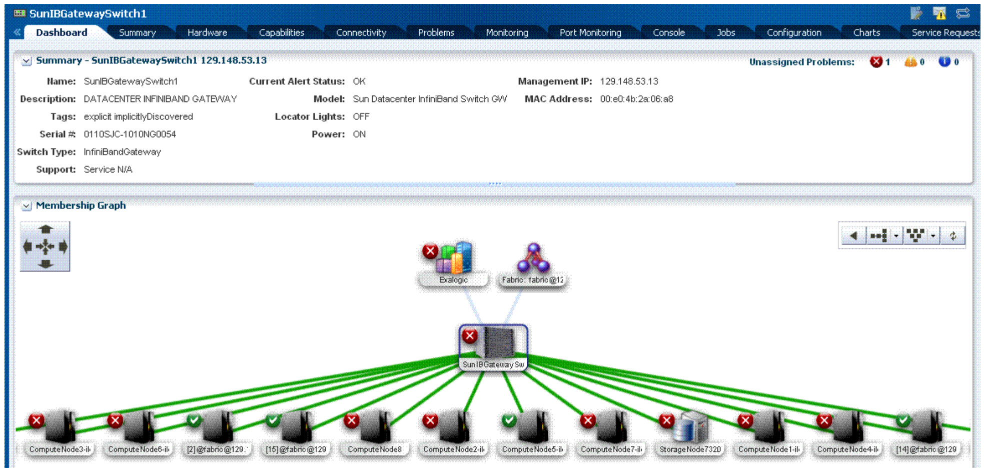

5. To monitor switch ports, click the **Port Monitoring** tab. Monitoring information about the switch ports is displayed, as shown in Figure 17-19.

Figure 17-19   Switch Port Monitoring

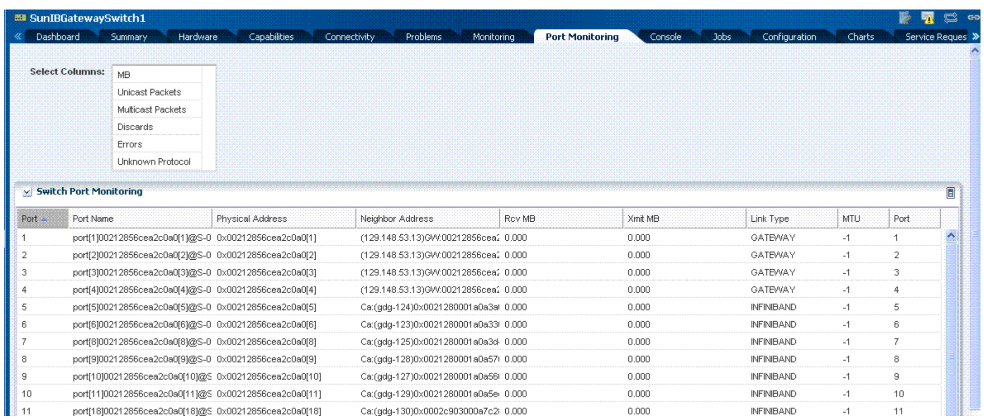

6. From this view, you can view detailed information about the fabric node, such as Summary, Hardware, Capabilities, Connectivity, Problems, Monitoring, Port Monitoring, Console, Jobs, Configuration, Charts, and Service Requests.

7. Similarly, on the left navigation pane, under the fabric drop-down list, click the storage appliance. Information about the storage appliance as a fabric node is shown, as shown in Figure 17-20.

Figure 17-20   Viewing the Storage Appliance in the Fabric

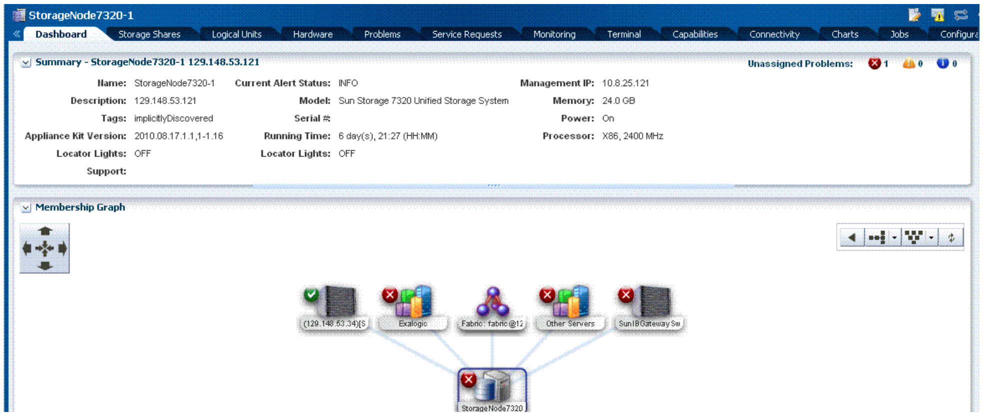

8. If you wish to launch the administration console for any of the fabric nodes (compute nodes, storage appliance, or gateway switches) from within the Oracle Enterprise Manager Ops Center interface, click the **Terminal** tab in the respective fabric dashboard pages.

17.14 About Problem Management

Problem management in Enterprise Manager Ops Center comprises several components that are designed to work together to simplify managing problems for the hardware assets in your Exalogic machine. The components include monitoring rules, suggested actions, and tools to automate problem identification and resolution.

When monitoring is enabled, it is connected with a problem management and notification system. Monitoring includes a standard set of monitoring rules and attributes, many of which are editable. In addition, you can add custom monitoring attributes and alert conditions.

A new alert is generated every time an attribute does not meet a monitoring rule. When an attribute for a managed asset or sub-asset type does not meet a monitoring rule, an alert is generated and is displayed as a problem in the Message Center. If an attribute exceeds a monitoring rule and then later meets the rule, the alert is automatically cleared. If the attribute does not meet the rule again, a new alert is generated.

> **Note:**
>
> For more information about problem management in Enterprise Manager Ops Center, see the "Problem Management" topic in the User section in the Oracle Enterprise Manager Ops Center 11g documentation. This topic describes the following tasks:
>
> - Viewing unresolved problems
> - Viewing problem details
> - Assigning a problem
> - Acknowledging problems
> - Adding an annotation
> - Displaying an annotation
> - Using maintenance mode
> - Taking action on a problem
> - Marking a problem repaired
> - Closing a problem

17.15 Using Monitoring Profiles and Rules

Monitoring rules, profiles, and plans detect components or attributes of a managed asset or resource that are not operating within specified parameters. A resource is a generic term for any resource managed through Enterprise Manager Ops Center, it can be an asset, a group, a network, or a library. An Enterprise Manager Ops Center administrator has permissions to edit and add monitoring rules and profiles.

The following are the main components of a complete monitoring configuration:

- **Monitoring Rules -** Express alerting conditions. You can apply one or more rules to an asset in order to monitor the asset and raise an alert when the monitoring rule condition is met.
- **Monitoring Profiles -** A set of monitoring rules targeted to a specific asset type. Default monitoring profiles contain a set of rules that are automatically applied. You can copy a profile and manually configure the rules in the profile.

> **Note:**
>
> For more information about monitoring profiles and rules, see the "Monitoring Profiles and Rules" topic in the Advanced User section in the Oracle Enterprise Manager Ops Center 11g documentation.

17.15.1 Creating a Monitoring Profile

You can create a new monitoring profile and edit the profile to add rules.

To create a monitoring profile, do the following:

1. Expand **Plan Management** in the navigation pane.
2. Click **Monitoring Profiles**.
3. Click **Create Profile** in the **Action** pane.
4. Provide a name and description for the monitoring profile, then select the resource type for the profile from the Subtype list.
5. Click **Finish** to save the profile. The new profile will appear in the center content pane.
6. (Optional) To add or remove rules or change monitoring parameters, double-click the profile in the center content pane.
7. (Optional) To make this profile the default monitoring profile, click the **Set as Default Profile** icon.

17.15.2 Adding a Monitoring Rule from the Asset View

To add a monitoring rule from the asset view, do the following:

1. Click **Assets** in the Navigation pane, expand the tree and click the asset to which you want to add the rule.
2. Click the **Monitoring** tab to see a list of all the monitoring rules.
3. Click the **Add Alert Monitoring Rule** icon in the center pane. The Add Alert Monitoring Rule Parameters window is displayed.
4. Select a Rule Type from the drop-down menu: **Threshold**, **Boolean Control**, **Enumerated Control** or **Expression**
5. Select an **Asset Type** from the drop-down menu.
6. Complete the **Monitored Attribute**. If you selected Expression, the Monitored Attribute option is not available.
7. Provide a name and description for the rule that will appear in the Profile Details page.
8. Define the monitoring schedule, either continuously or for a specific time period. The start and end times are based on the monitored asset's time zone.
9. Define how long the alerting condition must last to be considered an alert. The default setting is 5 minutes. You can change the amount of time and the unit of measurement to be either minutes, hours, or days.
10. Complete the Alert parameters for the different severity levels.
11. Use the **Immediate Action** field to define what action should take place when a problem is detected.
12. Click **Apply** to save the rule. The new rule will appear in the profile.

17.16 Using Reports in Oracle Enterprise Manager Ops Center

Reports in provide you with insight into all phases of the asset lifecycle. You can gather more detailed information about job history, firmware, OS updates, and then export that information to CSV or PDF output. Problem reports export to HTML.

You can create the following reports in Enterprise Manager Ops Center:

- Problem reports

  Problem reports summarize problem details for a specific managed asset or detailed information about specific problems. You can export these reports into an HTML output. The Summary Reports provide you with an historical account of the detected problems. You can create a report for a specific time period, for a specific severity level, status, type of problem, or the asset groups affected by the problem. These reports are invaluable in trend analysis and identifying patterns that you can then take steps to mitigate.

- Firmware compliance reports

  Firmware Compliance Reports enable you to maintain consistent firmware versions across your data center. You can associate one of your firmware profiles with the report, then run the Firmware Compliance Report to determine if the firmware on the asset complies with your firmware profile's specifications. If assets do not contain the firmware version identified in the profile, you can update the firmware from the report.

- Server provisioning reports

  Obtain a report of historical server provisioning actions. Run this simple ad-hoc report to obtain details about the Deployment Plan provisioning activities that occurred over a specified time period. Get specific information about the activity, including who ran the provisioning job, which profiles were selected, and the final outcome.

- System information reports

  Run a system information report to obtain the information about different resource types such as compute nodes, chassis, logical domains, global zone, and non-global zone. The system information report displays information about the current configuration of the targets.

> **Note:**
>
> For information about using reports, see the "Reports" topic in the User section in Oracle Enterprise Manager Ops Center 11g documentation. This topic describes the following tasks:
>
> – Creating charts
>
> – Creating a system information report
>
> – Updating reports
>
> – Creating problem reports

17.17 Using Oracle Services in Oracle Enterprise Manager Ops Center

Oracle Services enable you to view contract or warranty information and service requests for a specific asset. You can also view service requests that were filed as the result of an alert or problem in Enterprise Manager Ops Center, view service request details, and file a service request. By using Ops Center, you can display current contract and warranty information for a specific asset, or you can view the entitlements associated with your Oracle online account. An alert appears in the UI when a contract or warranty is about to expire.

You can quickly and easily create new service requests in the Oracle Enterprise Manager Ops Center UI. When you create a service request in the UI, you can view your requests and the requests of others in the UI.

> **Note:**
>
> You cannot display service requests filed outside of Oracle Enterprise Manager Ops Center. If you enter a service request outside of the UI, you must go to the Service Requests Home page on My Oracle Support to view the request and status.

This section discusses the following topics:

- Prerequisites for Using Oracle Services in Oracle Enterprise Manager Ops Center
- Viewing Service Requests
- Filing a Service Request

17.17.1 Prerequisites for Using Oracle Services in Oracle Enterprise Manager Ops Center

The following are the prerequisites for using Oracle Services in Oracle Enterprise Manager Ops Center:

- Register your assets with My Oracle Support
- Register your user as a My Oracle Support user in order to have access to the My Oracle Support database
- Run Enterprise Manager Ops Center in Connected Mode

To access the database, your user must be registered as a My Oracle Support user. This is the same account that is used to access My Oracle Support at http://support.oracle.com.To register an Oracle.com account, go to https://myprofile.oracle.com/EndUser/faces/profile/createUser.jspx?tid=minimal&nextURL=https://support.oracle.com:443/CSP/ui/flash.html?login.

To determine if you are running in Connected Mode and have access to My Oracle Support, view the following connection icons in the upper right corner of the UI:

- World icon indicates Internet connection
- Shield icon indicates connectivity to the Oracle Knowledge Base
- Phone icon indicates the connection status to My Oracle Support Services

  If the icons do not contain color, you are not connected. The following graphic shows how the icons appear when you are connected.

17.17.2 Viewing Service Requests

To view the service requests, complete the following steps:

1. In the left navigation pane, click **Message Center**.
2. Click **Open Service Requests**, **My Service Requests**, or **Service Requests Opened by Others** to display a list of requests.
3. To view details of a particular service request, highlight a row, then click the **View Service Request** icon. The following graphic shows an example service request.

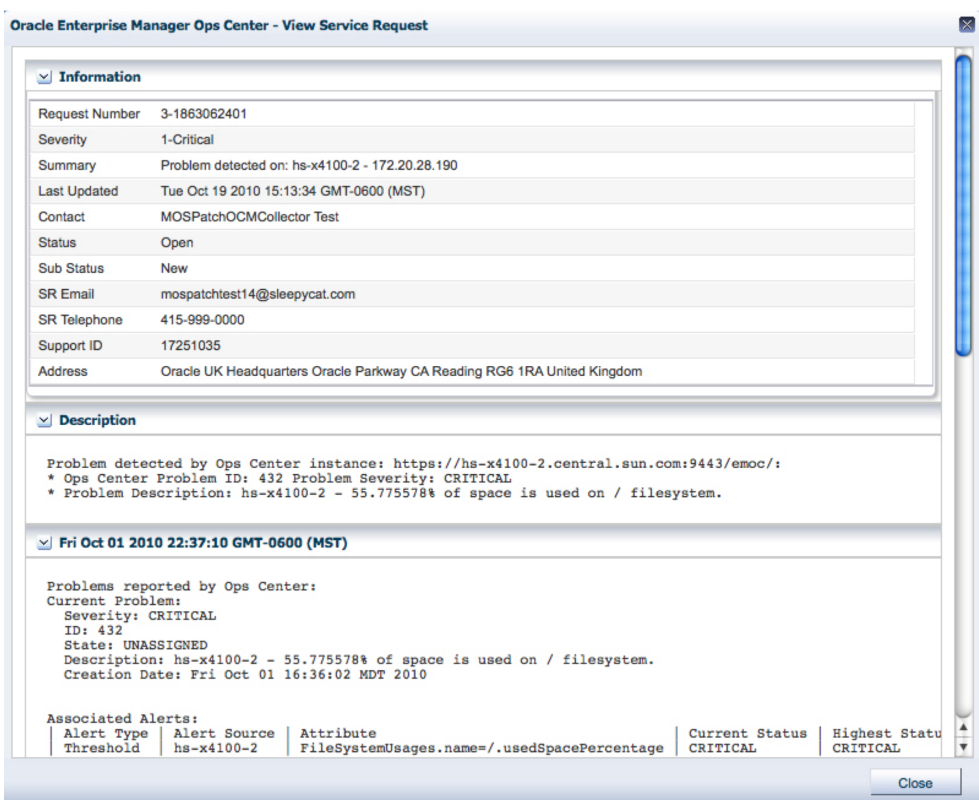

Note:

You can also view all open service requests associated with a particular asset. To do so, select an asset under **Assets** in the navigation pane, click the **Service Requests** tab in the content pane.

17.17.3 Filing a Service Request

When your assets are associated with a contract and registered in the Oracle database, you can create and file a service request in Ops Center from a problem or from an asset. See Prerequisites for Using Oracle Services in Oracle Enterprise Manager Ops Center for requirements that must be met before successfully filing a service request ticket.

If the asset is not registered in My Oracle Support, the service request job will fail.

To file a service request from a problem, do the following:

1. Click **Message Center** in the Navigation pane.
2. Click **My Problems** or **Unassigned Problems**.
3. Highlight the problem, then click the phone icon in the center pane.

To file a service request from an asset, do the following:

1. Click the server or hardware in the **Assets** section of the Navigation pane.
2. Click **Open Service Request** in the Actions pane. The action is disabled if the connection cannot be established.

18 Install the ASR Software

This chapter describes how to install Auto Service Request (ASR) for Exalogic machines.
It contains the following topics:

18.1 About Oracle Auto Service Request (ASR)

Auto Service Request (ASR) is a secure, scalable, customer-installable software feature of warranty and Oracle Support Services that provides auto-case generation when common hardware component faults occur. ASR is designed to enable faster problem resolution by eliminating the need to initiate contact with Oracle Support Services for common hardware component failures, reducing both the number of phone calls needed and overall phone time required. ASR also simplifies support operations by using electronic diagnostic data. Easily installed and deployed, ASR is completely controlled by you, the customer, to ensure security. ASR is applicable only for component faults. Not all component failures are covered, though the most common components (such as disk, fan, and power supplies) are covered.

