SPARC SuperCluster T4-4 Owner's Guide

Part No: E21660-10
February 2016

Part No: E21660-10

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.

Référence: E21660-10

Accessibilité de la documentation

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité à la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info ou le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs si vous êtes malentendant.

Contents

Using This Documentation

This document provides an overview of SPARC SuperCluster T4-4, and describes configuration options, site preparation specifications, installation information, and administration tools.

- **Overview** – Describes how to configure, install, tune, and monitor the system.
- **Audience** – Technicians, system administrators, and authorized service providers.
- **Required knowledge** – Advanced experience in system installation and administration.

Product Documentation Library

Documentation and resources for this product and related products are available at: http://docs.oracle.com/cd/E21659_01

Feedback

Provide feedback about this documentation at:

http://www.oracle.com/goto/docfeedback

Understanding the System

These topics describe the features and hardware components of the SPARC SuperCluster T4-4. These topics also describe the different software configurations that are available.

SPARC SuperCluster T4-4 Overview

The SPARC SuperCluster T4-4 is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and widely varied workloads. The SPARC SuperCluster T4-4 is intended for large-scale, performance-sensitive, mission-critical application deployments. The SPARC SuperCluster T4-4 combines industry-standard hardware and clustering software, such as optional Oracle Database 11g Real Application Clusters (Oracle RAC) and optional Oracle Solaris Cluster software. This combination enables a high degree of isolation between concurrently deployed applications, which have varied security, reliability, and performance requirements. The SPARC SuperCluster T4-4 enables customers to develop a single environment that can support end-to-end consolidation of their entire applications portfolio.

The SPARC SuperCluster T4-4 provides an optimal solution for all database workloads, ranging from scan-intensive data warehouse applications to highly concurrent online transaction processing (OLTP) applications. With its combination of smart Oracle Exadata Storage Server Software, complete and intelligent Oracle Database software, and the latest industry-standard hardware components, the SPARC SuperCluster T4-4 delivers extreme performance in a highly-available, highly-secure environment. Oracle provides unique clustering and workload management capabilities so the SPARC SuperCluster T4-4 is well-suited for consolidating multiple databases into a single grid. Delivered as a complete pre-optimized, and pre-configured package of software, servers, and storage, the SPARC SuperCluster T4-4 is fast to implement, and it is ready to tackle your large-scale business applications.

The SPARC SuperCluster T4-4 does not include any Oracle software licenses. Appropriate licensing of the following software is required when the SPARC SuperCluster T4-4 is used as a database server:

- Oracle Database Software
- Oracle Exadata Storage Server Software

In addition, Oracle recommends that the following software is licensed:

- Oracle Real Application Clusters
- Oracle Partitioning

The SPARC SuperCluster T4-4 is designed to fully leverage an internal InfiniBand fabric that connects all of the processing, storage, memory, and external network interfaces within the SPARC SuperCluster T4-4 to form a single, large computing device. Each SPARC SuperCluster T4-4 is connected to data center networks through 10-GbE (traffic) and 1-GbE (management) interfaces.

You can integrate SPARC SuperCluster T4-4 systems with Exadata or Exalogic machines by using the available InfiniBand expansion ports and optional data center switches. The InfiniBand technology used by the SPARC SuperCluster T4-4 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 the SPARC SuperCluster T4-4. You can connect many SPARC SuperCluster T4-4 systems, or a combination of SPARC SuperCluster T4-4 systems and Oracle Exadata Database Machines, to develop a single, large-scale environment. You can integrate SPARC SuperCluster T4-4 systems with their current data center infrastructure using the available 10 GbE ports in each SPARC T4-4 server.

Spares Kit Components

The SPARC SuperCluster T4-4 includes a spares kit that includes the following components:

- 1 x 600 GB 15 K RPM High Performance SAS disk or 1 x 3 TB 7.2 K RPM High Capacity SAS disk
- 1 x 96 GB Exadata Smart Flash Cache card

SPARC SuperCluster T4-4 Restrictions

The following restrictions apply to hardware and software modifications to the SPARC SuperCluster T4-4. Violating these restrictions can result in loss of warranty and support.

- SPARC SuperCluster T4-4 hardware cannot be modified or customized. There is one exception to this. The only allowed hardware modification to the SPARC SuperCluster

T4-4 is to the administrative 48-port Cisco 4948 Gigabit Ethernet switch included with the SPARC SuperCluster T4-4. Customers may choose to the following:

- Replace the Gigabit Ethernet switch, at customer expense, with an equivalent 1U 48-port Gigabit Ethernet switch that conforms to their internal data center network standards. This replacement must be performed by the customer, at their expense and labor, after delivery of the SPARC SuperCluster T4-4. If the customer chooses to make this change, then Oracle cannot make or assist with this change given the numerous possible scenarios involved, and it is not included as part of the standard installation. The customer must supply the replacement hardware, and make or arrange for this change through other means.
- Remove the CAT5 cables connected to the Cisco 4948 Ethernet switch, and connect them to the customer's network through an external switch or patch panel. The customer must perform these changes at their expense and labor. In this case, the Cisco 4948 Ethernet switch in the rack can be turned off and unconnected to the data center network.

- The Oracle Exadata Storage Expansion Rack can only be connected to a SPARC SuperCluster T4-4 or an Oracle Exadata Database Machine, and only supports databases running on the Oracle Database (DB) Domains in the SPARC SuperCluster T4-4 or on the database servers in the Oracle Exadata Database Machine.
- Standalone Exadata Storage Servers can only be connected to a SPARC SuperCluster T4-4 or an Oracle Exadata Database Machine, and only support databases running on the Database Domains in the SPARC SuperCluster T4-4 or on the database servers in the Oracle Exadata Database Machine. The standalone Exadata Storage Servers must be installed in a separate rack.
- Earlier Oracle Database releases can be run in Application Domains running Oracle Solaris 10. Non-Oracle databases can be run in either Application Domains running Oracle Solaris 10 or Oracle Solaris 11, depending on the Oracle Solaris version they support.
- Oracle Exadata Storage Server Software and the operating systems cannot be modified, and customers cannot install any additional software or agents on the Exadata Storage Servers.
- Customers cannot update the firmware directly on the Exadata Storage Servers. The firmware is updated as part of an Exadata Storage Server patch.
- Customers may load additional software on the Database Domains on the SPARC T4-4 servers. However, to ensure best performance, Oracle discourages adding software except for agents, such as backup agents and security monitoring agents, on the Database Domains. Loading non-standard kernel modules to the operating system of the Database Domains is allowed but discouraged. Oracle will not support questions or issues with the non-standard modules. If a server crashes, and Oracle suspects the crash may have been caused by a non-standard module, then Oracle support may refer the customer to the vendor of the non-standard module or ask that the issue be reproduced without the non-standard module. Modifying the Database Domain operating system other than by applying official patches and upgrades is not supported. InfiniBand-related packages should always be maintained at the officially supported release.
- The SPARC SuperCluster T4-4 supports separate domains dedicated to applications, with high throughput/low latency access to the database domains through InfiniBand.

Since Oracle Database is by nature a client server, applications running in the Application Domains can connect to database instances running in the Database Domain. Applications can be run in the Database Domain, although it is discouraged.

- Customers cannot connect USB devices to the Exadata Storage Servers except as documented in *Oracle Exadata Storage Server Software User's Guide* and this guide. In those documented situations, the USB device should not draw more than 100 mA of power.
- The network ports on the SPARC T4-4 servers can be used to connect to external non-Exadata Storage Servers using iSCSI or NFS. However, the Fibre Channel Over Ethernet (FCoE) protocol is not supported.
- Only switches specified for use in the SPARC SuperCluster T4-4, Oracle Exadata Rack and Oracle Exalogic Elastic Cloud may be connected to the InfiniBand network. It is not supported to connect third-party switches and other switches not used in the SPARC SuperCluster T4-4, Oracle Exadata Rack and Oracle Exalogic Elastic Cloud.

Identifying Hardware Components

The SPARC SuperCluster T4-4 consists of SPARC T4-4 servers, Exadata Storage Servers, and Sun ZFS Storage 7320 appliances, as well as required InfiniBand and Ethernet networking components.

This section contains the following topics:

Full Rack Components

FIGURE 1 SPARC SuperCluster T4-4 Full Rack Layout, Front View

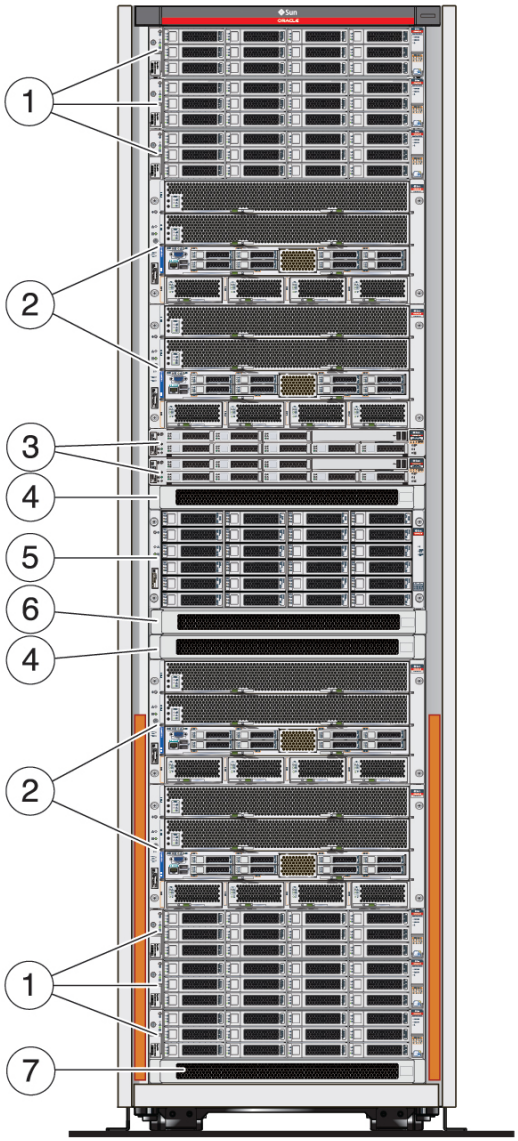

Figure Legend

1	Exadata Storage Servers (6)
2	SPARC T4-4 servers (4)
3	Sun ZFS Storage 7320 storage controllers (2)
4	Sun Datacenter InfiniBand Switch 36 leaf switches (2)
5	Sun Disk Shelf
6	Cisco Catalyst 4948 Ethernet management switch
7	Sun Datacenter InfiniBand Switch 36 spine switch

You can expand the amount of disk storage for your system using the Oracle Exadata Storage Expansion Rack. See "Oracle Exadata Storage Expansion Rack Components" on page 209 for more information.

You can connect up to eight SPARC SuperCluster T4-4s together, or a combination of SPARC SuperCluster T4-4s and Oracle Exadata or Exalogic machines on the same InfiniBand fabric, without the need for any external switches. See "Connecting Multiple SPARC SuperCluster T4-4 Systems" for more information.

Half Rack Components

FIGURE 2 SPARC SuperCluster T4-4 Half Rack Layout, Front View

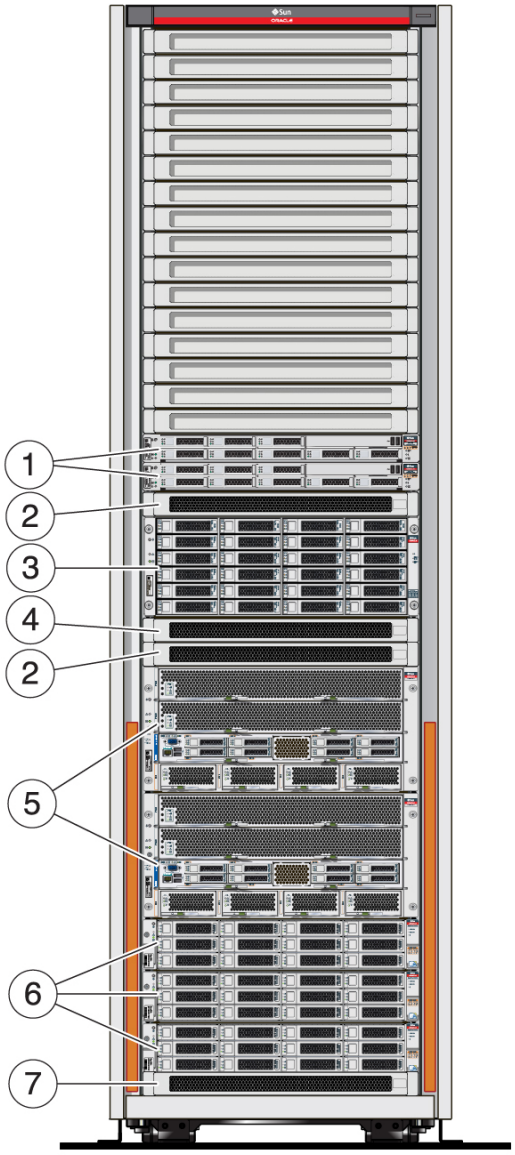

Figure Legend

1 Sun ZFS Storage 7320 storage controllers (2)
2 Sun Datacenter InfiniBand Switch 36 leaf switches (2)

3 Sun Disk Shelf
4 Cisco Catalyst 4948 Ethernet management switch
5 SPARC T4-4 servers (2)
6 Exadata Storage Servers (3)
7 Sun Datacenter InfiniBand Switch 36 spine switch

You can expand the amount of disk storage for your system using the Oracle Exadata Storage Expansion Rack. See "Oracle Exadata Storage Expansion Rack Components" on page 209 for more information.

You can connect up to eight SPARC SuperCluster T4-4s together, or a combination of SPARC SuperCluster T4-4s and Oracle Exadata or Exalogic machines on the same InfiniBand fabric, without the need for any external switches. See "Connecting Multiple SPARC SuperCluster T4-4 Systems" for more information.

Understanding the Hardware Components and Connections

These topics describe how the hardware components and connections are configured to provide full redundancy for high performance or high availability in the SPARC SuperCluster T4-4, as well as connections to the various networks:

Understanding the Hardware Components

The following SPARC SuperCluster T4-4 hardware components provide full redundancy, either through physical connections between components within the system, or through the components:

SPARC T4-4 Servers

The full rack version of the SPARC SuperCluster T4-4 contains four SPARC T4-4 servers. The half rack version of the SPARC SuperCluster T4-4 contains two SPARC T4-4 servers.

Redundancy in the SPARC T4-4 servers is achieved two ways:

- Through connections between the servers (described in "SPARC T4-4 Server Physical Connections" on page 24)
- Through components within the SPARC T4-4 servers:
  - **Power supplies** – Each SPARC T4-4 server contains four power supplies. The SPARC T4-4 server can continue to operate normally if one or two of those power supplies fails, or if one of the power distribution units fail.
  - **Fan modules** – Each SPARC T4-4 server contains five fan modules. The SPARC T4-4 server will continue to operate at full capacity if one of the fan modules fails.
  - **AC power connectors** – Each SPARC T4-4 server contains four AC power connectors at the rear of the server. Only two power connections are required for operation.
  - **Disk drives** – Each SPARC T4-4 server contains a combination of six hard disk drives and solid-state drives. The SPARC SuperCluster T4-4 software provides redundancy between the six disk drives.
  - **Processor modules** – Each SPARC T4-4 server contains two processor modules. In the case of a processor module failure, the system will recover and reconfigure around the failed module, preserving access to all disk drives and express modules.