> **Note:**
>
> ASR is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts, within your 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.

18.2 Recommended Configuration

The recommended configuration is to install the ASR Manager, which receives fault telemetry information from Exalogic compute nodes, on a system outside of Exalogic. This machine can also be a virtual machine that runs Oracle Solaris or Oracle Linux.

Although it is possible to install ASR Manager on one of the Exalogic compute nodes, it is recommended that you install it outside of Exalogic.

18.3 Before You Begin

Before installing ASR, ensure that the following conditions are met:

- Ensure that you have access to My Oracle Support (`http://support.oracle.com`) and your contact information is correct and current.
- Make sure that all of your Exalogic machine assets have a Contact assigned and that the contact is correct and current.
- Identify and designate a system, outside of Exalogic, to serve as ASR Manager.
- Identify and verify ASR assets.
- Ensure connectivity to the Internet using HTTPS.

> **Note:**
>
> For Exalogic compute nodes, ASR only covers what the ILOM monitors.

18.4 Prerequisites for Installing ASR Manager

The following are the prerequisites for installing ASR Manager on a system outside of Exalogic:

- Root access to the system on which ASR Manager will be installed
- Operating system: Oracle Linux 5.3 and later, or Oracle Solaris 10, Update 6 (10u6) and later
- Java version: at least JRE/JDK 1.6.0_04

18.5 Install ASR Manager on a Standalone System

Installing ASR Manager on a standalone system (a machine outside of Exalogic running Oracle Linux or Oracle Solaris) involves the following steps:

1. Install Service Tags for Oracle Linux
2. Install SASM Package
3. Install ASR Package

18.5.1 Install Service Tags for Oracle Linux

To install service tags for Oracle Linux, complete the following steps:

1. If necessary, download Sun Service Tags 1.1.5 from the following URL, and unzip:

   `https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=SYSTEMDOC&id=1185493.1`

2. Run the following commands to install the service tags:

```
rpm -i sun-servicetag-1.1.5-1.i386.rpm
```

```
rpm -i sun-hardware-reg-1.0.0-1.i386.rpm
```

18.5.2 Install SASM Package

To install the SASM package, complete the following steps:

1. Run the following command, as a `root` user, to verify you have version 1.2.1 or later:
   - On Oracle Linux:

     `# rpm -q SUNWsasm`
   - On Oracle Solaris:

     `# pkginfo -l SUNWsasm`

   If required, download SASM 1.2.1 from the following URL:

   https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=SYSTEMDOC&id=1185493.1
2. As `root`, run the following command to install the SASM package:
   - On Oracle Linux:

     `# rpm -i SUNWsasm.`*version_num*`.rpm`
   - On Oracle Solaris:

     `# pkgadd -d SUNWsasm.`*version_num*`.pkg`

18.5.3 Install ASR Package

To install the ASR package, complete the following steps:

1. Download the ASR package from the following URL, and unzip:

   https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=SYSTEMDOC&id=1185493.1
2. As `root`, run the following commands to install the ASR package:
   - On Oracle Linux:

     `# rpm -i SUNWswasr.`*version_num*`.rpm`
   - On Oracle Solaris:

     `# pkgadd -d SUNWswasr.`*version_num*`.pkg`
3. Add the `asr` command to the `PATH` (update to the root's `.profile`, `.cshrc`, `.kshrc`, or `.bashrc` as needed):

   ```
   PATH=$PATH:/opt/SUNWswasr/bin/asr export PATH
   ```

18.6 Register the ASR Manager

To register the ASR Manager, see the topic Register the ASR Manager in *Oracle® Auto Service Request Installation and Operations Guide*.

18.7 Activate ILOM for Exalogic Compute Nodes

To activate the ILOM for Exalogic compute nodes, complete the following steps:

1. Activate the ASR Manager host machine by running the following command:

   `asr activate_asset -i` *IP_address_of_ASM_Manager_host_machine*

2. Activate ILOM for Exalogic compute nodes individually by running the following command:

   `asr activate_asset -i` *ILOM_IP_address_of_compute_node*

> **Note:**
>
> For a list of default ILOM IP addresses assigned to Exalogic compute nodes, see Default IP Addresses and Ports.

18.8 Activate the Storage Appliance

To activate the storage appliance included in your Exalogic machine, complete the following steps:

1. Ensure that you have a valid My Oracle Support Account.
2. Ensure that you can connect to the Internet from your storage appliance using either a direct connection or a web proxy.
3. In a web browser, enter the IP address or host name you assigned to the `NET0` port of either storage head as follows:

   ```
   https://ipaddress:215
   ```

   or

   ```
   https://hostname:215
   ```

   The login screen appears.
4. Type `root` into the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen is displayed.
5. Click the **Configuration** tab, and click **SERVICES**. On the left navigation pane, click **Services** to display the list of services.
6. Click **Phone Home**, as shown in Figure 18-1.

Figure 18-1 Phone Home Tab on the List of Services

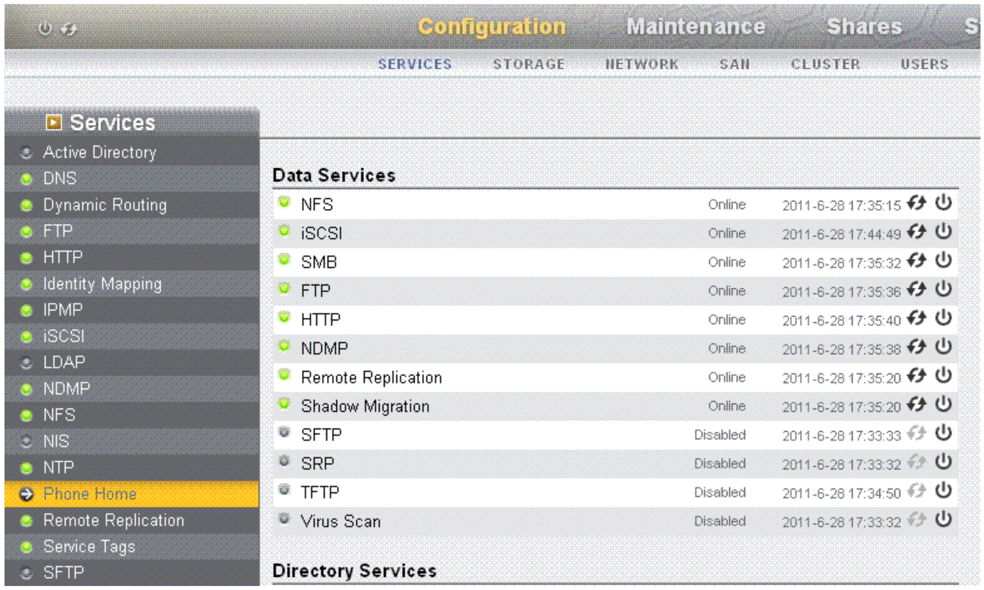

When you click **Phone Home**, the Phone Home page is displayed, as shown in Figure 18-2.

Figure 18-2 Phone Home Settings Page

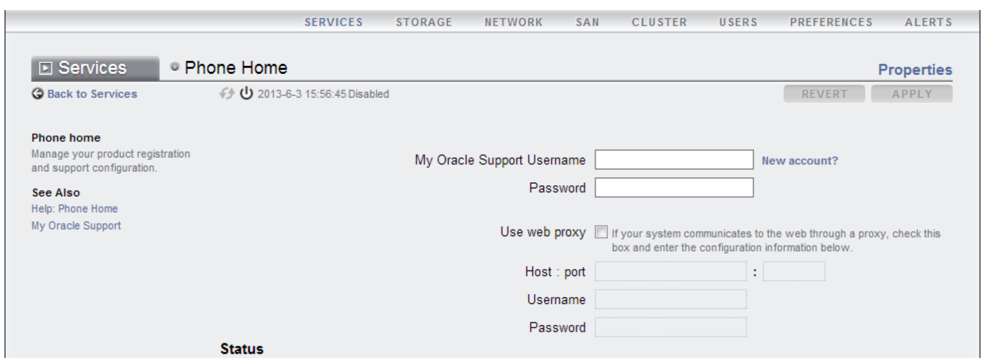

7. In the **My Oracle Support Username** field, enter your username.
8. In the **Password** field, enter the password for your My Oracle Support account.
9. If you are using a web proxy to connect to the Internet from the storage appliance, select the **Use web proxy** option, and enter the following information:
   - In the **Host: port** field, enter the complete host name of your web proxy server and the port.
   - In the **Username** field, enter your user name for the accessing the web proxy server.
   - In the **Password** field, enter your password.
10. Click the **Apply** button.
11. When the Service Enable / Disable popup is presented, select the **Enable** option.

> **Note:**
>
> For information about network configuration for the storage appliance, see Network Configuration Options.
>
> ASR activation for the storage appliance covers ILOM on the storage server heads.

18.9 Approve and Verify ASR Activation for Exalogic Machine Assets

This section contains the following topics:

- Approve Exalogic Machine Assets in My Oracle Support
- View and Verify ASR Assets

18.9.1 Approve Exalogic Machine Assets in My Oracle Support

To approve Exalogic machine assets in My Oracle Support, complete the following steps:

1. Log in to My Oracle Support.
2. On the home page, click the **More...** tab, and click **Settings**.
3. In the Settings left-hand section, click **Pending ASR Activations**.
4. Select the Exalogic machine hardware component you want to activate and review the information on the ASR Activation – Asset screen, as shown in Figure 18-3. You may need to update information about your Exalogic machine, as required.

Figure 18-3 ASR Activation - Asset Screen in My Oracle Support

5. Click **Approve** to complete activation.

> **Note:**
>
> An Exalogic machine asset must be in an active ASR state in My Oracle Support in order for Service Request autocreate to work.

18.9.2 View and Verify ASR Assets

To view and verify ASR assets (Exalogic machine components whose ASR is active), complete the following steps:

1. On the standalone system where ASM Manager is running, run the following command to verify the status of your Exalogic machine assets:

   ```
   list_asset
   ```

   This command should list ASR assets in your Exalogic machine, including compute nodes and storage server heads.
2. Log in to My Oracle Support. You must see the same set of ASR assets with their correct serial numbers.

A

Configure the Exalogic Machine Using ECU

This appendix provides an overview of the Exalogic Configuration Utility (ECU).
It contains the following topics:

- Overview
- Important Notes Before You use the ECU
- Configuration Tasks

A.1 Overview

The ECU enables you to perform the initial configuration of an Exalogic machine and to connect the machine to your existing network. The ECU does not perform all of the initial configuration steps. For more information, see Initial Network Configuration of Exalogic Machine.

The ECU verifies the current state of your Exalogic machine before performing any initial configuration.

During manufacturing, an Exalogic machine is configured with the following default network topology:

- Compute nodes and storage heads with `NET0`, `ILOM`, and `BOND0` (IPoIB) interfaces
- Sun Network QDR InfiniBand Gateway Switches with the `ILOM` interface

> **Note:**
>
> Sun Datacenter InfiniBand Switch 36, which is included in Exalogic full- and half-rack machines, is not connected or configured, by default. This switch is used in multirack configuration scenarios only. For more information, see Use the Sun Datacenter InfiniBand Switch 36 in Multirack Configurations.
>
> In addition, the Cisco Ethernet management switch is not configured, by default.
>
> For information about configuring VNICs for Ethernet connectivity, see Configure Ethernet Over InfiniBand.

A.2 Important Notes Before You use the ECU

Before using the ECU to configure the Exalogic machine, keep the following points in mind:

- Ensure that you have connected a laptop to the Exalogic machine

For information about connecting a laptop to your Exalogic machine, see Connect a Laptop to the Exalogic Machine.

- Validate the desired network configuration

  Review Default State of the Exalogic Machine Network Configuration and Understand Network Requirements and Configuration to identify the desired network topology of your Exalogic machine. In addition, review the network topology with your network administrator.

- Ensure that all required network parameters are available

  After determining the desired network configuration, you should determine a full set of network parameters, such as gateway addresses, netmask, IP addresses for `NET0`, `BOND0`, `BOND1` (if needed), ILOMs, and so on. You are required to enter them in the configurator spreadsheet, which is used as input by Exalogic Configuration Utility.

- Ensure that all the required cabling is completed

  You must have connected your Exalogic machine to the data center's Ethernet device, such as a switch. For more information, see Connectivity Between Exalogic Machine and External LAN Through Sun Network QDR InfiniBand Gateway Switch.

- Set up VNICs, as necessary, for Ethernet over InfiniBand (EoIB) connectivity. For more information, see Configure Ethernet Over InfiniBand.

A.3 Configuration Tasks

Oracle recommends strongly that the configuration tasks using ECU be performed by fully trained, qualified Oracle personnel or by formally accredited Oracle partners. For more information, contact Oracle Advanced Customer Support (ACS):

Email: acsdirect_us@oracle.com

Website: `http://www.oracle.com/acs`

B
Site Checklists

Complete the checklists in this appendix to ensure that the site is prepared for installing and deploying Exalogic machines.
This appendix contains the following topics:

- System Components Checklist
- Data Center Room Checklist
- Data Center Environment Checklist
- Access Route Checklist
- Facility Power Checklist
- Power Checklist
- Safety Checklist
- Logistics Checklist
- Network Specification Checklist
- Reracking Checklists

B.1 System Components Checklist

Complete the following checklist to ensure that the systems component considerations have been addressed.

System Components Checklist Items	Yes	No	N/A	Comment
Is it an Exalogic machine full rack, Exalogic machine half rack, Exalogic machine quarter rack, or Exalogic machine eighth rack?				
How many racks will be installed?				
If more than one rack will be installed, then was the cable upgrade kit ordered?				
Is the rack installation a new system or an addition to an existing system?				
Is the system going to connect to another Exalogic machine?				
Was the multirack cabling service purchased?				
Are all Exalogic machine racks adjacent to each other?				

System Components Checklist Items	Yes	No	N/A	Comment
If the connecting racks are not within the specified proximity, then have the following been done: • Purchased approved longer InfiniBand cables from an approved third-party provider? • Asked Oracle Support Services to provide and schedule the custom multirack cabling service with the installation?				
Will the Ethernet switch be replaced? If yes, then have you arranged for the replacement switch installation?				

B.2 Data Center Room Checklist

Complete the following checklist to ensure that the data center room requirements are met. For information about the data center requirements, see Flooring Requirements.

Data Center Room Checklist Items	Yes	No	N/A	Comment
Has the Exalogic machine location been allocated?				
Is there a vacant location for the new equipment?				
Does the floor layout meet the equipment maintenance access requirements?				
Is there adequate space available for maintenance?				
Will the equipment be positioned so that the exhaust air of one rack does not enter the air inlet of another rack?				
Have cabinet stabilization measures been considered?				
Does the raised floor satisfy the weight requirements for the new hardware?				
Can floor tiles be removed without permission to accommodate service?				
Are there cable routing channels or cutouts?				
Are you providing any additional hardware?				
Is the hardware you are providing fully compatible with the Exalogic machine?				
Will the new hardware location require any non-standard cable lengths?				
Is the floor to ceiling height a minimum of 2914 mm or 2.9 m (9.6 feet)?				

Data Center Room Checklist Items	Yes	No	N/A	Comment
Is the depth of the raised floor a minimum of 46 cm (18 inches)?				

B.3 Data Center Environment Checklist

Complete the following checklist to ensure that the data center environment requirements are met. For information about environment requirements, see Temperature and Humidity Requirements.

Data Center Environment Considerations	Yes	No	N/A	comment
Does the computer room air handling meet temperature and humidity requirements?				
Does the installation floor layout satisfy the ventilation requirements?				
Will the equipment be positioned so the exhaust air of one rack does not enter the air intake of another rack?				
Are the perforated floor tiles each rated at 400 CFM or greater?				
Do the data center air conditioners provide sufficient front to back airflow?				
Is airflow adequate to prevent hot spots?				
Can the data center continuously satisfy environmental requirements?				
Can more vented floor tiles be obtained if required?				

B.4 Access Route Checklist

Complete the following checklist to ensure that the access route requirements are met.

Access Route Considerations	Yes	No	N/A	Comment
Has the access route been checked for clearances of the packaged equipment?				
Do all the doors and entry ways conform to the width and height requirements for transportation, including the width of the unpacked unit?				
Do all the doors meet the height requirement of minimum 86 inches for packaged delivery?				
Does the access route provide sufficient space for transport of the packed devices?				