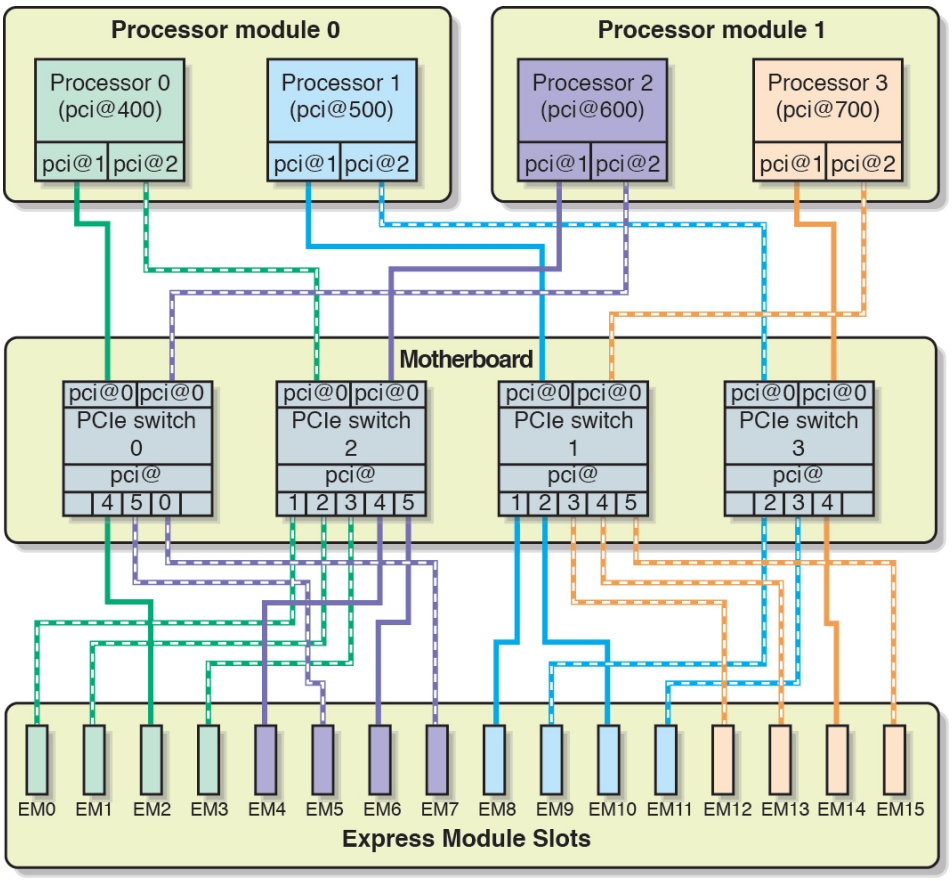

Exadata Storage Servers

The full rack version of the SPARC SuperCluster T4-4 contains six Exadata Storage Servers. The half rack version of the SPARC SuperCluster T4-4 contains three Exadata Storage Servers. Redundancy in the Exadata Storage Servers is achieved two ways:

- Through connections between the Exadata Storage Servers. For more information, see "Exadata Storage Server Physical Connections" on page 28.
- Through components within the Exadata Storage Servers:
  - Power supplies – Each Exadata Storage Server contains two power supplies. The Exadata Storage Server can continue to operate normally if one of the power supplies fail, or if one of the power distribution units fail.

- Disk drives – Each Exadata Storage Server contains 12 disk drives, where you can choose between disk drives designed for either high capacity or high performance when you first order the SPARC SuperCluster T4-4. The SPARC SuperCluster T4-4 software provides redundancy between the 12 disk drives within each Exadata Storage Server. For more information, see "Cluster Software for the Database Domain" on page 38.

Sun ZFS Storage 7320 Appliance

Each SPARC SuperCluster T4-4 contains one Sun ZFS Storage 7320 appliance, in either the full rack or half rack version. The Sun ZFS Storage 7320 appliance consists of the following:

- Two Sun ZFS Storage 7320 storage controllers
- One Sun Disk Shelf

Redundancy in the Sun ZFS Storage 7320 appliance is achieved two ways:

- Through connections from the two Sun ZFS Storage 7320 storage controllers to the Sun Disk Shelf. For more information, see "Sun ZFS Storage 7320 Appliance Physical Connections" on page 29.
- Through components within the Sun ZFS Storage 7320 appliance itself:
  - Power supplies – Each Sun ZFS Storage 7320 storage controller and Sun Disk Shelf contains two power supplies. Each Sun ZFS Storage 7320 storage controller and the Sun Disk Shelf can continue to operate normally if one of those power supplies fails.
  - Disk drives – Each Sun ZFS Storage 7320 storage controller contains two mirrored boot drives, so the controller can still boot up and operate normally if one boot drive fails. The Sun Disk Shelf contains 20 hard disk drives that are used for storage in the SPARC SuperCluster T4-4, and 4 solid-state drives that are used as write-optimized cache devices, also known as logzillas. The SPARC SuperCluster T4-4 software provides redundancy between the disk drives. For more information, see "Understanding the SPARC SuperCluster T4-4 Configurations" on page 34.

Sun Datacenter InfiniBand Switch 36 Switches

Each SPARC SuperCluster T4-4 contains three Sun Datacenter InfiniBand Switch 36 switches, in either the full rack or half rack version, two of which are leaf switches (the third is used as a spine switch to connect two racks together). The two leaf switches are connected to each other to provide redundancy should one of the two leaf switches fail. In addition, each SPARC T4-4 server, Exadata Storage Server, and Sun ZFS Storage 7320 storage controller has connections to both leaf switches to provide redundancy in the InfiniBand connections should one of the two leaf switches fail. For more information, see "Understanding the Physical Connections" on page 24.

Cisco Catalyst 4948 Ethernet Management Switch

The Cisco Catalyst 4948 Ethernet management switch contains two power supplies. The Cisco Catalyst 4948 Ethernet management switch can continue to operate normally if one of those power supplies fails.

Power Distribution Units

Each SPARC SuperCluster T4-4 contains two power distribution units, in either the full rack or half rack version. The components within the SPARC SuperCluster T4-4 connect to both power distribution units, so that power continues to be supplied to those components should one of the two power distribution units fail. For more information, see “Power Distribution Units Physical Connections” on page 33.

Understanding the Physical Connections

The following topics describe the physical connections between the components within the SPARC SuperCluster T4-4:

SPARC T4-4 Server Physical Connections

These topics provide information on the location of the cards and ports that are used for the physical connections for the SPARC T4-4 server, as well as information specific to the four sets of physical connections for the server:

Card and Port Locations (SPARC T4-4 Servers)

The following figure shows the cards and ports that will be used for the physical connections for the SPARC T4-4 servers.

FIGURE 3 Card and Port Locations on the SPARC T4-4 Server

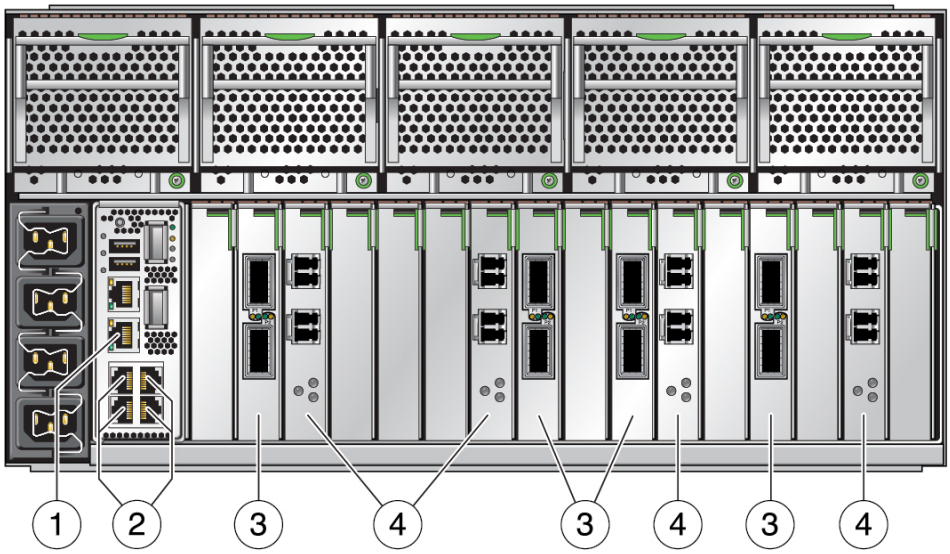

Figure Legend

1. NET MGT port, for connection to Oracle ILOM management network (see "Oracle ILOM Management Network Physical Connections (SPARC T4-4 Servers)" on page 27)
2. NET0-NET3 ports, for connection to 1 GbE host management network (see "1 GbE Host Management Network Physical Connections (SPARC T4-4 Servers)" on page 27)
3. Dual-port InfiniBand host channel adapters, for connection to the InfiniBand network (see "InfiniBand Private Network Physical Connections (SPARC T4-4 Servers)" on page 25)
4. Dual-port 10 GbE network interface cards, for connection to the 10 GbE client access network (see "10 GbE Client Access Network Physical Connections (SPARC T4-4 Servers)" on page 27)

InfiniBand Private Network Physical Connections (SPARC T4-4 Servers)

Each SPARC T4-4 server contains four dual-ported Sun QDR InfiniBand PCIe Low Profile host channel adapters (HCAs). The InfiniBand HCAs are installed in the following express module ports in every SPARC T4-4 server:

- Express module slot 1

- Express module slot 7
- Express module slot 9
- Express module slot 12

See "Card and Port Locations (SPARC T4-4 Servers)" on page 25 for more information on the location of the InfiniBand HCAs.

The two ports in each InfiniBand HCA (ports 1 and 2) connect to a different leaf switch to provide redundancy between the SPARC T4-4 servers and the leaf switches. The following figure shows how redundancy is achieved with the InfiniBand connections between the SPARC T4-4 servers and the leaf switches in a half rack configuration.

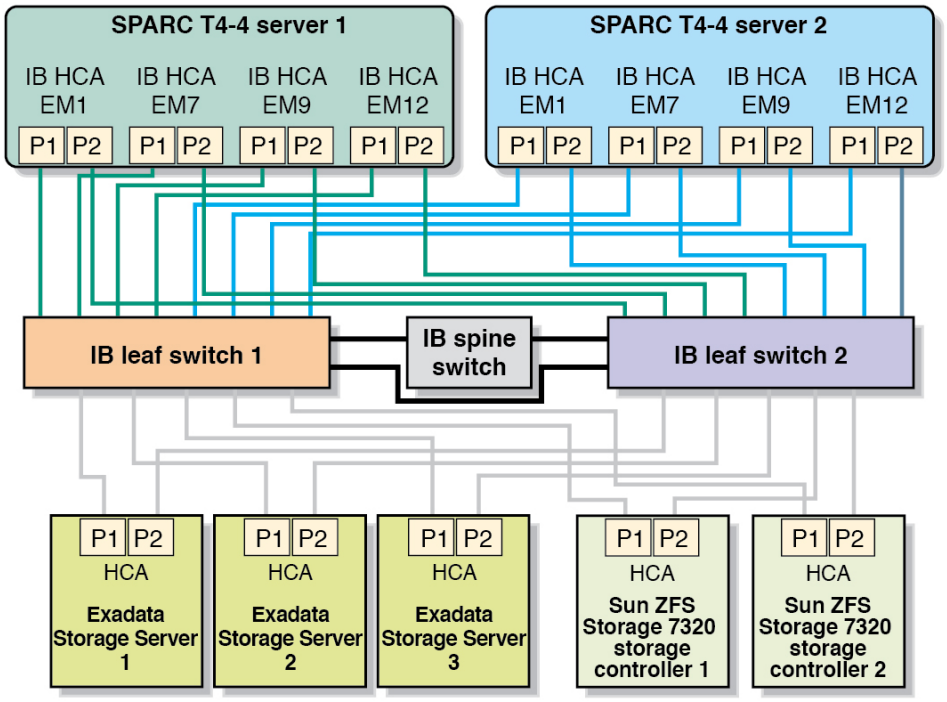

Note that only the physical connections for the InfiniBand private network are described in this section. Once the logical domains are created for each SPARC T4-4 server, the InfiniBand private network will be configured differently depending on the type of domain created for the SPARC T4-4 servers. The number of IP addresses needed for the InfiniBand network will also vary, depending on the type of domains created on each SPARC T4-4 server. For more information, see "Understanding the SPARC SuperCluster T4-4 Configurations" on page 34.

Oracle ILOM Management Network Physical Connections (SPARC T4-4 Servers)

Each SPARC T4-4 server connects to the Oracle Integrated Lights Out Manager (ILOM) management network through a single Oracle ILOM network port (NET MGT port) at the rear of each SPARC T4-4 server. One IP address is required for Oracle ILOM management for each SPARC T4-4 server.

See "Card and Port Locations (SPARC T4-4 Servers)" on page 25 for more information on the location of the NET MGT port.

1 GbE Host Management Network Physical Connections (SPARC T4-4 Servers)

Each SPARC T4-4 server connects to the 1 GbE host management network through the four 1 GbE host management ports at the rear of each SPARC T4-4 server (NET0 - NET3 ports). However, the way the 1 GbE host management connections are used differs from the physical connections due to logical domains. For more information, see "Understanding the SPARC SuperCluster T4-4 Configurations" on page 34.

See "Card and Port Locations (SPARC T4-4 Servers)" on page 25 for more information on the location of the 1 GbE host management ports.

10 GbE Client Access Network Physical Connections (SPARC T4-4 Servers)

Each SPARC T4-4 server contains four dual-ported Sun Dual 10 GbE SFP+ PCIe 2.0 Low Profile network interface cards (NICs). The 10 GbE NICs are installed in the following express module ports in every SPARC T4-4 server.

- Express module slot 2
- Express module slot 6
- Express module slot 10
- Express module slot 14

See "Card and Port Locations (SPARC T4-4 Servers)" on page 25 for the location of the 10 GbE NICs.

Depending on the configuration, one or two of the ports on the 10 GbE NICs (ports 0 and 1) will be connected to the client access network. In some configurations, both ports on the same 10 GbE NIC will be part of an IPMP group to provide redundancy and increased bandwidth. In other configurations, one port from two separate 10 GbE NICs will be part of an IPMP group.

The number of physical connections to the 10 GbE client access network varies, depending on the type of domains created on each SPARC T4-4 server. For more information, see "Understanding the SPARC SuperCluster T4-4 Configurations" on page 34.

Exadata Storage Server Physical Connections

Each Exadata Storage Server contains three sets of physical connections:

InfiniBand Private Network Physical Connections (Exadata Storage Servers)

Each Exadata Storage Server contains one dual-ported Sun QDR InfiniBand PCIe Low Profile host channel adapter (HCA). The two ports in the InfiniBand HCA are bonded together to increase available bandwidth. When bonded, the two ports appear as a single port, with a single IP address assigned to the two bonded ports, resulting in one IP address for InfiniBand private network connections for each Exadata Storage Server.

The two ports in the InfiniBand HCA connects to a different leaf switch to provide redundancy between the Exadata Storage Servers and the leaf switches. The following figure shows how redundancy is achieved with the InfiniBand connections between the Exadata Storage Servers and the leaf switches in a half rack configuration.

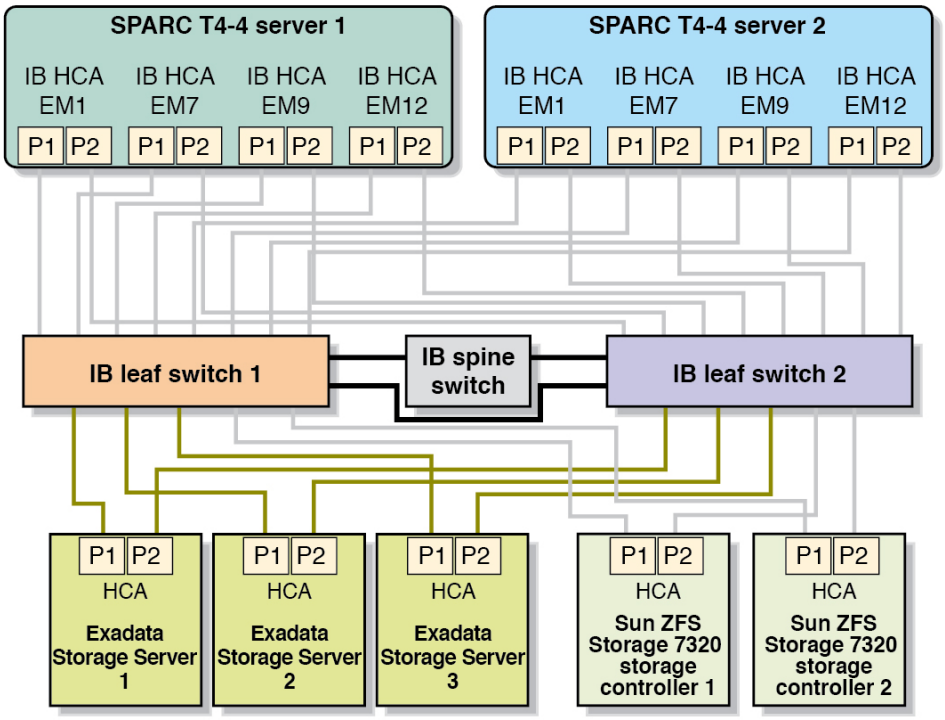

Oracle ILOM Management Network Physical Connections (Exadata Storage Servers)

Each Exadata Storage Server connects to the Oracle ILOM management network through a single Oracle ILOM network port (NET MGT port) at the rear of each Exadata Storage Server. One IP address is required for Oracle ILOM management for each Exadata Storage Server.