Access Route Considerations	Yes	No	N/A	Comment
Are there any ramps or thresholds that are of concern? If yes, then provide details.				
Are there any stairs or ramps in the moving patch for the new hardware?				
Have you confirmed that all route incline angles are within the permitted range?				
Have you confirmed that the access route is free of any obstacles that would expose the device to shock?				
Are all the surfaces acceptable for rolling the new unpacked and packed equipment?				
If a pallet jack is to be use, then have you confirmed the following: • The pallet jack supports the device weight? • The pallet jack tines are compatible with the shipping pallet?				
If there are stairs, then is a loading elevator accessible for the equipment?				
If an elevator is to be used, then have you confirmed the following: • The elevator car is wide enough for the device to be carried into it? • The elevator car is high enough for the device to be carried into it? • The load limit of the elevator is greater than the device weight? • Are elevators available to handle up to 1049.09 kg (2308 lbs) fully-loaded rack capacity? • The elevator door meets the minimum height requirement of 86 inches for packaged rack delivery?				
Does the path from the receiving location to the designated data center area support the weight of the unpacked equipment?				
Is the path onto the raised floor rated for dynamic loading of the server? Refer to Flooring Requirements for requirements.				

B.5 Facility Power Checklist

Complete the following checklist to ensure that the facility power requirements are met. For information about power requirements, see Electrical Power Requirements.

Facility Power Considerations	Yes	No	N/A	Comment
Do you know the required operating voltage and electric current level of the device and peripherals?				
Will you be using single-phase (low-voltage or high-voltage) or 3-phase (low-voltage or high-voltage) power?				
Are enough power outlets provided within 2 meters for each rack?				
Do the power outlets have appropriate socket receptacles for the PDU option ordered? Options are low voltage or high voltage, single-phase or 3-phase.				
Will optional ground cables be attached to the rack?				
Are the circuit breakers for the equipment suitable in terms of voltage and current-carrying capacities?				
Does the power frequency meet the equipment specifications?				
Are power outlets available for the new equipment at the designated location?				
Will system power be delivered from two separate grids?				
Is there a UPS to power the equipment?				
Do you have the minimum required power sources to support the power load for the new hardware? Use kilowatt (kW) /kilovolt (kVA) to express power load.				

B.6 Power Checklist

Complete the following checklist to ensure that the power requirements are met for Sun Oracle Database Machine. For information about power requirements, see Electrical Power Requirements.

Power Checklist Considerations	Yes	No	N/A	Comment
Do you have the minimum required power sources?				
Are power outlets available for the new equipment at the designated location?				
Does the power frequency meet the equipment specifications?				
Is there a UPS to power the equipment?				
Is the capacity of the UPS sufficient for the Exalogic machine?				

Power Checklist Considerations	Yes	No	N/A	Comment
Does the power source support the power load for the new hardware? Use kilowatt (kW)/kilovolt (kVA) to express power load.				

B.7 Safety Checklist

Complete the following checklist to ensure that the safety requirements are met. For information about safety, see Operational Procedures for Exalogic Machines, and Temperature and Humidity Requirements.

Safety Checklist Considerations	Yes	No	N/A	Comment
Is there an emergency power shut off?				
Is there a fire protection system in the data center room?				
Is the computer room adequately equipped to extinguish a fire?				
Is antistatic flooring installed?				
Is the floor below the raised floor free of obstacles and blockages?				

B.8 Logistics Checklist

Complete the following checklist to ensure that the logistics requirements are met. For information about unpacking and space requirements, see Space Requirements.

Logistics Checklist Considerations	Yes	No	N/A	Comment
Do you have contact information for the data center personnel?				
Is there security or access control for the data center?				
Are there any security background checks or security clearances required for vendor personnel to access the data center? If yes, then do you have a recommended agency?				
How many days in advance must background checks be completed?				
Are there any additional security access issues?				
Is computer room access available for installation personnel?				
Are laptops, cell phones, and cameras allowed in the data center?				
Does the building have a delivery dock?				

Logistics Checklist Considerations	Yes	No	N/A	Comment
Is there a delivery/unpacking/staging area?				
Is the delivery inside?				
If the delivery is not inside, then is the site prepared for uncrating?				
Is the unpacking/staging area protected from the elements?				
Does the building have adequate receiving space?				
Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?				
Will sufficient moving personnel be available to install the hardware?				
Is union labor required for any part of the delivery or installation?				
Are you prepared for uncrating and trash removal?				
Is uncrating of cabinet and cabinet trash removal required?				
Are there any restrictions on delivery truck length, width or height?				
Does the customer allow cardboard boxes and other packing material in the computer room? If no, then do ground level deliveries require a truck with a side rail lift?				
Is there a time constraint on dock access? If yes, then provide time constraints.				
Is tail lift required on delivery carrier to unload the equipment at the delivery dock?				
Will any of the following be required to place equipment in computer room? • Stair walkers • Lifters • Ramps • Steel plates • Floor covers				
Does the delivery carrier require any special equipment, such as non-floor damaging rollers, transport dollies, pallet jacks or fork lifts?				

B.9 Network Specification Checklist

Complete the following checklist to ensure that the network specification requirements are met. For information about IP addresses, see Understand Network Requirements and Configuration.

> **Note:**
>
> - If you are deploying multiple Exalogic machines that will be connected to form a single system, then contact your Oracle representative for instructions.
> - By default, there is one default InfiniBand partition at the Exalogic machine level. All compute nodes in the Exalogic machine are members of the default InfiniBand partition. The most common model for application isolation involves multiple IP subnetting, in which the most mission-critical applications are assigned their own IP subnets layered above the default IPoIB link. In this model, some subnets may also contain applications that have less stringent or otherwise different resource requirements.

Network Specification Considerations	Yes	No	N/A	Comment
Did you complete all the networking worksheets, which is included in Network Configuration Worksheets, and provide it to your Oracle technical representative?				
Have you received the site-specific installation template from your Oracle technical representative?				
Did you review the installation template and consult with your Oracle technical representative regarding site-specific changes, if any?				
Did you verify the IP addresses in the installation template are currently not in use?				
Have you performed the required configuration within your network infrastructure to allow the Exalogic machine to use the IP addresses specified in the template?				
Have you registered IP addresses in the installation template with DNS?				
Did you run the required network cables from your network equipment to the location where the Exalogic machine will be installed?				

Network Specification Considerations	Yes	No	N/A	Comment
Did you label the network cables that will connect to the Exalogic machine?				

B.10 Reracking Checklists

Reracking of Exalogic half racks, Exalogic quarter racks, or Exalogic eighth racks is allowed when the customer site requires a specialized infrastructure. Exalogic full racks cannot be reracked due to potential air flow and cooling issues.

Note:

- Customer must purchase both the Oracle Reracking Service and Oracle Installation Service.
- Oracle does not provide support for customer-supplied equipment.

Complete the following checklist prior to reracking an Exalogic machine:

Reracking Considerations	Yes	No	N/A	Comment
Has the customer purchased Oracle Reracking Service?				
Has the customer purchased the Oracle Installation Service?				
Is there a cart capable of carrying the weight of the servers to move the components and associated cabling from the supplied rack to the customer supplied rack?				
Do the target rack dimensions meet the following requirements? • Height: 42RU (shorter racks can be supported with customer provided PDUs) • Width: 600mm • Depth: 1112mm without front and rear doors				
Is the target rack at least 30 RU tall and does it include the customer provided PDUs?				
If the rack is less than 42 RU, then the rack must be at least 30 RU tall and the customer must provide compatible PDUs to install in the target rack.				

Reracking Considerations	Yes	No	N/A	Comment
Is the distance between the front and rear mounting planes between the minimum of 610 mm and the maximum 915 mm (24 inches to 36 inches)?				
Is the clearance depth in the front of the front mounting plane (distance to the front cabinet door) at least 25.4 mm (1 inch)?				
Does the target rack meet the following minimum load capacity? • 19kg/RU • 785 kg total				
Is the rack a four-post rack (mounting at both front and rear). **Note**: Two-post racks are not compatible.				
Does the target rack's horizontal opening and unit vertical pitch conform to ANSI/EIA 310-D-1992 or IEC 60927 standards?				
Does the target rack have RETMA rail support? **Note**: Exalogic racks require 19 inches for RETMA rail spacing width. The minimum rack width of 600 mm is recommended to accommodate the PDU and cable harnesses on the side. If the rack is less than 600 mm wide, then it must have additional depth to accommodate mounting behind the server CMAs.				
Does the target rack support Oracle cable management arms (CMAs)?				
Does the target rack support installation of Oracle vented and solid filler panels?				
Can the target rack provide tie-downs along the left rear side of the rack to support the InfiniBand cables? The side is left rear side when viewed from the front of the rack.				
Can the target rack provide tie-downs for the Ethernet wiring harness?				
Is there sufficient space for the cable harnesses and the PDUs in the target rack?				
Can a label with the Exalogic rack serial number be printed and attached to the target rack?				

The following checklist is specific to the power distribution units (PDUs) when using the reracking service.

PDU Reracking Considerations	Yes	No	N/A	Comment
Does the target rack support installation of standard Oracle PDUs? If not, then complete this checklist.				
Can the customer provide an equivalent pair of PDUs?				
Can the customer provide two PDUs with capacity of 10kVA per PDU?				
Can the customer provide at least the following number of 10A C13 plugs per PDU: • Half rack: 46 cables • Quarter rack: 28 cables • Eighth rack: 20 cables				
Can the customer provide a single PDU and its circuits to support the power requirements in case one PDU fails?				
Can the customer ensure power loads are evenly distributed across all circuits of a single PDU?				
Can the customer provide appropriate power drops for the PDUs?				

B.10.1 Rack Recycling

The following checklist is specific to recycling the rack:

Recycling Considerations	Yes	No	N/A	Comment
Does the customer want Oracle to take back and recycle the empty rack? If yes, then the shipping company can pack the empty rack for recycling.				
If Oracle is recycling the rack, then has the customer completed the appropriate application at the Oracle Electronic Waste and Product Return website? The website is at `http://www.oracle.com/us/products/applications/green/waste-and-product-returns-185031.html`				

C

Cabling Diagrams

This appendix contains cabling diagrams for each of the Exalogic machine configurations. This appendix has the following sections:

- Exalogic Eighth Rack
- Exalogic Quarter Rack
- Exalogic Half Rack
- Exalogic Full Rack

NOT_SUPPORTEDNOT_SUPPORTED

For information about connecting an Exalogic machine with one or more Exalogic or Exadata machines, see the *Exalogic Elastic Cloud Multirack Cabling Guide.*

C.1 Exalogic Eighth Rack

Figure C-1 illustrates the cabling for a single Exalogic eighth rack.

Figure C-1 Exalogic Eighth Rack Network Connectivity

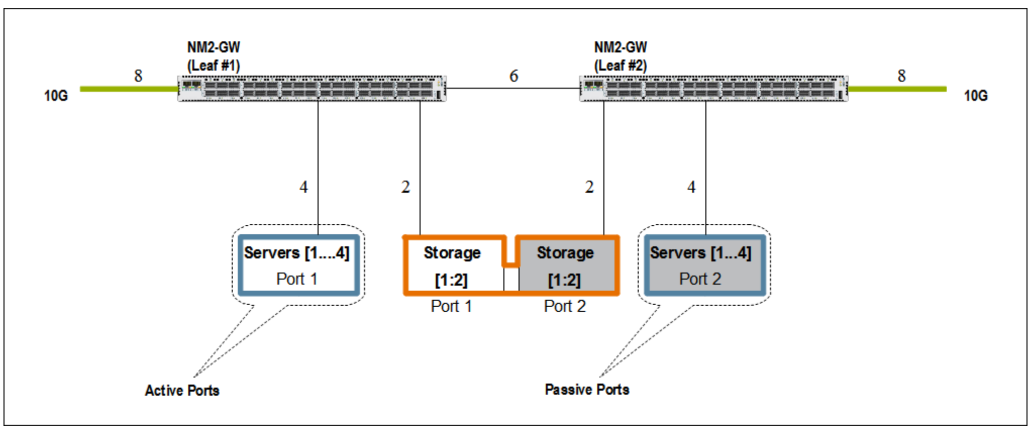

In Figure C-1, **NM2-GW** represents a Sun Network QDR InfiniBand Gateway Switch (leaf switches). **Servers** refers to compute nodes in the Exalogic machine. Half of the compute nodes are connected with their active ports to the first gateway switch and their passive ports to the second gateway switch. The remaining half of the compute nodes are connected with their active ports to the second gateway switch and their passive ports to the first gateway switch. This connection is for high availability and load distribution purposes.

Exalogic machine eighth racks do not contain the Sun Datacenter InfiniBand Switch 36.

C.2 Exalogic Quarter Rack

Figure C-2 illustrates the cabling for a single Exalogic quarter rack.

Figure C-2 Exalogic Quarter Rack Network Connectivity

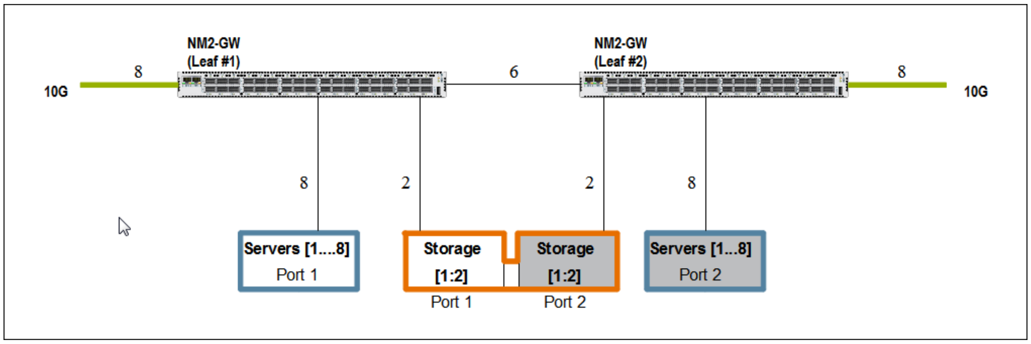

In Figure C-2, **NM2-GW** represents a Sun Network QDR InfiniBand Gateway Switch (leaf switches). **Servers** refers to compute nodes in the Exalogic machine. Half of the compute nodes are connected with their active ports to the first gateway switch and their passive ports to the second gateway switch. The remaining half of the compute nodes are connected with their active ports to the second gateway switch and their passive ports to the first gateway switch. This connection is for high availability and load distribution purposes.

Exalogic machine quarter racks do not contain the Sun Datacenter InfiniBand Switch 36.

C.3 Exalogic Half Rack

Figure C-3 illustrates the cabling for a single Exalogic half rack.

Figure C-3 Exalogic Half Rack Network Connectivity

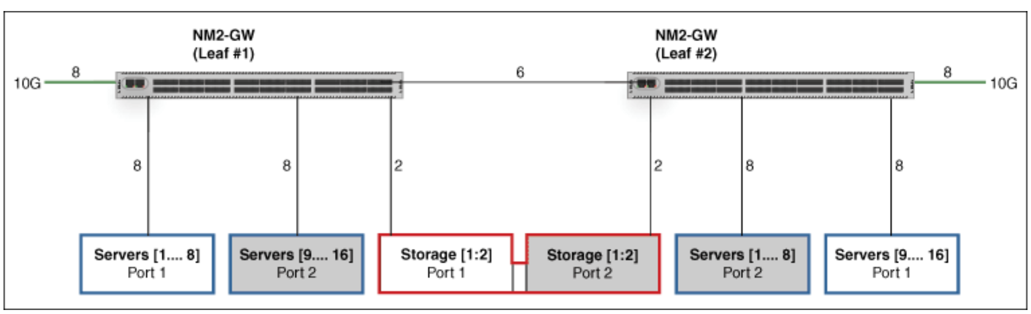

In Figure C-3, **NM2-GW** represents a Sun Network QDR InfiniBand Gateway Switch (leaf switches). **Servers** refers to compute nodes in the Exalogic machine.

C.4 Exalogic Full Rack

Figure C-4 illustrates the cabling for a single Exalogic full rack.

Figure C-4 Exalogic Full Rack Network Connectivity

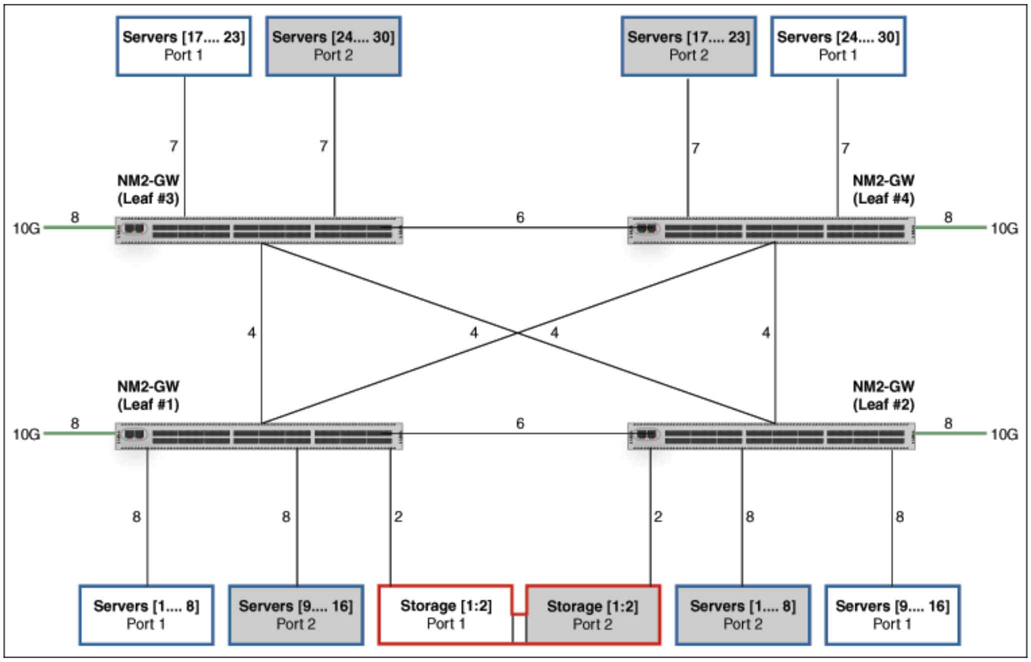

In Figure C-4, **NM2-GW** represents a Sun Network QDR InfiniBand Gateway Switch (leaf switches). **Servers** refers to compute nodes in the Exalogic machine.