1 GbE Host Management Network Physical Connections (Exadata Storage Servers)

Each Exadata Storage Server connects to the 1 GbE host management network through the 1 GbE host management port (NET 0 port) at the rear of each Exadata Storage Server. One IP address is required for 1 GbE host management for each Exadata Storage Server.

Sun ZFS Storage 7320 Appliance Physical Connections

The Sun ZFS Storage 7320 appliance has five sets of physical connections:

InfiniBand Private Network Physical Connections (Sun ZFS Storage 7320 Appliance)

The Sun ZFS Storage 7320 appliance connects to the InfiniBand private network through one of the two Sun ZFS Storage 7320 storage controllers. The Sun ZFS Storage 7320 storage controller contains one Sun Dual Port 40GB InfiniBand QDR HCA. The two ports in each InfiniBand HCA are bonded together to increase available bandwidth. When bonded, the two ports appear as a single port, with a single IP address assigned to the two bonded ports, resulting in one IP address for InfiniBand private network connections for the Sun ZFS Storage 7320 storage controller.

The two ports in the InfiniBand HCA connect to a different leaf switch to provide redundancy between the Sun ZFS Storage 7320 storage controller and the leaf switches. The following figure shows how redundancy is achieved with the InfiniBand connections between the Sun ZFS Storage 7320 storage controller and the leaf switches in a half rack configuration.

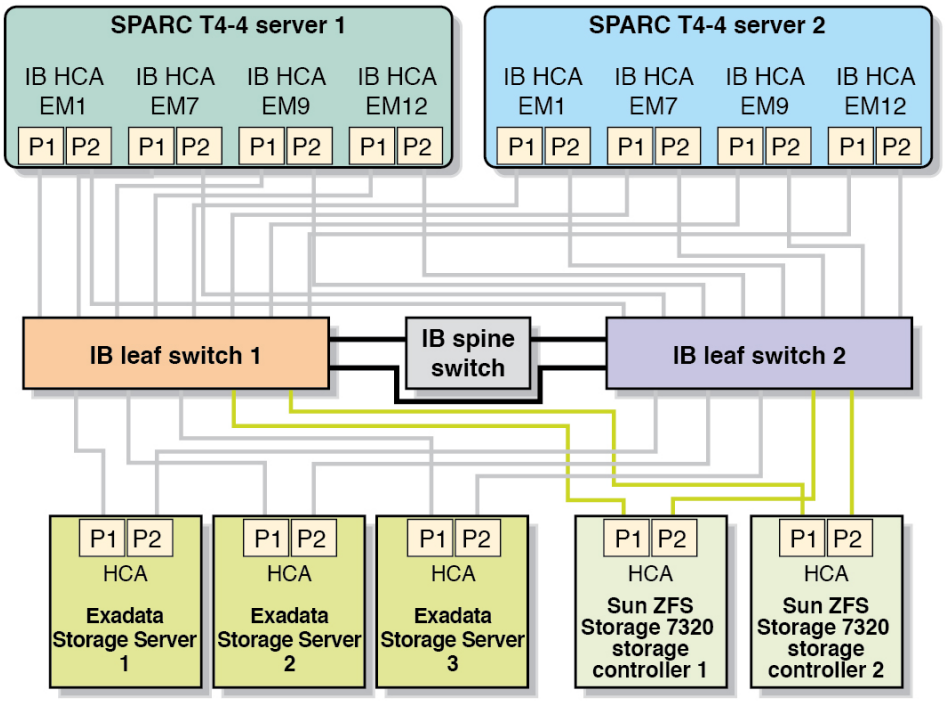

Oracle ILOM Management Network Physical Connections (Sun ZFS Storage 7320 Appliance)

The Sun ZFS Storage 7320 appliance connects to the Oracle ILOM management network through the two Sun ZFS Storage 7320 storage controllers. Each storage controller connects to the Oracle ILOM management network through the NET0 port at the rear of each storage controller using sideband management. One IP address is required for Oracle ILOM management for each storage controller.

1 GbE Host Management Network Physical Connections (Sun ZFS Storage 7320 Appliance)

The Sun ZFS Storage 7320 appliance connects to the 1 GbE host management network through the two Sun ZFS Storage 7320 storage controllers. The storage controllers connect to the 1 GbE host management network through the following ports at the rear of each storage controller:

- NET0 on the first storage controller (installed in slot 25 in the rack)
- NET1 on the second storage controller (installed in slot 26 in the rack)

One IP address is required for 1 GbE host management for each storage controller.

SAS Physical Connections (Sun ZFS Storage 7320 Appliance)

Each Sun ZFS Storage 7320 storage controller is populated with a dual-port SAS-2 HBA card. The Sun Disk Shelf also has two SIM Link In and two SIM Link Out ports. The two storage controllers connect to the Sun Disk Shelf in the following manner:

- **Storage controller 1** – Both ports from the SAS-2 HBA card to the SIM Link Out ports on the Sun Disk Shelf.
- **Storage controller 2** – Both ports from the SAS-2 HBA card to the SIM Link In ports on the Sun Disk Shelf.

The following figure show the SAS connections between the two storage controllers and the Sun Disk Shelf.

FIGURE 4 SAS Connections for the Sun ZFS Storage 7320 Appliance

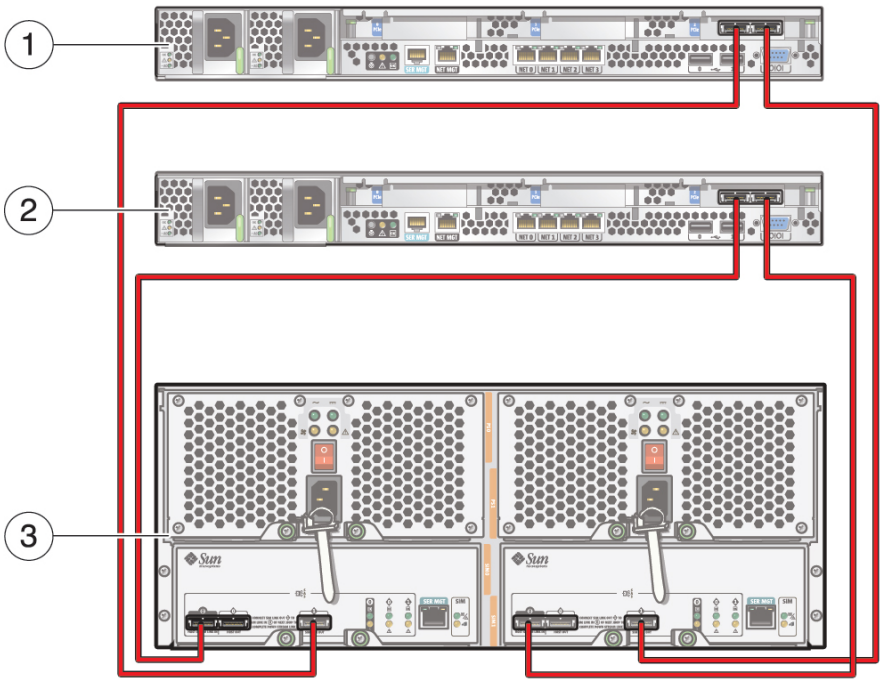

Figure Legend

1 Storage controller 1
2 Storage controller 2
3 Sun Disk Shelf

Cluster Physical Connections (Sun ZFS Storage 7320 Appliance)

Each Sun ZFS Storage 7320 storage controller contains a single cluster card. The cluster cards in the storage controllers are cabled together as shown in the following figure. This allows a heartbeat signal to pass between the storage controllers to determine if both storage controllers are up and running.