D
Replacement Units

This appendix lists the replacement units for the Exalogic machine.
There are two types of replacement units, FRUs (field replaceable units), and CRUs (customer replaceable units). FRUs are installed by trained Oracle field technicians. CRUs are installed by the customer.

> **Note:**
>
> The Exalogic machine is accompanied by a spares kit containing additional parts and accessories (cables, for example) that Oracle Services personnel will use to replace non-working and broken parts quickly. When Oracle Services personnel visit your site for repairing and replacing hardware, you should make the spares kit available to them. Note that any parts used by Oracle Service personnel from the spares kit will be replenished by Oracle.

This appendix contains the following topics:

- Rack-Level FRUs for Exalogic X5-2
- Rack-Level FRUs for Exalogic X4-2
- Rack-Level FRUs for Exalogic X2-2 and X3-2
- Parts for ZS3-ES Storage Appliance for Exalogic X5-2
- Parts for ZS3-ES Storage Appliance for Exalogic X4-2
- Parts for Sun ZFS Storage 7320 Appliance
- Parts for Oracle Server X5-2 Compute Nodes
- Parts for Sun Server X4-2 Compute Nodes
- Parts for Sun Server X3-2 Compute Nodes
- Parts for X4170 M2 Compute Nodes
- Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X5-2
- Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X4-2
- Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X2-2 and X3-2
- Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X5-2
- Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X4-2
- Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X2-2 and X3-2
- Parts for the Cisco Catalyst 4948E-F-S Switch
- Parts for the Cisco Catalyst 4948 Switch
- Parts for the Gari DE2-24C Disk Enclosure for Exalogic X5-2

- Parts for the Gari DE2-24C Disk Enclosure

D.1 Rack-Level FRUs for Exalogic X5-2

Table D-1 lists the rack-level replaceable parts for Exalogic X5-2 machines.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-1 Rack-Level Replacement Parts for EL X5-2

Part Number	Description	Coverage under Basic Warranty
7015399	Pwrcord, Jmpr, SR2, 1m, C14RA, 10A, C13	FRU
7015400	Pwrcord, Jmpr, SR2, 2m, C14RA, 10A, C13	FRU
530-4436-01	CBL, 7' ETHRNT(CAT5/CAT5E) BLACK	FRU
530-4502-01	CBL, 10' ETHRNT(CAT5/CAT5E) BLUE	FRU
530-4446-01	5M, 10GE QSFP Passive Copp, CAB	FRU
350-1546-01	ASSY, 1U CMA KIT, EXADATA, X4170	FRU
371-4919-01	ASY, SLIDE-RAIL, LONG, 1U-2U, STAN	FRU
135-1204-01	XCVR, 40GBps, SR, QSFP-MTO	FRU

D.2 Rack-Level FRUs for Exalogic X4-2

Table D-2 lists the rack-level replaceable parts for Exalogic X4-2 machines.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-2 Rack-Level Replacement Parts for EL X4-2

Part Number	Description	Coverage under Basic Warranty
7079238	Base Rack	FRU
7015400	JUMPER CABLE, 2 M, C13-C14	FRU
530-4445-01	3M, 10GE QSFP Passive Copp, CBL	FRU
530-4502-01	CBL, SHIELD RJ45 CAT5E, 1M	FRU
530-4436-01	CBL, 7FT ETHERNET(CAT5/CAT5E) BLACK	FRU

Table D-2 (Cont.) Rack-Level Replacement Parts for EL X4-2

Part Number	Description	Coverage under Basic Warranty
530-4526-01	CBL, 7FT ETHERNET(CAT5/CAT5E) GREEN	FRU
530-4527-01	CBL, 7FT ETHERNET(CAT5/CAT5E) YELLOW	FRU
7069031	CBL, BLACK, 10 FT, 10G, CAT6A	FRU
7069032	CBL, BLUE, 10 FT, 10G, CAT6A	FRU
371-4919-01	ASY, SLIDE-RAIL, LONG, 1U-2U, STAN	FRU
7042273	CBL, MGT, ARM, SLIM-RAIL, 1U, 2U, SNAP-IN	FRU
135-1203	XCVR, 1000BASE-T, COPPER SFP	FRU
530-4444	1M, QSFP Passive Copper, CBL	FRU

D.3 Rack-Level FRUs for Exalogic X2-2 and X3-2

Table D-3 lists the rack-level replaceable parts for Exalogic X2-2 and X3-2 machines.

Note:

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-3 Rack-Level Replacement Parts for EL X2-2 and X3-2

Part Number	Description	Coverage under Basic Warranty
F530-4502	CBL,10FT ETHERNET (CAT5/CAT5e),BLUE	FRU
F530-4526	CBL,7FT ETHERNET (CAT5/CAT5e),GREEN	FRU
F530-4527	CBL,7FT ETHERNET (CAT5/CAT5e),YELLOW	FRU
F530-4432	CBL,7FT ETHERNET (CAT5/CAT5e),BLUE	FRU
F530-4433	CBL,10FT ETHERNET (CAT5/CAT5e),RED	FRU
F530-4434	CBL,7FT ETHERNET (CAT5/CAT5e),RED	FRU
F530-4435	CBL,10FT ETHERNET (CAT5/CAT5e),BLACK	FRU
F530-4436	CBL,7FT ETHERNET (CAT5/CAT5e),BLACK	FRU
F530-4437	CBL,10FT ETHERNET (CAT5/CAT5e),ORANGE	FRU
F530-4438	CBL,7FT ETHERNET (CAT5/CAT5e),ORANGE	FRU
F350-1519	ASSY,SERIAL CABLE SET(3) NM (2 Serial Cables + 1 Null Modem)	FRU
F310-0307	ASSY,COOLING FAN FOR CISCO SWITCH	FRU
F371-4784	SNET,WS-C4948-S,CISCO CATALYST (Cisco Switch)	FRU

Table D-3 (Cont.) Rack-Level Replacement Parts for EL X2-2 and X3-2

Part Number	Description	Coverage under Basic Warranty
F371-4785	SPWR,PWR-C49-300AC,CISCO C4948 (Second Power Supply for Cisco Switch)	FRU
F350-1521	BOLT-ON vented FILLER PNL,SUN RACKII,1U,VENT	FRU
F350-1546	1U CMA FOR QUICKSILVER,X4170	FRU

D.4 Parts for ZS3-ES Storage Appliance for Exalogic X5-2

Table D-4 lists the replaceable parts for the ZS3-ES storage appliance, which is included in the Exalogic X5-2 machine:

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-4 Replacement Parts for Oracle ZS3-ES Appliance

Part Number	Description	Coverage under Basic Warranty
7015351	CPU, INTEL, E5-2658, 2.1G, 95W, 8-core	FRU
7020774	ASSY, HEATSINK, CPU 1U Nashua, 60 MM Fin	FRU
7018701	LV DIMM, 16GB,DDR3,1 600, 1.3V	FRU
7097377	ASSY, SSD, 1.6TB, 2.5, SAS2,Readzilla	FRU
7066874	FRU, DR, 900GB, SAS2, 10K Marlin	FRU
541-2732-03	ASSY, DVD FILLER W/USB, BZL, TRAY	FRU
7046442	DUAL PORT CX3 4XQDR PCI-E, LP	FRU
7048710	ASSY, RISER, X8, 1U, PCI-E	FRU
7048707	ASSY, RISER, X16, 1U, PCI-E	FRU
7027478	ASSY, 2.5 DISK BP	FRU
7013526	ASSY, FAN MODULE, CR, 40X56	FRU
7048712	ASSY, MOTHERBOARD, LYNX2	FRU
150-3993-01	BATTERY, 3V, 390MAH, LITH, COIN	FRU
7079385	PS, AC, A247A, F, 12V, 760W	FRU
7056175	ASSY, GEN-2 CLUSTER CRD, BKT, R13	FRU
7056272	PCA, THEBE-EXT, SAS PCIE 6GBS 8 PORT	FRU
530-3510-01	CBL, ASSY, SWITCH, INTRLCK, A77	FRU
530-3883-01	CABLE, MINI-SAS, 36POS, LONG, LYNX	FRU

For more information, see the *ZS3-ES Storage Appliance Customer Service Manual*.

D.5 Parts for ZS3-ES Storage Appliance for Exalogic X4-2

Table D-5 lists the replaceable parts for the ZS3-ES storage appliance, which is included in the Exalogic X4-2 machine:

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-5 Replacement Parts for Oracle ZS3-ES Appliance

Part Number	Description	Coverage under Basic Warranty
7015351	CPU, INTEL, E5-2658, 2.1G, 95W, 8-core	FRU
7020774	ASSY, HEATSINK, CPU 1U Nashua, 60 MM Fin	FRU
7018701	LV DIMM, 16GB,DDR3,1 600, 1.3V	FRU
7046052	SSD, 1.6TB, 2.5" SAS Readzilla, Marlin	FRU
7066874	FRU, DR, 900GB, SAS2, 10K Marlin	FRU
541-2732-03	ASSY, DVD FILLER W/USB, BZL, TRAY	FRU
375-3696-01	DUAL PORT CX2 4XQDR PCI-E, LP	FRU
7038488	Spare: 4GB USB Flash Memory	FRU
7048710	ASSY, RISER, X8, 1U, PCI-E	FRU
7048707	ASSY, RISER, X16, 1U, PCI-E	FRU
7027478	ASSY, 2.5 DISK BP	FRU
7020903	Assy, 2M, 4X Mini SAS Cbl, Shld	FRU
7013526	ASSY, FAN MODULE, CR, 40X56	FRU
7048712	ASSY, MOTHERBOARD, LYNX2	FRU
150-3993-01	BATTERY, 3V, 390MAH, LITH, COIN	FRU
7060951	FRU, PS, AC, A256, F,12V, 600W	FRU
7056175	ASSY, GEN-2 CLUSTER CRD, BKT, R13	FRU
7056272	PCA, THEBE-EXT, SAS PCIE 6GBS 8 PORT	FRU
7043562	FRU, ASSY, Indctr, 3.5" 1U, Nashua	FRU
7039841	FRU, ASM,SIS_PCB_2.5in. Bonnie	FRU

For more information, see the *ZS3-ES Storage Appliance Customer Service Manual*.

D.6 Parts for Sun ZFS Storage 7320 Appliance

Table D-6 lists the replaceable parts for the Sun ZFS Storage 7320 appliance, which is included in the Exalogic machine.

Note:

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-6 Replacement Parts for Sun ZFS Storage 7320 appliance

Part Number	Description	Coverage under Basic Warranty
F371-4885	1 x Intel Westmere E5620	FRU
F541-2545	Heatsink, CPU, X4170, ATO	FRU
F371-4898	4GB DDR3-1333 Registered ECC Memory (1x4GB)	FRU
F371-4743	4GB,SLC,USB-2,SBC (USB Thumb Drive)	FRU
F542-0371	500GB 7.2K 2.5 inch SATA HDD (Silver Marlin Bracket, OS drive)	FRU
7042768	500GB 10K 2.5 inch SATA HDD (Silver Marlin Bracket, OS drive)	FRU
F542-0330	512GB SSD, 2.5" SATA Marlin bracket (ReadZilla) (FW-AGXA0201)	FRU
F375-3609	8 Port LSI SAS-2 6G HBA	FRU
F375-3606	Dual port 4X QDR InfiniBand PCI-E HCA (CX1)	FRU
F511-1496	Generation 2 Cluster Heatbeat PCI-E Card	FRU
F300-2233	PS,AC,A247,F,12V,760W	FRU
F541-2732	DVD Drive Bay Panel	FRU

For more information, see the *Sun ZFS Storage 7x20 Appliance Customer Service Manual*.

D.7 Parts for Oracle Server X5-2 Compute Nodes

Table D-7 lists the replaceable parts for the Oracle Server X5-2 compute nodes used in the Exalogic machine.

Note:

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-7 Replacement Parts for Oracle Server X5-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7089073	FAN MODULE	FRU

Table D-7 (Cont.) Replacement Parts for Oracle Server X5-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7094132	SSD, 400GB ME, 2.5" SAS3, Marlin	FRU
7078072	DIMM, 32GB, DDR3, 1600, 4Rx4, 1.35V	FRU
7067728	HEAT SINK CPU 1U	FRU
7095106	CPU, E5-2699, 2.3G, 145W, 18C	FRU
7097205	ASSY, 2.5 DISK BP	FRU
7085209	ASY, 12G SAS RAID HBA, INT, 8P	FRU
7083433	ASSY, RISER, DOUBLE_STACK	FRU
7083430	ASSY, RISER	FRU
7098505	MOTHER BOARD ASSEMBLY	FRU
7050794	ASSY, BBU08 BATTERY MODULE, 55 C	FRU
7057184	ASSY, MOUNT-BATTERY, REMOTE	FRU
7079395	PS, AC, A256, F, 12V, 600W	FRU
7086632	ASM, SIS_PCB_2.5IN, BONNIE	FRU
530-4445-01	3M, 10GE QSFP Passive Copp, CBL	FRU
7094264	Cable Kit	FRU
7300399	ASM, FIM, Disk, 1U	FRU

For more information, see the *Oracle Server X5-2 Service Manual*.

D.8 Parts for Sun Server X4-2 Compute Nodes

Table D-8 lists the replaceable parts for the Sun Server X4-2 compute nodes used in the Exalogic machine.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-8 Replacement Parts for Sun Server X4-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7058153	Motherboard Assembly	FRU
7069240	FRU, CPU, Intel E5-2697, 12Core, 2.7 GHz	FRU
7060951	FRU, PS,AC, A256, F, 12V, 600W	FRU
7013526	FRU, ASM, FAN_MOD, NASHUA	FRU
7048710	FRU ASSY, RISER, NASHUA	FRU

Table D-8 (Cont.) Replacement Parts for Sun Server X4-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7048707	FRU ASSY, RISER, DOUBLE_STACK, NASHUA	FRU
7020774	FRU ASSY, HEAT SINK CPU 1U NASHUA, 60 MM FIN	FRU
7027478	ASM, SIS_PCB_2.5IN, BONNIE Rear 2 disk bkpln	FRU
7019614	FRU, Cable Kit, Nashua	FRU
7046442	CX3	FRU
7027478	ASSY, 2.5 DISK BP	FRU
7018701	DIMM, 16GB, DDR3, 1600, 2Rx4, 1.35V	FRU
7038488	FLASH DRIVE, TLC, OSA ID, USB, 4GB	FRU
7047503	ASY, 6G SAS RAID HBA, INT, B4, 8P	FRU
7057184	ASSY, MOUNT-BATTERY, REMOTE	FRU
7077233	400 GB SSD in Marlin Brkt	FRU

For more information, see the *Sun Server X4-2 Service Manual*.

D.9 Parts for Sun Server X3-2 Compute Nodes

Table D-9 lists the replaceable parts for the Sun Server X3-2 compute nodes used in the Exalogic machine.

Note:

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-9 Replacement Parts for Sun Server X3-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7026878	1X 2.9 Ghz Intel 8-Core Xeon E5-2690, 135 W	FRU
7020774	Heatsink, CPU, X4170 M3	FRU
7018701	16 GB DIMMs	FRU
F375-3696	Dual port 4X QDR InfiniBand PCI-E HCA (CX2)	FRU
F375-3701	6 GB/s SAS RAID HBA,Internal	FRU
7105627	ASSY Battery Carrier Remote	FRU
7106158	Universal Upgrade Remote Battery Kit	FRU
542-0388	300 GB Disk	FRU
7047410	A256 600 W AC Input Power Supply	FRU
7048712	System Board Assembly	FRU

Table D-9 (Cont.) Replacement Parts for Sun Server X3-2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
7027478	4-Slot 2.5" Disk Backplane Assembly	FRU
7048710	1 Slot PCI Express Riser Assembly	FRU
7048707	2 Slot PCI Express Riser Assembly	FRU
7013526	Dual Counter Rotating Fan Module	FRU
F150-3993	CR2032 Battery (for M/B)	FRU
350-1287	Cable Management Arm	FRU
350-1719	Snap-In Slide Rail Rackmount Kit	FRU
7039841	2.5" Disk Cage Front Indicator Module	FRU
7019614	Cable Kit (includes the following)	FRU
530-4537	Disk Drive Flex Cable	FRU
7014918	Dual Backplane Power Cable	FRU
7014919	Single Backplane Power Cable	FRU
7019726	Interlock Switch Cable	FRU
7020903	Disk Data Cable, SFF-8087 to SFF-8643, 620 mm	FRU
7027302	Front Indicator Module Cable	FRU
7041996	Disk Data Cable, SFF-8087 to SFF 8643, 650mm	FRU

For more information, see the *Sun Server X3-2 Service Manual*.