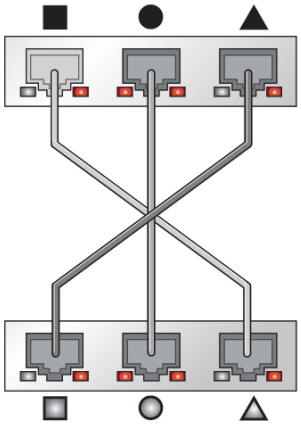

Power Distribution Units Physical Connections

The SPARC SuperCluster T4-4 rack contains two power distribution units. Each component in the SPARC SuperCluster T4-4 has redundant connections to the two power distribution units:

- **SPARC T4-4 servers** – Each SPARC T4-4 server has four AC power connectors. Two AC power connectors connect to one power distribution unit, and the other two AC power connectors connect to the other power distribution unit.
- **Exadata Storage Servers** – Each Exadata Storage Server has two AC power connectors. One AC power connector connects to one power distribution unit, and the other AC power connector connects to the other power distribution unit.
- **Sun ZFS Storage 7320 storage controllers** – Each Sun ZFS Storage 7320 storage controller has two AC power connectors. One AC power connector connects to one power distribution unit, and the other AC power connector connects to the other power distribution unit.
- **Sun Disk Shelf** – The Sun Disk Shelf has two AC power connectors. One AC power connector connects to one power distribution unit, and the other AC power connector connects to the other power distribution unit.
- **Sun Datacenter InfiniBand Switch 36 switches** – Each Sun Datacenter InfiniBand Switch 36 switch has two AC power connectors. One AC power connector connects to one power

distribution unit, and the other AC power connector connects to the other power distribution unit.

- **Cisco Catalyst 4948 Ethernet management switch** – The Cisco Catalyst 4948 Ethernet management switch has two AC power connectors. One AC power connector connects to one power distribution unit, and the other AC power connector connects to the other power distribution unit.

Understanding the SPARC SuperCluster T4-4 Configurations

The SPARC SuperCluster T4-4 is set up with logical domains (LDoms), which provide users with the flexibility to create different specialized virtual systems within a single hardware platform.

Each SPARC T4-4 server is configured with at least one logical domain. When domains are created on the SPARC SuperCluster T4-4, the following configurations are set for each domain:

- Percentage of CPU resources allocated to the domain – Set between 25% and 100% at installation for each domain, depending on the type of domain configuration. It may be possible to change the percentage of CPU resources allocated to the domain after installation, depending on the type of domain configuration. However, you will have to reboot the SPARC T4-4 server after you have reallocated the CPU resources for the changes to take affect.
- Percentage of memory resources allocated to the domain – Set between 25% and 100% at installation for each domain, depending on the type of domain configuration. It may be possible to change the percentage of memory resources allocated to the domain after installation, depending on the type of domain configuration. However, you will have to reboot the SPARC T4-4 server after you have reallocated the memory resources for the changes to take affect.
- Number of PCI root complexes allocated to the domain – Set from one to four PCI root complexes allocated to each domain, depending on the type of domain configuration. The number of PCI root complexes allocated to each domain cannot be changed after the initial installation.

Note - For more information on the PCI root complexes, see "SPARC T4-4 Servers" on page 20.

LDom configurations supported on the SPARC SuperCluster T4-4 have the following characteristics:

- From one to four LDoms on each SPARC T4-4 server in the SPARC SuperCluster T4-4
- Each LDom can be one of the following three types:

- Database Domain
- Application Domain running Oracle Solaris 10
- Application Domain running Oracle Solaris 11

In order to fully understand the different configuration options that are available for the SPARC SuperCluster T4-4, you must first understand the basic concepts for the express module slots and the different networks that are used for the system.

- "Understanding Logical Domains and the Express Module Slots" on page 35
- "Understanding the Management Network" on page 35
- "Understanding the 10 GbE Client Access Network" on page 36
- "Understanding the InfiniBand Network" on page 36

Understanding Logical Domains and the Express Module Slots

Each SPARC T4-4 server in the SPARC SuperCluster T4-4 has sixteen express module slots. The following cards are installed in certain express module slots and are used to connect to these networks:

- 10 GbE network interface cards (NICs) – Used to connect to the 10 GbE client access network
- InfiniBand host channel adapters (HCAs) – Used to connect to the private InfiniBand network

The express module slots used for each configuration varies, depending on the type and number of logical domains that are used for that configuration.

Understanding the Management Network

The management network connects to your existing management network, and is used for administrative work. Each SPARC T4-4 server provides access to the following management networks:

- Oracle Integrated Lights Out Manager (ILOM) management network – Connected through the Oracle ILOM Ethernet interface on each SPARC T4-4 server. Connections to this network are the same, regardless of the type of configuration that is set up on the SPARC T4-4 server.
- 1 GbE host management network – Connected through the four 1 GbE host management interfaces (NET0 - NET3) on each SPARC T4-4 server. Connections to this network will vary, depending on the type of configuration that is set up on the system. In most cases, the four 1 GbE host management ports at the rear of the SPARC T4-4 servers use IP network

multipathing (IPMP) to provide redundancy for the management network interfaces to the logical domains. However, the ports that are grouped together, and whether IPMP is used, varies depending on the type of configuration that is set up on the SPARC T4-4 server.

Understanding the 10 GbE Client Access Network

This required 10 GbE network connects the SPARC T4-4 servers to your existing client network and is used for client access to the servers. 10 GbE NICs installed in the express module slots are used for connection to this network. The number of 10 GbE NICs and the express module slots that they are installed in varies depending on the type of configuration that is set up on the SPARC T4-4 server.

Understanding the InfiniBand Network

The InfiniBand network connects the SPARC T4-4 servers, Sun ZFS Storage 7320 appliance, and Exadata Storage Servers using the InfiniBand switches on the rack. This non-routable network is fully contained in SPARC SuperCluster T4-4, and does not connect to your existing network.

When the SPARC SuperCluster T4-4 is configured with the appropriate types of domains, the InfiniBand network is partitioned to define the data paths between the SPARC T4-4 servers, and between the SPARC T4-4 servers and the storage appliances.

The defined InfiniBand data path coming out of the SPARC T4-4 servers varies, depending on the type of domain created on each SPARC T4-4 server:

- “InfiniBand Network Data Paths for a Database Domain” on page 36
- “InfiniBand Network Data Paths for an Application Domain” on page 37

InfiniBand Network Data Paths for a Database Domain

When a Database Domain is created on a SPARC T4-4 server, the Database Domain has the following InfiniBand paths:

- SPARC T4-4 server to both Sun Datacenter InfiniBand Switch 36 leaf switches
- SPARC T4-4 server to each Exadata Storage Server, through the Sun Datacenter InfiniBand Switch 36 leaf switches
- SPARC T4-4 server to the Sun ZFS Storage 7320 appliance, through the Sun Datacenter InfiniBand Switch 36 leaf switches

Each SPARC T4-4 server contains four InfiniBand HCAs. The number of InfiniBand HCAs that are assigned to the Database Domain varies, depending on the type of configuration that is set up on the SPARC T4-4 server.

For the InfiniBand HCAs assigned to a Database Domain, the following InfiniBand private networks are used:

- One InfiniBand private network for the Database Domains to communicate with each other and with the Application Domains running Oracle Solaris 10, and with the Sun ZFS Storage 7320 appliance
- One InfiniBand private network for the Oracle Database 11g Real Application Clusters (Oracle RAC) interconnects, and for communication between the Database Domains and the Exadata Storage Servers

The two ports on each InfiniBand HCA will connect to different Sun Datacenter InfiniBand Switch 36 leaf switches to provide redundancy between the SPARC T4-4 servers and the leaf switches. For more information on the physical connections for the SPARC T4-4 servers to the leaf switches, see "InfiniBand Private Network Physical Connections (SPARC T4-4 Servers)" on page 25.

InfiniBand Network Data Paths for an Application Domain

When an Application Domain is created on a SPARC T4-4 server (an Application Domain running either Oracle Solaris 10 or Oracle Solaris 11), the Application Domain will have the following InfiniBand paths:

- SPARC T4-4 server to both Sun Datacenter InfiniBand Switch 36 leaf switches
- SPARC T4-4 server to the Sun ZFS Storage 7320 appliance, through the Sun Datacenter InfiniBand Switch 36 leaf switches

Note that the Application Domain would not access the Exadata Storage Servers, which are used only for the Database Domain.

Every SPARC T4-4 server contains four InfiniBand HCAs. The number of InfiniBand HCAs that are assigned to the Application Domain varies, depending on the type of configuration that is set up on the SPARC T4-4 server.

For the InfiniBand HCAs assigned to an Application Domain, the following InfiniBand private networks are used:

- One InfiniBand private network for Application Domains to communicate with each other and with the Database Domains, and with the Sun ZFS Storage 7320 appliance
- Two InfiniBand private networks for the optional Oracle Solaris Cluster interconnects

The two ports on each InfiniBand HCA will connect to different Sun Datacenter InfiniBand Switch 36 leaf switches to provide redundancy between the SPARC T4-4 servers and the leaf switches. For more information on the physical connections for the SPARC T4-4 servers to the leaf switches, see "InfiniBand Private Network Physical Connections (SPARC T4-4 Servers)" on page 25.