D.10 Parts for X4170 M2 Compute Nodes

Table D-10 lists the replaceable parts for the X4170 M2 compute nodes used in the Exalogic machine.

Note:

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-10 Replacement Parts for X4170 M2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
F371-4889	1 x Intel Westmere X5670	FRU
F541-2545	Heatsink, CPU, X4170, ATO	FRU
F371-4966	8 GB DDR3-1333 Low Voltage Registered ECC Memory (1 x 8GB)	FRU

Table D-10 (Cont.) Replacement Parts for X4170 M2 Compute Nodes

Part Number	Description	Coverage with Basic Warranty
F375-3696	Dual port 4X QDR InfiniBand PCI-E HCA (CX2)	FRU
F375-3701	6 GB/s SAS RAID HBA,Internal	FRU
F371-4982	Battery for 6 GB/s SAS RAID HBA,Internal (BBU08)	FRU
F540-7841	32 GB SATA SSD	FRU
F300-2233	PS,AC,A247,F,12V,760W	FRU
F541-4081	Motherboard (FRU'd in removable tray assembly.)	FRU
F511-1489	PDB	FRU
F511-1548	8 Disk BP	FRU
F541-2739	PCI_BAR_1U	FRU
F541-2883	PCI-E Riser, X8, 1U	FRU
F541-2885	PCI-E Riser, X16, 1U	FRU
F541-4275	Paddle Card, CR_FANs	FRU
F530-4228	CBL,MINI-SAS,36POS,LONG,LY	FRU
F530-4431	CABLE,FAN_DATA,1U	FRU
F530-3927	CBL,PDB,MB,1U+2U,RIBBON	FRU
F530-4417	CABLE,FAN-PADDLE BD,P	FRU
F541-4276	FAN MODULE,CR,40X56	FRU
F541-4274	FAN_DECK,1U,CR_FANS	FRU
F541-2075	BUSS_BAR,PWR,1U+2U	FRU
F150-3993	CR2032 Battery (for M/B)	FRU

For more information, see the *Sun Fire X4170 M2 Server Service Manual*.

D.11 Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X5-2

Table D-11 lists the replaceable parts for the Sun Network QDR InfiniBand Gateway Switches used in Exalogic X5-2 machines.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-11 Replacement Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X5-2 Machines

Part Number	Description	Coverage with Basic Warranty
7054724	ASSY, SYSTEM, NM2-GW, LF	FRU
7061031	FAN MODULE, 1RU, REVERSE-FLOW	FRU
371-2210-01	ASSY, CR2032 BATTERY, A84/85/88/93	FRU
7065505	PS, AC, A247B	FRU

For more information, see the *Sun Network QDR InfiniBand Gateway Switch Service Manual*.

D.12 Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X4-2

Table D-12 lists the replaceable parts for the Sun Network QDR InfiniBand Gateway Switches used in Exalogic X4-2 machines.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-12 Replacement Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X4-2 Machines

Part Number	Description	Coverage with Basic Warranty
7054724	NM2-GW system FRU	FRU
300-2299-01	PS, AC, A247A, F, 12V, 760W	FRU
7061031	ASSY, FAN MOD, 1RU, SANACE, REVERSE-FLOW	FRU

For more information, see the *Sun Network QDR InfiniBand Gateway Switch Service Manual*.

D.13 Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X2-2 and X3-2

Table D-13 lists the replaceable parts for the Sun Network QDR InfiniBand Gateway Switches used in Exalogic X2-2 and X3-2 machines.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-13 Replacement Parts for Sun Network QDR InfiniBand Gateway Switch for Exalogic X2-2 and X3-2 Machines

Part Number	Description	Coverage with Basic Warranty
F350-1566-01	FRU, ASSY, FAN MOD, 1RU, SANACE, REVERSE	FRU
F300-2233-02	FRU, PS, AC, A247, F, 12V, 760W	FRU
F371-2210-01	FRU, CR2032, BATTERY, A84/85/88	FRU
F541-4188-01	FRU, ASSY, SYSTEM, NM2-GATEWAY	FRU

For more information, see the *Sun Network QDR InfiniBand Gateway Switch Service Manual*.

D.14 Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X5-2

Table D-14 lists the replaceable parts for the Sun Datacenter InfiniBand Switch 36 used in the Exalogic X5-2 machine full rack and half rack configurations.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-14 Replacement Parts for Sun Datacenter InfiniBand Switch 36

Part Number	Description	Coverage with Basic Warranty
7052970	Assy, System, NM2-36P Managed	FRU
7065505	PS, AC, A247B, F, 12V, 760W	FRU
7061031	Fan Module, 1RU, Reverse-Flow	FRU
371-2210	ASSY, CR2032 Battery	FRU

For more information, see the *Sun Datacenter InfiniBand Switch 36 User's Guide*.

D.15 Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X4-2

Table D-15 lists the replaceable parts for the Sun Datacenter InfiniBand Switch 36 used in Exalogic X4-2 machine full rack and half rack configurations.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-15 Replacement Parts for Sun Datacenter InfiniBand Switch 36

Part Number	Description	Coverage with Basic Warranty
7052970	ASSY, SYSTEM, NM2-36P MANAGED	FRU
300-2299-01	PS, AC, A247A, F, 12V, 760W	FRU
7061031	ASSY, FAN MOD, 1RU, SANACE, REVERSE-FLOW	FRU

For more information, see the *Sun Datacenter InfiniBand Switch 36 User's Guide*.

D.16 Parts for Sun Datacenter InfiniBand Switch 36 for Exalogic X2-2 and X3-2

Table D-16 lists the replaceable parts for the Sun Datacenter InfiniBand Switch 36 used in Exalogic X2-2 and X3-2 machines full rack and half rack configurations.

> **Note:**
>
> All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-16 Replacement Parts for Sun Datacenter InfiniBand Switch 36

Part Number	Description	Coverage with Basic Warranty
F541-3495-03	FRU, ASSY, SYSTEM, NM2-36P MANAGE	FRU
F350-1312-04	FRU, FAN MODULE, 1RU, REVERSE-FLO	FRU
F300-2233-02	FRU, PS, AC, A247, F, 12V, 760W	FRU

For more information, see the *Sun Datacenter InfiniBand Switch 36 User's Guide*.

D.17 Parts for the Cisco Catalyst 4948E-F-S Switch

Table D-17 lists the replacement parts for the Cisco Catalyst 4948E-F-S switch.

Table D-17 Replacement Parts for the Cisco Catalyst 4948E-F-S Switch

Part Number	Description	Repair Category	Coverage with Basic Warranty
7023685	Cisco C4948E-F-S switch without the power supply and fan	IR	FRU
7024423	Power supply for Cisco 4948E-F-S switch	HS	FRU
7024424	Cooling fan for Cisco 4948E-F-S switch	HS	FRU

For repair information, see the *Cisco Catalyst 4948E Switch Installation Guide*.

D.18 Parts for the Cisco Catalyst 4948 Switch

Table D-18 lists the replacement parts for the Cisco Catalyst 4948 switch.

Table D-18 Replacement Parts for Cisco Catalyst 4948 Switch

Part Number	Description	Repair Category	Coverage with Basic Warranty
F371-4784	Cisco Catalyst 4948 switch	IR	FRU
F371-4785	Power supply for Cisco Catalyst 4948 switch	HS	FRU
F310-0307	Cooling fan for Cisco Catalyst 4948 switch	HS	FRU

For repair information, see the *Cisco Catalyst 4948E Switch Installation Guide*.

D.19 Parts for the Gari DE2-24C Disk Enclosure for Exalogic X5-2

Table D-19 lists the replacement parts for the Gari DE2-24C disk enclosure for Exalogic X5-2.

Table D-19 Replacement Parts for Gari DE2-24C Disk Enclosure

Part Number	Description	Coverage with Basic Warranty
7094120	ASSY,SSD,200GB WI, SAS3 Logz	FRU
7043627	ASSY, DE2-24P&C 580W AC PCM	FRU
7043628	ASSY, DE2-24P&C I/O Module (IOM)	FRU
7044319	ASSY, 4U Chassis with Midplane	FRU
7063682	ASSY, DE2-24C rail kit for rack shipments	FRU

Table D-19 (Cont.) Replacement Parts for Gari DE2-24C Disk Enclosure

Part Number	Description	Coverage with Basic Warranty
7066831	ASSY, Drive, 4TB. 3.5" SAS	FRU
530-3883-01	ASSY, 2M, 4X MINI SAS CBL, SHLD	FRU

D.20 Parts for the Gari DE2-24C Disk Enclosure

Table D-20 lists the replacement parts for the Gari DE2-24C disk enclosure.

Table D-20 Replacement Parts for Gari DE2-24C Disk Enclosure

Part Number	Description	Coverage with Basic Warranty
7043627	ASSY, DE2-24P&C 580W AC Power Cooling Module	FRU
7043628	ASSY, DE2-24P&C I/O Module (IOM)	FRU
7043630	ASSY, 2U Chassis with Midplane	FRU
530-3883	ASSY, 2M, 4X MINI SAS CBL, SHLD	FRU
530-3886	ASSY, 0.5M, 4X MINI SAS CBL, SHLD	FRU
7044619	ASSY, DE2-24P UNIVERSAL RAIL KIT	FRU
7044396	ASSY DE2-24P SAS 73GB SSD, LogZ	FRU
7066831	Assy, DRIVE, 4T 3.5", SAS/7200 rpm, DE2-24C	FRU

E

Cabling Tables

The tables in this appendix show the cable layouts for Exalogic machines.
This appendix contains the following topics:

- Administrative Gigabit Ethernet Network-Cabling Tables
- InfiniBand Network Cabling Tables
- Power Distribution Unit Cabling Tables

> **Note:**
>
> The eighth-, quarter-, and half-rack X4-2 and newer machines are pre-cabled with a complete set (equivalent in number to a full rack) of InfiniBand (92), Ethernet (42), and power cables. All the cables are connected to the appropriate switches and routed to the correct rack-unit location. The unconnected ends of the cables are tied off to lacer bars. At a later time, when the machine is upgraded to a larger rack, the filler panels and lacer bars are removed, the X4-2 and newer servers are installed with rack rails and cable-management arms, and the cables (power, InfiniBand, and Ethernet) are connected to the newly installed compute nodes.

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.
U*n*	Unit height in rack, where *n* is the number, such as U20.
P*n*	InfiniBand port *n*, where n is port number, such as P8A.

> **Note:**
>
> For Exalogic machine rack layouts, see Exalogic Machine Rack Layout.

E.1 Administrative Gigabit Ethernet Network-Cabling Tables

This section contains the tables for the administrative Gigabit Ethernet network cabling. The Gigabit Ethernet switch is located in rack unit 25.

Table E-1 shows the cable connections from the compute nodes to the Gigabit Ethernet switch (Cisco switch) in an Exalogic machine full rack:

Table E-1 Gigabit Ethernet Cabling for Exalogic Machine Full Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
Not applicable	External LAN	NET	1	10' Blue Cat 5
Not applicable	PDU-A	NET	2	1M Grey Cat 5
Not applicable	PDU-B	NET	3	1M Grey Cat 5
U42	Compute node	NET0	4	10' Blue Cat 5
U41	Compute node	NET0	5	10' Blue Cat 5
U40	Compute node	NET0	6	10' Blue Cat 5
U39	Compute node	NET0	7	10' Blue Cat 5
U38	Compute node	NET0	8	10' Blue Cat 5
U37	Compute node	NET0	9	10' Blue Cat 5
U36	Compute node	NET0	10	7' Blue Cat 5
U35	Compute node	NET0	11	7' Blue Cat 5
U34	Compute node	NET0	12	7' Blue Cat 5
U33	Compute node	NET0	13	7' Blue Cat 5
U32	Compute node	NET0	14	7' Blue Cat 5
U31	Compute node	NET0	15	7' Blue Cat 5
U30	Compute node	NET0	16	7' Blue Cat 5
U29	Compute node	NET0	17	7' Blue Cat 5
U28	Compute node	NET0	18	7' Blue Cat 5
U27	Compute node	NET0	19	7' Blue Cat 5
U26	Sun Network QDR InfiniBand Gateway Switch	NET0	38	7' Black Cat 5
U24	Sun Network QDR InfiniBand Gateway Switch	NET0	39	7' Black Cat 5
U23	Sun Network QDR InfiniBand Gateway Switch	NET0	40	7' Black Cat 5
U22	Storage appliance server head	NET0	20	7' Blue Cat 5
U22	Storage appliance server head	NET1	21	7' Blue Cat 5
U22	Storage appliance server head	NET2	43	7' Blue Cat 5
U21	Storage appliance server head	NET0	22	7' Blue Cat 5
U21	Storage appliance server head	NET1	23	7' Blue Cat 5
U21	Storage appliance server head	NET2	44	7' Blue Cat 5
U16	Sun Network QDR InfiniBand Gateway Switch	NET0	41	7' Black Cat 5
U15	Compute node	NET0	24	7' Blue Cat 5

Table E-1 (Cont.) Gigabit Ethernet Cabling for Exalogic Machine Full Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
U14	Compute node	NET0	25	7' Blue Cat 5
U13	Compute node	NET0	26	7' Blue Cat 5
U12	Compute node	NET0	27	10' Blue Cat 5
U11	Compute node	NET0	28	10' Blue Cat 5
U10	Compute node	NET0	29	10' Blue Cat 5
U9	Compute node	NET0	30	10' Blue Cat 5
U8	Compute node	NET0	31	10' Blue Cat 5
U7	Compute node	NET0	32	10' Blue Cat 5
U6	Compute node	NET0	33	10' Blue Cat 5
U5	Compute node	NET0	34	10' Blue Cat 5
U4	Compute node	NET0	35	10' Blue Cat 5
U3	Compute node	NET0	36	10' Blue Cat 5
U2	Compute node	NET0	37	10' Blue Cat 5

Table E-2 shows the cable connections from the compute nodes to the Gigabit Ethernet switch (Cisco switch) in an Exalogic machine half rack:

Table E-2 Gigabit Ethernet Cabling for Exalogic Machine Half Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
Not applicable	External LAN	NET	1	10' Blue Cat 5
Not applicable	PDU-A	NET	2	1M Grey Cat 5
Not applicable	PDU-B	NET	3	1M Grey Cat 5
U28	Compute node	NET0	18	7' Blue Cat 5
U27	Compute node	NET0	19	7' Blue Cat 5
U23	Sun Network QDR InfiniBand Gateway Switch	NET0	40	7' Black Cat 5
U22	Storage appliance server head	NET0	20	7' Blue Cat 5
U22	Storage appliance server head	NET1	21	7' Blue Cat 5
U22	Storage appliance server head	NET2	43	7' Blue Cat 5
U21	Storage appliance server head	NET0	22	7' Blue Cat 5
U21	Storage appliance server head	NET1	23	7' Blue Cat 5

Table E-2 (Cont.) Gigabit Ethernet Cabling for Exalogic Machine Half Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
U21	Storage appliance server head	NET2	44	7' Blue Cat 5
U16	Sun Network QDR InfiniBand Gateway Switch	NET0	41	7' Black Cat 5
U15	Compute node	NET0	24	7' Blue Cat 5
U14	Compute node	NET0	25	7' Blue Cat 5
U13	Compute node	NET0	26	7' Blue Cat 5
U12	Compute node	NET0	27	10' Blue Cat 5
U11	Compute node	NET0	28	10' Blue Cat 5
U10	Compute node	NET0	29	10' Blue Cat 5
U9	Compute node	NET0	30	10' Blue Cat 5
U8	Compute node	NET0	31	10' Blue Cat 5
U7	Compute node	NET0	32	10' Blue Cat 5
U6	Compute node	NET0	33	10' Blue Cat 5
U5	Compute node	NET0	34	10' Blue Cat 5
U4	Compute node	NET0	35	10' Blue Cat 5
U3	Compute node	NET0	36	10' Blue Cat 5
U2	Compute node	NET0	37	10' Blue Cat 5

Table E-3 shows the cable connections from the compute nodes to the Gigabit Ethernet switch (Cisco switch) in an Exalogic machine quarter rack:

Table E-3 Gigabit Ethernet Cabling for Exalogic Machine Quarter Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
Not applicable	External LAN	NET	1	10' Blue Cat 5
Not applicable	PDU-A	NET	2	1M Grey Cat 5
Not applicable	PDU-B	NET	3	1M Grey Cat 5
U23	Sun Network QDR InfiniBand Gateway Switch	NET0	40	7' Black Cat 5
U22	Storage appliance server head	NET0	20	7' Blue Cat 5
U22	Storage appliance server head	NET1	21	7' Blue Cat 5
U22	Storage appliance server head	NET2	43	7' Blue Cat 5
U21	Storage appliance server head	NET0	22	7' Blue Cat 5

Table E-3 (Cont.) Gigabit Ethernet Cabling for Exalogic Machine Quarter Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
U21	Storage appliance server head	NET1	23	7' Blue Cat 5
U21	Storage appliance server head	NET2	44	7' Blue Cat 5
U16	Sun Network QDR InfiniBand Gateway Switch	NET0	41	7' Black Cat 5
U9	Compute node	NET0	30	10' Blue Cat 5
U8	Compute node	NET0	31	10' Blue Cat 5
U7	Compute node	NET0	32	10' Blue Cat 5
U6	Compute node	NET0	33	10' Blue Cat 5
U5	Compute node	NET0	34	10' Blue Cat 5
U4	Compute node	NET0	35	10' Blue Cat 5
U3	Compute node	NET0	36	10' Blue Cat 5
U2	Compute node	NET0	37	10' Blue Cat 5

Table E-4 shows the cable connections from the compute nodes to the Gigabit Ethernet switch (Cisco switch) in an Exalogic machine eighth rack:

Table E-4 Gigabit Ethernet Cabling for Exalogic Machine Eighth Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
Not applicable	External LAN	NET	1	10' Blue Cat 5
Not applicable	PDU-A	NET	2	1M Grey Cat 5
Not applicable	PDU-B	NET	3	1M Grey Cat 5
U23	Sun Network QDR InfiniBand Gateway Switch	NET0	40	7' Black Cat 5
U22	Storage appliance server head	NET0	20	7' Blue Cat 5
U22	Storage appliance server head	NET1	21	7' Blue Cat 5
U22	Storage appliance server head	NET2	43	7' Blue Cat 5
U21	Storage appliance server head	NET0	22	7' Blue Cat 5
U21	Storage appliance server head	NET1	23	7' Blue Cat 5
U21	Storage appliance server head	NET2	44	7' Blue Cat 5
U16	Sun Network QDR InfiniBand Gateway Switch	NET0	41	7' Black Cat 5
U5	Compute node	NET0	34	10' Blue Cat 5
U4	Compute node	NET0	35	10' Blue Cat 5

Table E-4 (Cont.) Gigabit Ethernet Cabling for Exalogic Machine Eighth Rack

From Rack Slot	Type of Equipment	From Equipment Port	To Cisco Switch Port	Description
U3	Compute node	`NET0`	36	10' Blue Cat 5
U2	Compute node	`NET0`	37	10' Blue Cat 5

E.2 InfiniBand Network Cabling Tables

This section contains the tables for the InfiniBand network cabling. The Sun Network QDR InfiniBand Gateway Switch is located in located in rack units 26, 24, 23, or 16 (depending on the rack configuration).

This section contains the following topics:

- Exalogic Machine Full Rack
- Exalogic Machine Half Rack
- Exalogic Machine Quarter Rack
- Exalogic Machine Eighth Rack

E.2.1 Exalogic Machine Full Rack

This section contains the following tables:

- Table E-5
- Table E-6
- Table E-7

Table E-5 lists the location, ports and cables for the InfiniBand connections for an Exalogic machine full rack.

Table E-5 InfiniBand Network Cabling for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U26	9A	U42	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U26	9B	U41	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U26	10A	U40	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-5 (Cont.) InfiniBand Network Cabling for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U26	10B	U39	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U26	11A	U38	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U26	11B	U37	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U26	12A	U36	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U26	12B	U35	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	13A	U34	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	13B	U33	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	14A	U32	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	14B	U31	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	15A	U30	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U26	15B	U29	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	9A	U42	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U24	9B	U41	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U24	10A	U40	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U24	10B	U39	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U24	11A	U38	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U24	11B	U37	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	12A	U36	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	12B	U35	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable

Table E-5 (Cont.) InfiniBand Network Cabling for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U24	13A	U34	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	13B	U33	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	14A	U32	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	14B	U31	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	15A	U30	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	15B	U29	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	7A	U28	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	7B	U27	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	8A	U22	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	8B	U21	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	9A	U15	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	9B	U14	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	10A	U13	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	10B	U12	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	11A	U11	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U23	11B	U10	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U23	12A	U9	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	12B	U8	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable

Table E-5 (Cont.) InfiniBand Network Cabling for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	13A	U7	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	13B	U6	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14A	U5	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14B	U4	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15A	U3	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15B	U2	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	7A	U28	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	7B	U27	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	8A	U22	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	8B	U21	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	9A	U15	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	9B	U14	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	10A	U13	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	10B	U12	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	11A	U11	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	11B	U10	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	12A	U9	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	12B	U8	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable

Table E-5 (Cont.) InfiniBand Network Cabling for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	13A	U7	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	13B	U6	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14A	U5	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14B	U4	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15A	U3	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15B	U2	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-6 lists the InfiniBand gateway switch inter-connections for an Exalogic machine full rack.

Table E-6 InfiniBand Gateway Switch Inter-Connections for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	4A	U23	Sun Network QDR InfiniBand Gateway Switch	4A	2 meter QDR InfiniBand cable
U16	4B	U23	Sun Network QDR InfiniBand Gateway Switch	4B	2 meter QDR InfiniBand cable
U16	5A	U23	Sun Network QDR InfiniBand Gateway Switch	5A	2 meter QDR InfiniBand cable
U16	5B	U23	Sun Network QDR InfiniBand Gateway Switch	5B	2 meter QDR InfiniBand cable

Table E-6 (Cont.) InfiniBand Gateway Switch Inter-Connections for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	6A	U23	Sun Network QDR InfiniBand Gateway Switch	6A	2 meter QDR InfiniBand cable
U16	6B	U23	Sun Network QDR InfiniBand Gateway Switch	6B	2 meter QDR InfiniBand cable
U24	4A	U26	Sun Network QDR InfiniBand Gateway Switch	4A	1 meter QDR InfiniBand cable
U24	4B	U26	Sun Network QDR InfiniBand Gateway Switch	4B	1 meter QDR InfiniBand cable
U24	5A	U26	Sun Network QDR InfiniBand Gateway Switch	5A	1 meter QDR InfiniBand cable
U24	5B	U26	Sun Network QDR InfiniBand Gateway Switch	5B	1 meter QDR InfiniBand cable
U24	6A	U26	Sun Network QDR InfiniBand Gateway Switch	6A	1 meter QDR InfiniBand cable
U24	6B	U26	Sun Network QDR InfiniBand Gateway Switch	6B	1 meter QDR InfiniBand cable

Table E-7 lists the InfiniBand gateway switch cross-connections for an Exalogic machine full rack.

Table E-7 InfiniBand Gateway Switch Cross-Connections for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	0A	U26	Sun Network QDR InfiniBand Gateway Switch	0A	2 meter QDR InfiniBand cable
U16	0B	U26	Sun Network QDR InfiniBand Gateway Switch	0B	2 meter QDR InfiniBand cable
U16	1A	U26	Sun Network QDR InfiniBand Gateway Switch	1A	2 meter QDR InfiniBand cable
U16	1B	U26	Sun Network QDR InfiniBand Gateway Switch	1B	2 meter QDR InfiniBand cable
U16	2A	U24	Sun Network QDR InfiniBand Gateway Switch	2A	2 meter QDR InfiniBand cable
U16	2B	U24	Sun Network QDR InfiniBand Gateway Switch	2B	2 meter QDR InfiniBand cable
U16	3A	U24	Sun Network QDR InfiniBand Gateway Switch	3A	2 meter QDR InfiniBand cable
U16	3B	U24	Sun Network QDR InfiniBand Gateway Switch	3B	2 meter QDR InfiniBand cable
U23	0A	U24	Sun Network QDR InfiniBand Gateway Switch	0A	1 meter QDR InfiniBand cable
U23	0B	U24	Sun Network QDR InfiniBand Gateway Switch	0B	1 meter QDR InfiniBand cable
U23	1A	U24	Sun Network QDR InfiniBand Gateway Switch	1A	1 meter QDR InfiniBand cable
U23	1B	U24	Sun Network QDR InfiniBand Gateway Switch	1B	1 meter QDR InfiniBand cable
U23	2A	U26	Sun Network QDR InfiniBand Gateway Switch	2A	1 meter QDR InfiniBand cable

Table E-7 (Cont.) InfiniBand Gateway Switch Cross-Connections for Exalogic Machine Full Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	2B	U26	Sun Network QDR InfiniBand Gateway Switch	2B	1 meter QDR InfiniBand cable
U23	3A	U26	Sun Network QDR InfiniBand Gateway Switch	3A	1 meter QDR InfiniBand cable
U23	3B	U26	Sun Network QDR InfiniBand Gateway Switch	3B	1 meter QDR InfiniBand cable

E.2.2 Exalogic Machine Half Rack

This section contains the following tables:

- Table E-8
- Table E-9

Table E-8 lists the location, ports and cables for the InfiniBand connections for an Exalogic machine half rack.

Table E-8 InfiniBand Network Cabling for Exalogic Machine Half Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	7A	U28	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	7B	U27	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	8A	U22	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	8B	U21	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	9A	U15	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable

Table E-8 (Cont.) InfiniBand Network Cabling for Exalogic Machine Half Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	9B	U14	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	10A	U13	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	10B	U12	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U23	11A	U11	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U23	11B	U10	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U23	12A	U9	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	12B	U8	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	13A	U7	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	13B	U6	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14A	U5	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14B	U4	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15A	U3	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15B	U2	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	7A	U28	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	7B	U27	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	8A	U22	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	8B	U21	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	9A	U15	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable

Table E-8 (Cont.) InfiniBand Network Cabling for Exalogic Machine Half Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	9B	U14	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	10A	U13	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	10B	U12	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	11A	U11	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	11B	U10	Compute node	PCIe 2, P2	2 meter QDR InfiniBand cable
U16	12A	U9	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	12B	U8	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	13A	U7	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	13B	U6	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14A	U5	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14B	U4	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15A	U3	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15B	U2	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-9 lists the InfiniBand gateway switch inter-connections for an Exalogic machine half rack.

Table E-9 InfiniBand Gateway Switch Inter-Connections for Exalogic Machine Half Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	4A	U23	Sun Network QDR InfiniBand Gateway Switch	4A	2 meter QDR InfiniBand cable
U16	4B	U23	Sun Network QDR InfiniBand Gateway Switch	4B	2 meter QDR InfiniBand cable
U16	5A	U23	Sun Network QDR InfiniBand Gateway Switch	5A	2 meter QDR InfiniBand cable
U16	5B	U23	Sun Network QDR InfiniBand Gateway Switch	5B	2 meter QDR InfiniBand cable
U16	6A	U23	Sun Network QDR InfiniBand Gateway Switch	6A	2 meter QDR InfiniBand cable
U16	6B	U23	Sun Network QDR InfiniBand Gateway Switch	6B	2 meter QDR InfiniBand cable

E.2.3 Exalogic Machine Quarter Rack

This section contains the following tables:

- Table E-10
- Table E-11

Table E-10 lists the location, ports and cables for the InfiniBand connections for an Exalogic machine quarter rack.

Table E-10 InfiniBand Network Cabling for Exalogic Machine Quarter Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	8A	U22	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable

Table E-10 (Cont.) InfiniBand Network Cabling for Exalogic Machine Quarter Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	8B	U21	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	12A	U9	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	12B	U8	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	13A	U7	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	13B	U6	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14A	U5	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14B	U4	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15A	U3	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15B	U2	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	8A	U22	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	8B	U21	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	12A	U9	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	12B	U8	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	13A	U7	Compute node	PCIe 2, P1	2 meter QDR InfiniBand cable
U16	13B	U6	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14A	U5	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14B	U4	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15A	U3	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-10 (Cont.) InfiniBand Network Cabling for Exalogic Machine Quarter Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	15B	U2	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-11 lists the InfiniBand gateway switch inter-connections for an Exalogic machine quarter rack.

Table E-11 InfiniBand Gateway Switch Inter-Connections for Exalogic Machine Quarter Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	4A	U23	Sun Network QDR InfiniBand Gateway Switch	4A	2 meter QDR InfiniBand cable
U16	4B	U23	Sun Network QDR InfiniBand Gateway Switch	4B	2 meter QDR InfiniBand cable
U16	5A	U23	Sun Network QDR InfiniBand Gateway Switch	5A	2 meter QDR InfiniBand cable
U16	5B	U23	Sun Network QDR InfiniBand Gateway Switch	5B	2 meter QDR InfiniBand cable
U16	6A	U23	Sun Network QDR InfiniBand Gateway Switch	6A	2 meter QDR InfiniBand cable
U16	6B	U23	Sun Network QDR InfiniBand Gateway Switch	6B	2 meter QDR InfiniBand cable

E.2.4 Exalogic Machine Eighth Rack

This section contains the following tables:

- Table E-12

- Table E-13

Table E-12 lists the location, ports and cables for the InfiniBand connections for an Exalogic machine eighth rack.

Table E-12 InfiniBand Network Cabling for Exalogic Machine Eighth Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U23	8A	U22	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	8B	U21	Storage appliance server head	PCIe 0, P2	2 meter QDR InfiniBand cable
U23	14A	U5	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	14B	U4	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15A	U3	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U23	15B	U2	Compute node	PCIe 2, P2	3 meter QDR InfiniBand cable
U16	8A	U22	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	8B	U21	Storage appliance server head	PCIe 0, P1	3 meter QDR InfiniBand cable
U16	14A	U5	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	14B	U4	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15A	U3	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable
U16	15B	U2	Compute node	PCIe 2, P1	3 meter QDR InfiniBand cable

Table E-13 lists the InfiniBand gateway switch inter-connections for an Exalogic machine eighth rack.

Table E-13 InfiniBand Gateway Switch Inter-Connections for Exalogic Machine Eighth Rack

From Rack Unit (Sun Network QDR InfiniBand Gateway Switch)	Port	To Rack Unit	Type of Equipment	Port	Description
U16	4A	U23	Sun Network QDR InfiniBand Gateway Switch	4A	2 meter QDR InfiniBand cable
U16	4B	U23	Sun Network QDR InfiniBand Gateway Switch	4B	2 meter QDR InfiniBand cable
U16	5A	U23	Sun Network QDR InfiniBand Gateway Switch	5A	2 meter QDR InfiniBand cable
U16	5B	U23	Sun Network QDR InfiniBand Gateway Switch	5B	2 meter QDR InfiniBand cable
U16	6A	U23	Sun Network QDR InfiniBand Gateway Switch	6A	2 meter QDR InfiniBand cable
U16	6B	U23	Sun Network QDR InfiniBand Gateway Switch	6B	2 meter QDR InfiniBand cable
U16	USB Serial	Not applicable	Open ended	Not applicable	USB Adapter
U16	USB Serial	Not applicable	Open ended	Not applicable	USB Adapter

E.3 Power Distribution Unit Cabling Tables

This section contains the tables for the Power Distribution Units (PDU) cabling for Exalogic Elastic Cloud.

This section contains the following topics for Single Phase PDUs:

- Single Phase PDU Power Cabling Half Rack
- Single Phase PDU Power Cabling Quarter Rack
- Single Phase PDU Power Cabling Eighth Rack

This section contains the following topics for Three Phase PDUs:

- Three Phase PDU Power Cabling Full Rack
- Three Phase PDU Power Cabling Half Rack

- Three Phase PDU Power Cabling Quarter Rack
- Three Phase PDU Power Cabling Eighth Rack

E.3.1 Single Phase PDU Power Cabling Half Rack

This section contains the following tables:

- Table E-14

Table E-14 lists the location, ports and cables for the Single Phase PDU connections for an Exalogic machine half rack.

Table E-14 Single Phase PDU Power Cabling for Exalogic Machine Half Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U28	Compute node	G3-3	G2-3	2M
U27	Compute node	G3-2	G2-4	2M
U25	Cisco Switch	G3-1	G2-5	2M
U23	NM2-GW-SW	G2-6	G3-0	2M
U22	ZS3-ES HN	G2-5	G3-1	2M
U21	ZS3-ES HN	G2-4	G3-2	2M
U17	DE2-24C DS	G2-3	G3-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U15	Compute node	G2-2	G3-4	2M
U14	Compute node	G2-1	G3-5	2M
U13	Compute node	G2-0	G3-6	2M
U12	Compute node	G1-4	G4-2	2M
U11	Compute node	G1-3	G4-3	2M
U10	Compute node	G1-2	G4-4	2M
U09	Compute node	G1-1	G4-5	2M
U08	Compute node	G1-0	G4-6	2M
U07	Compute node	G0-6	G5-0	2M
U06	Compute node	G0-5	G5-1	2M
U05	Compute node	G0-4	G5-2	2M
U04	Compute node	G0-3	G5-3	2M
U03	Compute node	G0-2	G5-4	2M
U02	Compute node	G0-1	G5-5	2M

E.3.2 Single Phase PDU Power Cabling Quarter Rack

This section contains the following tables:

- Table E-15

Table E-15 lists the location, ports and cables for the Single Phase PDU connections for an Exalogic machine quarter rack.