Understanding Clustering Software

Clustering software is typically used on multiple interconnected servers so that they appear as if they are one server to end users and applications. For the SPARC SuperCluster T4-4, clustering software is used to cluster certain logical domains on the SPARC T4-4 servers together with the same type of domain on other SPARC T4-4 servers. The benefits of clustering software include the following:

- Reduce or eliminate system downtime because of software or hardware failure
- Ensure availability of data and applications to end users, regardless of the kind of failure that would normally take down a single-server system
- Increase application throughput by enabling services to scale to additional processors by adding nodes to the cluster and balancing the load
- Provide enhanced availability of the system by enabling you to perform maintenance without shutting down the entire cluster

The SPARC SuperCluster T4-4 uses the following clustering software:

- "Cluster Software for the Database Domain" on page 38
- "Cluster Software for the Oracle Solaris Application Domains" on page 39

Cluster Software for the Database Domain

Oracle Database 11g Real Application Clusters (Oracle RAC) enables the clustering of the Oracle Database on the Database Domain. Oracle RAC uses Oracle Clusterware for the infrastructure to cluster the Database Domain on the SPARC T4-4 servers together.

Oracle Clusterware is a portable cluster management solution that is integrated with the Oracle database. The Oracle Clusterware is also a required component for using Oracle RAC. The Oracle Clusterware enables you to create a clustered pool of storage to be used by any combination of single-instance and Oracle RAC databases.

Single-instance Oracle databases have a one-to-one relationship between the Oracle database and the instance. Oracle RAC environments, however, have a one-to-many relationship between the database and instances. In Oracle RAC environments, the cluster database instances access one database. The combined processing power of the multiple servers can provide greater throughput and scalability than is available from a single server. Oracle RAC is the Oracle Database option that provides a single system image for multiple servers to access one Oracle database.

Oracle RAC is a unique technology that provides high availability and scalability for all application types. The Oracle RAC infrastructure is also a key component for implementing the Oracle enterprise grid computing architecture. Having multiple instances access a single database prevents the server from being a single point of failure. Applications that you deploy on Oracle RAC databases can operate without code changes.

Cluster Software for the Oracle Solaris Application Domains

The Oracle Solaris Cluster software is an optional clustering tool used for the Oracle Solaris Application Domains. On the SPARC SuperCluster T4-4, the Oracle Solaris Cluster software is used to cluster the Oracle Solaris Application Domain on the SPARC T4-4 servers together.

Example SPARC SuperCluster T4-4 Configuration

This example shows how domains might be configured on a SPARC SuperCluster T4-4 and how the physical components and resources are allocated based on those domain configurations.

This example uses the following configurations:

- A SPARC SuperCluster T4-4 half rack containing the following components:
  - Two SPARC T4-4 servers
  - Three Exadata Storage Servers
  - One Sun ZFS Storage 7320 appliance, consisting of two Sun ZFS Storage 7320 storage controllers and one Sun Disk Shelf
  - Three Sun Datacenter InfiniBand Switch 36 switches
  - One 48-port Cisco Catalyst 4948 Ethernet management switch
- The following domain configurations on each SPARC T4-4 server:
  - One Database Domain
  - One Application Domain running Oracle Solaris 10
- 50% of the CPU resources allocated to each domain on each SPARC T4-4 server
- 50% of the memory resources allocated to each domain on each SPARC T4-4 server
- Two PCI root complexes allocated to each domain on each SPARC T4-4 server

The following sections describe how the physical components and resources are allocated on each of the domains based on their configurations:

- "Database Domain" on page 39
- "Application Domain Running Oracle Solaris 10" on page 43

Database Domain

The following hardware and resource allocations would be set for the Database Domain in both SPARC T4-4 servers in this example scenario.

CPU, Memory, and PCI Root Complex Resource Allocation

These resources are allocated to the Database Domain in both SPARC T4-4 servers:

- CPU resources – All the cores from processors 0 and 1 in processor module 0
- Memory resources – All the memory associated with the two lgroups associated with processors 0 and 1 in processor module 0
- PCI root complexes – PCI root complexes pci@400 and pci@500 in processor module 0

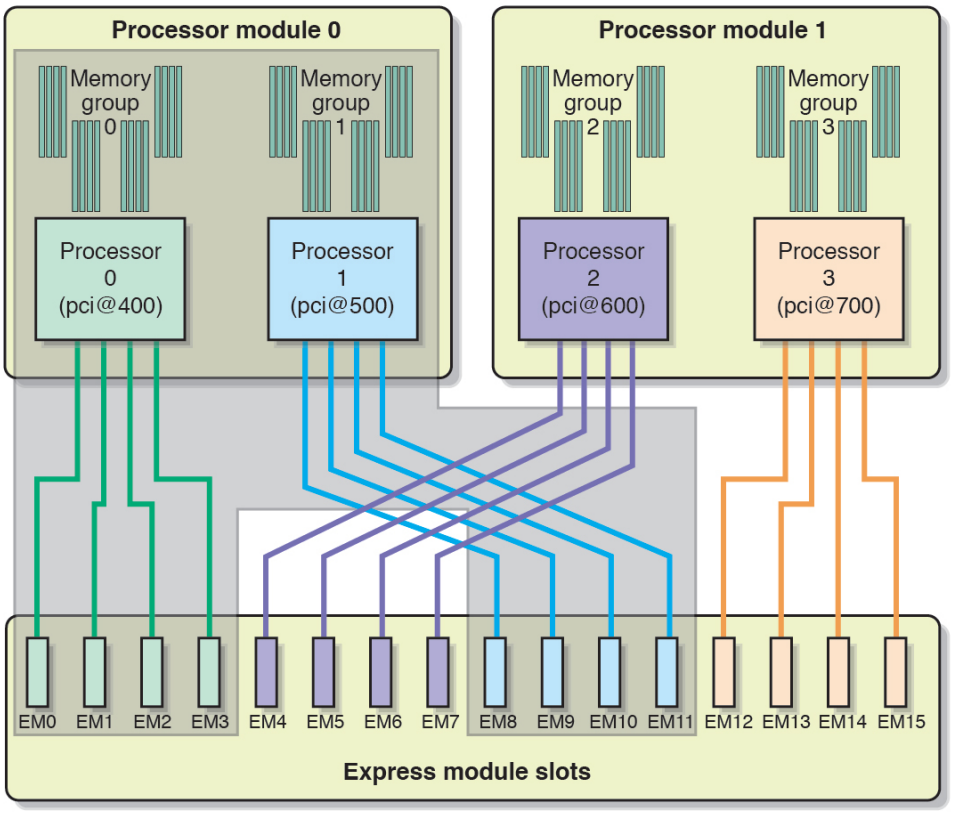

InfiniBand Connections

The following InfiniBand connections are set up for the Database Domain in both SPARC T4-4 servers in this example scenario.

At the SPARC T4-4 Server Level

Two of the four dual-ported Sun QDR InfiniBand PCIe Low Profile HCAs in the SPARC T4-4 server are assigned to the Database Domain, installed in these slots in the SPARC T4-4 server:

- Express module slot 1
- Express module slot 9

Both ports in each InfiniBand HCA are bonded together, resulting in two IPoIB interfaces for the Database Domain (one IPoIB interface for each InfiniBand HCA). The two ports in each InfiniBand HCA connect to a different leaf switch to provide redundancy between the SPARC T4-4 servers and the leaf switches. See “InfiniBand Private Network Physical Connections (Exadata Storage Servers)” on page 28 for more information.

At the SPARC SuperCluster T4-4 Level

The Database Domains in both SPARC T4-4 servers have InfiniBand connections to the InfiniBand leaf switches, Exadata Storage Servers, and the shared storage on the Sun ZFS Storage 7320 appliance.