Table E-15 Single Phase PDU Power Cabling for Exalogic Machine Quarter Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U25	Cisco Switch	G3-1	G2-5	2M
U23	NM2-GW-SW	G2-6	G3-0	2M
U22	ZS3-ES HN	G2-5	G3-1	2M
U21	ZS3-ES HN	G2-4	G3-2	2M
U17	DE2-24C DS	G2-3	G3-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U09	Compute node	G1-1	G4-5	2M
U08	Compute node	G1-0	G4-6	2M
U07	Compute node	G0-6	G5-0	2M
U06	Compute node	G0-5	G5-1	2M
U05	Compute node	G0-4	G5-2	2M
U04	Compute node	G0-3	G5-3	2M
U03	Compute node	G0-2	G5-4	2M
U02	Compute node	G0-1	G5-5	2M

E.3.3 Single Phase PDU Power Cabling Eighth Rack

This section contains the following tables:

- Table E-16

Table E-16 lists the location, ports and cables for the Single Phase PDU connections for an Exalogic machine eighth rack.

Table E-16 Single Phase PDU Power Cabling for Exalogic Machine Eighth Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U25	Cisco Switch	G3-1	G2-5	2M
U23	NM2-GW-SW	G2-6	G3-0	2M
U22	ZS3-ES HN	G2-5	G3-1	2M
U21	ZS3-ES HN	G2-4	G3-2	2M
U17	DE2-24C DS	G2-3	G3-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U05	Compute node	G0-4	G5-2	2M
U04	Compute node	G0-3	G5-3	2M
U03	Compute node	G0-2	G5-4	2M
U02	Compute node	G0-1	G5-5	2M

E.3.4 Three Phase PDU Power Cabling Full Rack

This section contains the following tables:

- Table E-17

Table E-17 lists the location, ports and cables for the Three Phase PDU connections for an Exalogic machine full rack.

Table E-17 Three Phase PDU Power Cabling for Exalogic Machine Full Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U42	Compute node	G5-5	G2-1	2M
U41	Compute node	G5-4	G2-2	2M
U40	Compute node	G5-3	G2-3	2M
U39	Compute node	G5-2	G2-4	2M
U38	Compute node	G5-1	G2-5	2M
U37	Compute node	G5-0	G2-6	2M
U36	Compute node	G4-6	G1-0	2M
U35	Compute node	G4-5	G1-1	2M
U34	Compute node	G4-4	G1-2	2M
U33	Compute node	G4-3	G1-3	2M
U32	Compute node	G4-2	G1-4	2M
U31	Compute node	G4-1	G1-5	2M
U30	Compute node	G3-6	G0-0	2M
U29	Compute node	G3-5	G0-1	2M
U28	Compute node	G3-4	G0-2	2M
U27	Compute node	G3-3	G0-3	2M
U26	NM2-GW SW	G3-2	G0-4	2M
U25	Cisco Switch	G3-1	G0-5	2M
U24	NM2-GW SW	G2-6	G5-0	2M
U23	NM2-GW-SW	G2-5	G5-1	2M
U22	ZS3-ES HN	G3-0	G0-6	2M
U21	ZS3-ES HN	G2-4	G5-2	2M
U17	DE2-24C DS	G2-3	G5-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U15	Compute node	G2-2	G5-4	2M
U14	Compute node	G2-1	G5-5	2M
U13	Compute node	G2-0	G5-6	2M
U12	Compute node	G1-4	G4-2	2M
U11	Compute node	G1-3	G4-3	2M
U10	Compute node	G1-2	G4-4	2M

Table E-17 (Cont.) Three Phase PDU Power Cabling for Exalogic Machine Full Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U09	Compute node	G1-1	G4-5	2M
U08	Compute node	G1-0	G4-6	2M
U07	Compute node	G0-6	G3-0	2M
U06	Compute node	G0-5	G3-1	2M
U05	Compute node	G0-4	G3-2	2M
U04	Compute node	G0-3	G3-3	2M
U03	Compute node	G0-2	G3-4	2M
U02	Compute node	G0-1	G3-5	2M

E.3.5 Three Phase PDU Power Cabling Half Rack

This section contains the following tables:

- Table E-18

Table E-18 lists the location, ports and cables for the Three Phase PDU connections for an Exalogic machine half rack.

Table E-18 Three Phase PDU Power Cabling for Exalogic Machine Half Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U28	Compute node	G3-4	G0-2	2M
U27	Compute node	G3-3	G0-3	2M
U25	Cisco Switch	G3-1	G0-5	2M
U23	NM2-GW-SW	G2-5	G5-1	2M
U22	ZS3-ES HN	G3-0	G0-6	2M
U21	ZS3-ES HN	G2-4	G5-2	2M
U17	DE2-24C DS	G2-3	G5-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U15	Compute node	G2-2	G5-4	2M
U14	Compute node	G2-1	G5-5	2M
U13	Compute node	G2-0	G5-6	2M
U12	Compute node	G1-4	G4-2	2M
U11	Compute node	G1-3	G4-3	2M
U10	Compute node	G1-2	G4-4	2M
U09	Compute node	G1-1	G4-5	2M
U08	Compute node	G1-0	G4-6	2M
U07	Compute node	G0-6	G3-0	2M

Table E-18 (Cont.) Three Phase PDU Power Cabling for Exalogic Machine Half Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U06	Compute node	G0-5	G3-1	2M
U05	Compute node	G0-4	G3-2	2M
U04	Compute node	G0-3	G3-3	2M
U03	Compute node	G0-2	G3-4	2M
U02	Compute node	G0-1	G3-5	2M

E.3.6 Three Phase PDU Power Cabling Quarter Rack

This section contains the following tables:

- Table E-19

Table E-19 lists the location, ports and cables for the Three Phase PDU connections for an Exalogic machine quarter rack.

Table E-19 Three Phase PDU Power Cabling for Exalogic Machine Quarter Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U25	Cisco Switch	G3-1	G0-5	2M
U23	NM2-GW-SW	G2-5	G5-1	2M
U22	ZS3-ES HN	G3-0	G0-6	2M
U21	ZS3-ES HN	G2-4	G5-2	2M
U17	DE2-24C DS	G2-3	G5-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U09	Compute node	G1-1	G4-5	2M
U08	Compute node	G1-0	G4-6	2M
U07	Compute node	G0-6	G3-0	2M
U06	Compute node	G0-5	G3-1	2M
U05	Compute node	G0-4	G3-2	2M
U04	Compute node	G0-3	G3-3	2M
U03	Compute node	G0-2	G3-4	2M
U02	Compute node	G0-1	G3-5	2M

E.3.7 Three Phase PDU Power Cabling Eighth Rack

This section contains the following tables:

- Table E-20

Table E-20 lists the location, ports and cables for the Three Phase PDU connections for an Exalogic machine eighth rack.

Table E-20 Three Phase PDU Power Cabling for Exalogic Machine Eighth Rack

From Rack Unit	Type of Equipment	PDU-A	PDU-B	Jumper Length
U25	Cisco Switch	G3-1	G0-5	2M
U23	NM2-GW-SW	G2-5	G5-1	2M
U22	ZS3-ES HN	G3-0	G0-6	2M
U21	ZS3-ES HN	G2-4	G5-2	2M
U17	DE2-24C DS	G2-3	G5-3	1M
U16	NM2-GW-SW	G1-6	G4-0	2M
U05	Compute node	G0-4	G3-2	2M
U04	Compute node	G0-3	G3-3	2M
U03	Compute node	G0-2	G3-4	2M
U02	Compute node	G0-1	G3-5	2M

F
Manage Solaris Zones on Exalogic

Oracle Solaris zones are an integral part of the Oracle Solaris operating system. Zones isolate software applications and services using flexible software-defined boundaries. This appendix describes how to manage Solaris zones on Exalogic. This appendix contains the following sections:

F.1 Requirements

Creating zones on an Exalogic machine has the following requirements:

- An Exalogic machine imaged to release 2.0.4.0 running Oracle Solaris.
- An Exalogic machine patched to the April 2013 Patch Set Update available in the My Oracle Support document ID 1545364.1.
- An Exalogic machine patched with the Solaris patch for Zones on Shared Storage (ZOSS) over iSCSI available in the My Oracle Support document ID 16514816.

F.2 Terminology

Table F-1 describes the terms used in this appendix.

Table F-1 Terminology

Term	Description
Logical Unit	A logical unit is a component of a storage system. A logical unit is uniquely numbered creating a Logical Unit Number (LUN). The storage appliance can contain many LUNs. LUNs when associated with one or more SCSI targets, form a unique SCSI device. This SCSI device can be accessed by one or more SCSI initiators.
Initiator	An initiator is an application or production system end-point that is capable of initiating a SCSI session, sending SCSI commands, and I/O requests. Initiators are also identified by unique addressing methods.
Initiator Group	A set of initiators. When an initiator group is associated with a LUN, only initiators from that group may access the LUN.
Target	A target is an end-point that provides a service of processing SCSI commands and I/O requests from an initiator. A target, once configured, consists of zero or more logical units.

Table F-1 (Cont.) Terminology

Term	Description
Target Group	A set of targets. LUNs are exported over all the targets in one specific target group.

F.3 Create a Solaris Zone

This section describes how to create a Solaris zone. This section contains the following topics:

- Prerequisites
- Set Up a Solaris Zone

F.3.1 Prerequisites

Before creating a Solaris zone, you should perform the following tasks:

- Create an iSCSI Target
- Create an iSCSI Initiator
- Create the Project and LUN
- Disable the Write Cache
- Format the LUN
- Set Up the Exclusive 10 GbE Network for the Zone

F.3.1.1 Create an iSCSI Target

You can create an iSCSI target by doing the following:

1. Log in to the storage appliance BUI as the `root` user.
2. Click the **Configuration** tab.
3. Click **SAN**.
4. Click **iSCSI Targets**.
5. To create a new iSCSI target, click the plus button next to **iSCSI Targets**.

   The New iSCSI Target dialog box appears.
6. For the Target IQN, select the **Auto-assign** option.
7. In the **Alias** field, enter a name for your iSCSI target.
8. For the **Initiator authentication mode**, select the authentication mode you are using for communication between the compute node and storage appliance. By default, no authentication is used.

> **Note:**
>
> For more information on setting up CHAP authentication between the compute node and storage, see the "Setting Up CHAP Authentication" topic in the following document: `http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/iscsi-quickstart-v1-2-051512-1641594.pdf`

9. From the **Network interfaces** list, select the interface that corresponds to your InfiniBand partition. You can identify the interface by logging in to the storage appliance and running `configuration net interfaces show`. If there are no partitions defined, identify the interface for the label `IB_Interface`.
10. Click **OK**.
11. You can add the iSCSI target to an iSCSI target group by dragging and dropping the target to a iSCSI target group in the iSCSI Target Groups panel on the right. If required, you can create a new iSCSI target group by dragging and dropping the target to the top of iSCSI Target Groups panel on the right.

F.3.1.2 Create an iSCSI Initiator

You can create an iSCSI initiator by doing the following:

1. Before you can create an iSCSI initiator, you must identify an initiator IQN. The initiator IQN is a unique reference number associated with a specific compute node. To find the initiator IQN for a compute node, do the following:
   a. Log in to an Exalogic compute node.
   b. Run the `iscsiadm list initiator-node` command as follows:

```
# iscsiadm list initiator-node
Initiator node name: iqn.1986-03.com.sun:01:e00000000000.51891a8b
Initiator node alias: el01cn01
        Login Parameters (Default/Configured):
                Header Digest: NONE/-
                Data Digest: NONE/-
                Max Connections: 65535/-
        Authentication Type: NONE
        RADIUS Server: NONE
        RADIUS Access: disabled
        Tunable Parameters (Default/Configured):
                Session Login Response Time: 60/-
                Maximum Connection Retry Time: 180/-
                Login Retry Time Interval: 60/-
        Configured Sessions: 1
```

   In this example, the initiator IQN is:

```
iqn.1986-03.com.sun:01:e00000000000.51891a8b
```

2. Log in to the storage appliance BUI as the `root` user.
3. Click the **Configuration** tab.
4. Click **SAN**.
5. Click **Initiators**.

6. Click **iSCSI Initiators**.
7. Click the plus button next to **iSCSI Initiators** to create a new iSCSI initiator.
8. In the **Initiator IQN** field, enter the initiator IQN you identified in step 1.
9. In the **Alias** field, enter a name for the iSCSI initiator you are creating.
10. If you are using CHAP authentication, select the **Use CHAP** check box and fill in the **Initiator CHAP name** and **Initiator CHAP secret** fields as you did in Create an iSCSI Target.

> **Note:**
>
> For more information on setting up CHAP authentication between the compute node and storage appliance, see the "Setting Up CHAP Authentication" topic in the following document: `http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/iscsi-quickstart-v1-2-051512-1641594.pdf`

11. Click **OK**.
12. Add the iSCSI initiator to an iSCSI initiator group by dragging and dropping the initiator.

    If required, you can create a new iSCSI initiator group.

F.3.1.3 Create the Project and LUN

You can create the project and LUN by doing the following:

1. Create a project as described in Create Custom Projects.
2. You can create the LUN by doing the following:
   a. Next to the **Project**, click **Shares**.
   b. Click **LUNs**.

      The list of LUNs appears.
   c. Click the plus button next to LUNs.
   d. In the **Project** field, select the project you created in step 1.
   e. In the **Name** field, enter a name for the LUN.
   f. Enter the size of the volume in GB.
   g. Select **Thin provisioned**.
   h. Set the **Volume block size** as 32k.
   i. In the **Target Group** field, select the target group you used in Create an iSCSI Target.
   j. In the **Initiator Group** field, select the initiator group you used in Create an iSCSI Initiator.
   k. Click **Apply**.
   l. Note the GUID of the LUN you created in the list of LUNs. For example, g600144f09c96cca900005190bfc4000a.

Note:

After creating the LUN, ensure that the **Write cache enabled** check box is deselected. You can find this check box in the **Protocols** tab of the LUN.

F.3.1.4 Disable the Write Cache

You must disable the write cache on the LUN permanently by doing the following:

1. Log in to the compute node for which you identified the initiator node name as described in Create an iSCSI Initiator.
2. Edit the `/kernel/drv/sd.conf` file.
3. Add the following to the `sd.conf` file:

```
sd-config-list="SUN     ZFS Storage 7120","write-cache-disable",
               "SUN     ZFS Storage 7320","write-cache-disable",
               "SUN     ZFS Storage 7420","write-cache-disable",
               "SUN     ZFS Storage 7335","write-cache-disable";
write-cache-disable=1,0x00008,0,0,0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0;
```

4. Restart the compute node by running the `reboot` command.

F.3.1.5 Format the LUN

Before the compute node can use the LUN, you must format the LUN. You can format the LUN by doing the following:

1. Log in to a compute node as the `root` user.
2. Run the `iscsiadm` commands to discover the iSCSI targets from the compute node:

```
# iscsiadm add discovery-address IPoIB_address_of_the_storage_appliance
# iscsiadm modify discovery -t enable
```

   In this example, `IPoIB_address_of_the_storage_appliance` is the IP address of the storage appliance on the IPoIB network.
3. Run the following command to load drivers, attach device instances, create logical links to device nodes, and load the device policy for iSCSI:

```
# devfsadm -c iscsi
```

4. Identify the disk you should format and label by running `echo | format` as follows:

```
# echo | format
Searching for disks...done


AVAILABLE DISK SELECTIONS:
       0. c0t600144F09C96CCA90000518CDEB10005d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdeb10005
       1. c0t600144F09C96CCA90000518CDF100006d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdf100006
       2. c0t600144F09C96CCA90000518CDFB60007d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdfb60007
```

```
       3. c0t600144F09C96CCA900005190BFC4000Ad0 <SUN-ZFS Storage 7320-1.0 cyl
8352 alt 2 hd 255 sec 63>
          /scsi_vhci/disk@g600144f09c96cca900005190bfc4000a
       4. c7t0d0 <LSI-MR9261-8i-2.12-28.87GB>
          /pci@0,0/pci8086,340a@3/pci1000,9263@0/sd@0,0
Specify disk (enter its number): Specify disk (enter its number):
```

The value after `/scsi_vhci/disk@g` is the GUID of the LUN you created in Create the Project and LUN. In this example, the disk `c0t600144F09C96CCA900005190BFC4000Ad0` with the GUID `g600144f09c96cca900005190bfc4000a` should be formatted and labelled.

5. Format the disk by doing the following:
   a. Run the **format** command to start formatting the disk as follows:

      ```
      # format -e c0t600144F09C96CCA900005190BFC4000Ad0
      selecting c0t600144F09C96CCA900005190BFC4000Ad0
      [disk formatted]
      ```

      The `format` prompt appears.
   b. Enter `fdisk` to manipulate the partition tables as follows:

      ```
      format> fdisk
      No fdisk table exists. The default partition for the disk is:

        a 100% "SOLARIS System" partition
      ```
   c. When prompted, enter **n** to edit the partition table.

      ```
      Type "y" to accept the default partition, otherwise type "n" to edit the
      partition table. n
      ```
   d. Enter 1 to set the partition type.
   e. Enter f to set the partition type as EFI (Protective) as follows:

      ```
      Select the partition type to create:
         1=SOLARIS2    2=UNIX      3=PCIXOS     4=Other        5=DOS12
         6=DOS16       7=DOSEXT    8=DOSBIG     9=DOS16LBA     A=x86 Boot
         B=Diagnostic  C=FAT32     D=FAT32LBA   E=DOSEXTLBA    F=EFI (Protective)
         G=EFI_SYS     0=Exit? f
      ```
6. Label the LUN by doing the following:
   a. Enter 6 to label the LUN.

      The `format` prompt appears.
   b. Enter `label` to label the disk as follows:

      ```
      format> label
      ```

      The list of label types appears.
   c. Enter 1 to specify the label type as an EFI label as follows:

      ```
      [0] SMI Label
      [1] EFI Label
      Specify Label type[1]: 1
      ```

      A confirmation message appears.
   d. Enter `y` to continue.
   e. Enter `quit` to exit the format prompt.

f. You can use the `format` command to ensure the disk is available and the same size you specified in the storage appliance BUI as follows:

```
# format
Searching for disks...done


AVAILABLE DISK SELECTIONS:
       0. c0t600144F09C96CCA90000518CDEB10005d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdeb10005
       1. c0t600144F09C96CCA90000518CDF100006d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdf100006
       2. c0t600144F09C96CCA90000518CDFB60007d0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca90000518cdfb60007
       3. c0t600144F09C96CCA900005190BFC4000Ad0 <SUN-ZFS Storage
7320-1.0-64.00GB>
          /scsi_vhci/disk@g600144f09c96cca900005190bfc4000a
       4. c7t0d0 <LSI-MR9261-8i-2.12-28.87GB>
          /pci@0,0/pci8086,340a@3/pci1000,9263@0/sd@0,0
Specify disk (enter its number):
```

F.3.1.6 Set Up the Exclusive 10 GbE Network for the Zone

The zone you want to create should be given access to an exclusive network. You should create the necessary VNICs for the zone by doing the following:

1. Create a VLAN and VNIC. Follow the procedure in Oracle Solaris: Creating VNICs and Associating Them with VLANs..
2. Log in to the compute node as the `root` user.
3. Run the `dladm show-phys` command to identify the physical links of the EoIB devices as in the following example:

```
# dladm show-phys
LINK              MEDIA                STATE      SPEED  DUPLEX    DEVICE
net6              Infiniband           up         32000  unknown   ibp1
net0              Ethernet             up         1000   full      igb0
net1              Ethernet             unknown    0      unknown   igb1
net3              Ethernet             unknown    0      unknown   igb3
net4              Ethernet             up         10     full      usbecm0
net8              Ethernet             up         10000  full      eoib1
net2              Ethernet             unknown    0      unknown   igb2
net5              Infiniband           up         32000  unknown   ibp0
net9              Ethernet             up         10000  full      eoib0
```

   In this example, `net8` and `net9` are the physical links of the EoIB devices.

4. Create a VNIC on the compute node for the first physical link using the `dladm create-vnic` command as follows:

```
# dladm create-vnic -l link_of_eoib0 -v VLAN_ID vnic1_name
```

   Example:

```
# dladm create-vnic -l net9 -v 1706 vnic3_1706
```

5. Create a VNIC for the second physical link using the `dladm create-vnic` command as follows:

```
# dladm create-vnic -l link_of_eoib1 -v VLAN_ID vnic2_name
```

Example:

```
# dladm create-vnic -l net8 -v 1706 vnic2_1706
```

F.3.2 Set Up a Solaris Zone

With the storage appliance prepared, you can store the zone on the storage appliance and set up additional bonded network on the 10 GbE Exalogic client network exclusively for the zone.

You can set up the solaris zone by doing the following:

1. Create a Zone
2. Install and Boot Up the Zone

F.3.2.1 Create a Zone

You can create a zone by doing the following:

1. Log in to the compute node as the `root` user.
2. Run the `zonecfg` command to configure the zone as follows:

   ```
   # zonecfg -z zone_name
   ```

   Example:

   ```
   # zonecfg -z zone04
   Use 'create' to begin configuring a new zone.
   ```

   In this example, the name of the zone you are creating is `zone04`.
3. Enter `create` to begin configuring the zone:

   ```
   zonecfg:zone04 create
   create: Using system default template 'SYSdefault'
   ```

4. Create the zone by running the following commands:

   ```
   zonecfg:zone04> set zonepath=/zones/zone04
   zonecfg:zone04> add rootzpool
   zonecfg:zone04:rootzpool> add storage iscsi://
   IPoIB_Address_of_the_storage_Appliance/luname.naa.LUNGUID
   zonecfg:zone04:rootzpool> end
   zonecfg:zone04> remove anet
   zonecfg:zone04> add net
   zonecfg:zone04:net> set physical=vnic1_name
   zonecfg:zone04:net> end
   zonecfg:zone04> add net
   zonecfg:zone04:net> set physical=vnic2_name
   zonecfg:zone04:net> end
   zonecfg:zone04> verify
   zonecfg:zone04> commit
   ```

   In this example:

   - `IPoIB_address_of_the_storage_appliance` is the IP address of the storage appliance on the IPoIB network.
   - `LUNGUID` is the GUID of the LUN you created in Create the Project and LUN.

- vnic1_name and vnic2_name are the VNICs you created in Set Up the Exclusive 10 GbE Network for the Zone.

Example:

```
zonecfg:zone04> set zonepath=/zones/zone04
zonecfg:zone04> add rootzpool
zonecfg:zone04:rootzpool> add storage iscsi://192.168.14.133/luname.naa.
600144f09c96cca900005190bfc4000a
zonecfg:zone04:rootzpool> end
zonecfg:zone04> remove anet
zonecfg:zone04> add net
zonecfg:zone04:net> set physical=vnic2_1706
zonecfg:zone04:net> end
zonecfg:zone04> add net
zonecfg:zone04:net> set physical=vnic3_1706
zonecfg:zone04:net> end
zonecfg:zone04> verify
zonecfg:zone04> commit
```

5. You can verify the details of the zone by running the `info` command as follows:

```
zonecfg:zone04> info
zonename: zone04
zonepath: /zones/zone04
brand: solaris
autoboot: false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
net:
    address not specified
    allowed-address not specified
    configure-allowed-address: true
    physical: vnic1_name
    defrouter not specified
net:
    address not specified
    allowed-address not specified
    configure-allowed-address: true
    physical: vnic2_name
    defrouter not specified
rootzpool:
    storage: iscsi://IPoIB_Address_of_the_Storage_Appliance/luname.naa.LUNGUID
zonecfg:zone04>
```

F.3.2.2 Install and Boot Up the Zone

Before installing the zone, ensure that you have a repository for the Solaris installation set up stored on the storage appliance. The zone creation uses this repository to store the operating system files for the zone.

1. Install the zone by running the command as follows:

```
# zoneadm -z zone04 install

Configured zone storage resource(s) from:
```

```
    iscsi://192.168.14.133/luname.naa.600144f09c96cca900005190bfc4000a
Created zone zpool: zone04_rpool
Progress being logged to /var/log/zones/zoneadm.20130513T104657Z.zone04.install
       Image: Preparing at /zones/zone04/root.

 AI Manifest: /tmp/manifest.xml.lPaGVo
  SC Profile: /usr/share/auto_install/sc_profiles/enable_sci.xml
    Zonename: zone04
Installation: Starting ...

              Creating IPS image
Startup linked: 1/1 done
              Installing packages from:
                  exa-family
                    origin:                                     http://localhost:
1008/exa-family/acbd22da328c302a86fb9f23d43f5d10f13cf5a6/
                  solaris
                      origin:  http://install1/release/solaris/
DOWNLOAD                                PKGS         FILES    XFER (MB)   SPEED
Completed                            185/185   34345/34345  229.7/229.7 10.6M/s

PHASE                                          ITEMS
Installing new actions                   48269/48269
Updating package state database                 Done
Updating image state                            Done
Creating fast lookup database                   Done
Installation: Succeeded

        Note: Man pages can be obtained by installing pkg:/system/manual

 done.

        Done: Installation completed in 81.509 seconds.


  Next Steps: Boot the zone, then log into the zone console (zlogin -C)

              to complete the configuration process.

Log saved in non-global zone as /zones/zone04/root/var/log/zones/zoneadm.
20130513T104657Z.zone04.install
```

2. Boot up the zone by running the following command:

```
# zoneadm -z zone04 boot
```

3. Once the zone has booted up, log in to the zone using the `zlogin` command as follows:

```
# zlogin zone04
[Connected to zone 'zone04' pts/7]
```

4. Bond the VNICs you created as described in step 5 of Oracle Solaris: Creating VNICs and Associating Them with VLANs..

5. Run the following command to display the bond you created in the previous step:

```
root@zone04:~# ipadm show-addr
ADDROBJ           TYPE     STATE        ADDR
lo0/v4            static   ok           127.0.0.1/8
bond1/v4          static   ok           138.3.51.2/22
lo0/v6            static   ok           ::1/128
```

Note the IP address of the bond you created.

6. Run `netstat -rn` to display the routing table as in the following example:

```
root@zone04:~# netstat -rn

Routing Table: IPv4
  Destination           Gateway            Flags  Ref     Use     Interface
-------------------- -------------------- ----- ----- ---------- ---------
127.0.0.1            127.0.0.1            UH        2          0 lo0
138.3.48.0           138.3.51.2           U         2          0 bond1

Routing Table: IPv6
  Destination/Mask            Gateway                    Flags Ref   Use    If
--------------------------- --------------------------- ----- --- ------- -----
::1                         ::1                         UH      2       0 lo0
```

7. Add the IP address of the bond you noted in step 5 by running the following command:

```
root@zone04:~# route -p add default IP_address_of_bond
```

Example:

```
root@zone04:~# route -p add default 138.3.48.1
add net default: gateway 138.3.48.1
add persistent net default: gateway 138.3.48.1
```

8. Display the routing table again to verify that the IP address of the bond was added as in the following example:

```
root@zone04:~# netstat -rn

Routing Table: IPv4
  Destination           Gateway            Flags  Ref     Use     Interface
-------------------- -------------------- ----- ----- ---------- ---------
default              138.3.48.1           UG        1          0
127.0.0.1            127.0.0.1            UH        2          0 lo0
138.3.48.0           138.3.51.2           U         2          0 bond1

Routing Table: IPv6
  Destination/Mask            Gateway                    Flags Ref   Use    If
--------------------------- --------------------------- ----- --- ------- -----
::1                         ::1                         UH      2       0 lo0
```

F.4 Migrate a Zone to a New Host

You can migrate a zone from one physical host to another by running the following procedure:

> **Note:**
>
> The zone is shut down during the migration process. If you require high availability, ensure you use a clustered software solution.

1. Log in to the compute node hosting the zone as the `root` user.
2. Shutdown the zone by running the following command:

```
# zoneadm -z name_of_zone shutdown
```

Example:

```
# zoneadm -z zone04 shutdown
```

3. Detach the zone by running the following command:

```
# zoneadm -z name_of_zone detach
```

Example:

```
# zoneadm -z zone04 detach
zoneadm: zone 'zone04': warning(s) occured during processing URI: 'iscsi://
192.168.14.133/luname.naa.600144f09c96cca900005190bfc4000a'
Could not remove one or more iSCSI discovery addresses because logical unit is
in use
Exported zone zpool: zone04_rpool
Unconfigured zone storage resource(s) from:
        iscsi://192.168.14.133/luname.naa.600144f09c96cca900005190bfc4000a
#
```

4. Create a directory on the storage appliance to which you can export the configuration of the zone:

```
# mkdir -p directory
```

Example:

```
# mkdir -p /u01/common/general/zone04
```

5. Export the configuration of the zone by running the following command:

```
# zonecfg -z name_of_zone export > directory/name_of_zone.cfg
```

Example:

```
# zonecfg -z zone04 export > /common/general/zone04/zone04.cfg
```

6. Log in to the compute node you want to migrate the zone to as the `root` user.

7. Import the zone from the configuration file you created in the previous step by running the following command:

```
# zonecfg -z name_of_zone -f directory/name_of_zone.cfg
```

Example:

```
# zonecfg -z zone04 -f /common/general/zone04/zone04.cfg
```

8. Attach the zone by running the following command:

```
# zoneadm -z name_of_zone attach
```

Example:

```
# zoneadm -z zone04 attach
Configured zone storage resource(s) from:
    iscsi://192.168.14.133/luname.naa.600144f09c96cca900005190bfc4000a
Imported zone zpool: zone04_rpool
Progress being logged to /var/log/zones/zoneadm.20130513T135704Z.zone04.attach
    Installing: Using existing zone boot environment
      Zone BE root dataset: zone04_rpool/rpool/ROOT/solaris
                     Cache: Using /var/pkg/publisher.
  Updating non-global zone: Linking to image /.
Processing linked: 1/1 done
```

```
  Updating non-global zone: Auditing packages.
No updates necessary for this image.

  Updating non-global zone: Zone updated.
                    Result: Attach Succeeded.
Log saved in non-global zone as /zones/zone04/root/var/log/zones/zoneadm.
20130513T135704Z.zone04.attach
```

9. Boot up the zone by running the following command:

```
# zoneadm -z name_of_zone boot
```

Example:

```
# zoneadm -z zone04 boot
```

Note:

In some situations, the process of detaching and attaching can cause the server to boot up with the system configuration wizard running.

You can resolve this issue, by logging in to the console and completing the wizard. You can use the following command to log in to the zone:

```
# zlogin -C name_of_zone
```

G
Customize Linux on the Compute Nodes

This appendix describes how to customize the Linux operating system on compute nodes in an Exalogic machine, to suit your business needs, by installing additional RPMs, and updating or removing RPMs that are installed by default. It applies to Exalogic in a physical configuration, with the Linux operating system installed on bare metal. The tasks described in this appendix are optional.
This appendix contains the following sections:

G.1 RPMs That Must Not Be Modified or Removed

Do not modify or delete the following RPMs outside of an Exalogic upgrade, patch set update (PSU), or patch.

```
kernel*
compat-dapl*
dapl*
ib-bonding*
ibacm*
ibutils*
ibsim*
infiniband-diags*
libibcm*
libibmad*
libibumad*
libibverbs*
libmlx4*
librdmacm*
libsdp*
mpi-selector*
mpitests_openmpi_gcc*
mstflint*
mvapich*
ofa*
ofed*
openmpi_gcc*
opensm*
perftest*
qperf*
rds-tools*
sdpnetstat*
srptools*
exalogic*
infinibus*
```

G.2 Prepare the Compute Nodes for Yum Updates

1. On a host outside the Exalogic machine, set up a local yum repository as described in the "Server Setup" section of the following document:

   http://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html

   While doing this, select *only* the **`ol5_x86_64_latest`** channel.

2. On *each compute node* on the Exalogic machine, do the following:

   a. Log in as the `root` user.

   b. Set up the compute node as a `yum` client of the repository you set up in step 1, by performing the steps in the "Client Setup" section of the following document:

      http://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html

   c. To ensure that essential RPMs are not modified during `yum` updates, append the following `exclude` directive, to the `/etc/yum.conf` file.

      Note:

      The entire directive must be on one line.

      ```
      exclude=kernel* compat-dapl* dapl* ib-bonding* ibacm* ibutils* ibsim*
      infiniband-diags* libibcm* libibmad* libibumad* libibverbs* libmlx4*
      librdmacm* libsdp* mpi-selector* mpitests_openmpi_gcc* mstflint* mvapich*
      ofa* ofed* openmpi_gcc* opensm* perftest* qperf* rds-tools* sdpnetstat*
      srptools* exalogic* infinibus*
      ```

      This exclusion list includes all of the RPMs listed in RPMs That Must Not Be Modified or Removed.

G.3 Install, Update, and Remove RPMs Using Yum

1. Ensure that the compute node on which you want to perform `yum` updates has been prepared, as described in Prepare the Compute Nodes for Yum Updates.

2. Install, update, or remove RPMs as follows:

   - To install an RPM from the `yum` repository, run the following command:

     ```
     # yum install rpm_name
     ```

   - To update a specific RPM, run the following command:

     ```
     # yum update rpm_name
     ```

   - To update all the RPMs, run the following command:

     ```
     # yum update
     ```

   - To remove an RPM, run the following command:

     ```
     # yum remove rpm_name
     ```

For the `yum update` or `yum remove` command, if you specify an RPM that is in the exclusion list defined earlier, the command will fail.