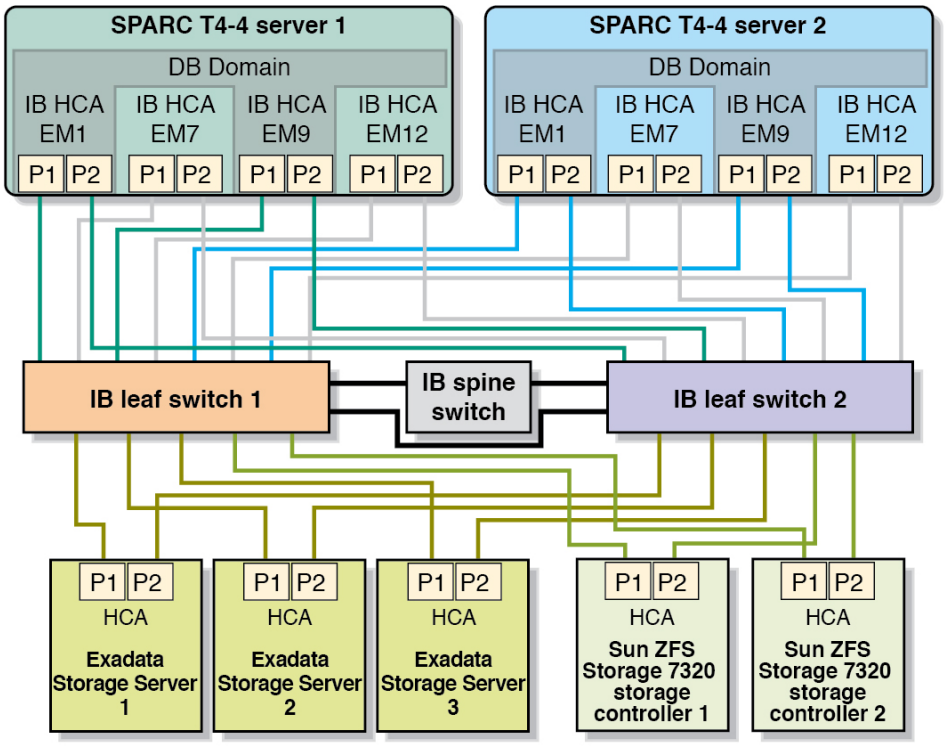

10 GbE Client Access Connections

Two of the four dual-ported Sun Dual 10 GbE SFP+ PCIe 2.0 Low Profile NICs in the SPARC T4-4 server are assigned to the Database Domain, installed in these slots in the SPARC T4-4 server:

- Express module slot 2
- Express module slot 10

One port from each dual-ported 10 GbE NIC is part of one IPMP group to provide redundancy. The IPMP grouping will be set during the initial configuration when the SPARC SuperCluster T4-4 is first installed at the customer site. In this example, the following ports are part of one IPMP group and are used for the physical connection to the 10 GbE network for the Database Domain:

- 10 GbE NIC in express module slot 2 – Port 0
- 10 GbE NIC in express module slot 10 – Port 1

The following figure gives a graphical representation of the 10 GbE client access connections for the Database Domains on the two SPARC T4-4 servers at the SPARC SuperCluster T4-4 level.

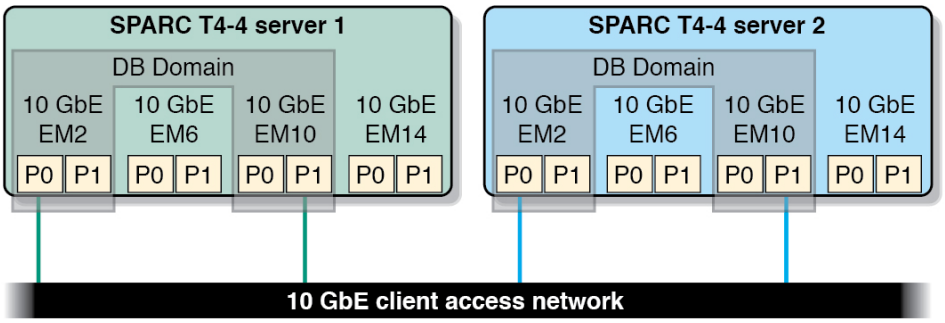

Application Domain Running Oracle Solaris 10

The following hardware and resource allocations are set for the Application Domain running Oracle Solaris 10 in both SPARC T4-4 servers in this example scenario.

- "CPU, Memory, and PCI Root Complex Resource Allocation" on page 43
- "InfiniBand Connections" on page 44
- "10 GbE Client Access Connections" on page 46

CPU, Memory, and PCI Root Complex Resource Allocation

The following resources are allocated to the Application Domain running Oracle Solaris 10 in both SPARC T4-4 servers in this example scenario:

- CPU resources – All the cores from processors 2 and 3 in processor module 1
- Memory resources – All the memory associated with the two lgroups associated with processors 2 and 3 in processor module 1
- PCI root complexes – PCI root complexes pci@600 and pci@700 in processor module 1

The following figure gives a graphical representation of how the CPU, memory, and PCI root complexes are allocated to the Application Domain running Oracle Solaris 10 in this example scenario.

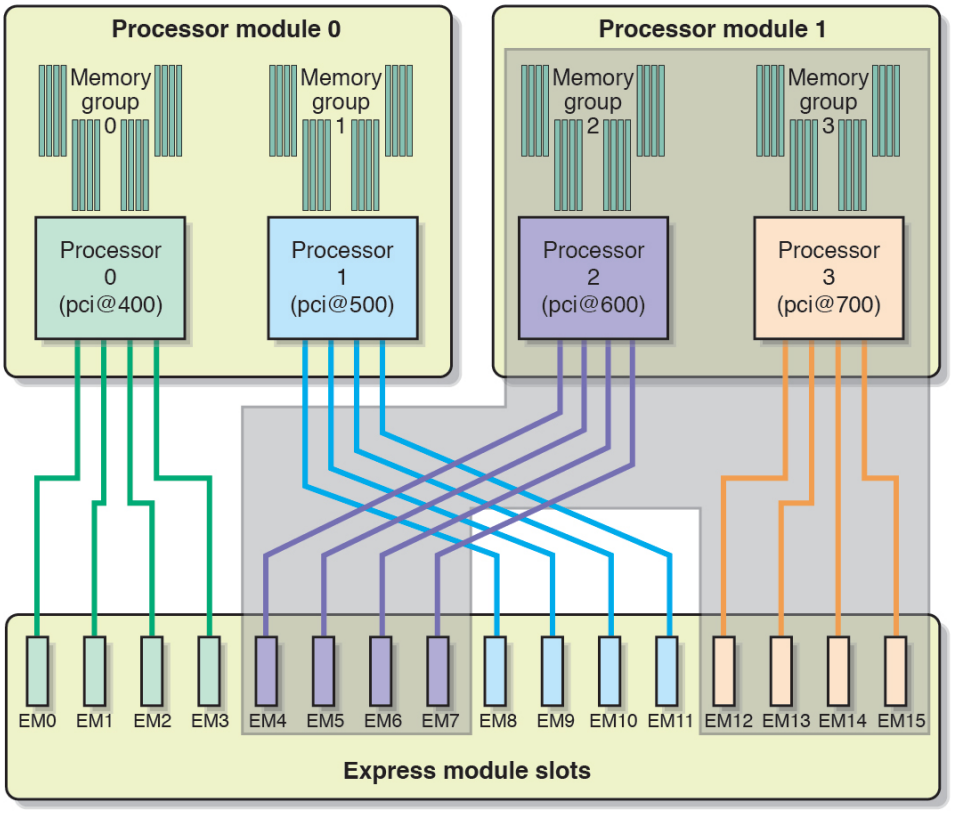

InfiniBand Connections

The following InfiniBand connections are set up for the Application Domain running Oracle Solaris 10 in both SPARC T4-4 servers in this example scenario.

At the SPARC T4-4 Server Level

Two of the four dual-ported Sun QDR InfiniBand PCIe Low Profile HCAs in the SPARC T4-4 server are assigned to the Application Domain running Oracle Solaris 10, installed in these slots in the SPARC T4-4 server:

- Express module slot 7
- Express module slot 12

However, the InfiniBand network setup for the Application Domain running Oracle Solaris 10 is set up differently from the Database Domain:

- One port from each InfiniBand HCA are bonded together, resulting in one IPoIB interface for the Application Domain running Oracle Solaris 10
- One port from each InfiniBand HCA are used for the optional Oracle Solaris Cluster interconnects

In this example, the following port are used for InfiniBand connections for the Application Domain running Oracle Solaris 10:

- InfiniBand HCA installed in express module slot 7 – Port 1
- InfiniBand HCA installed in express module slot 12 – Port 2

At the SPARC SuperCluster T4-4 Level

The Application Domains running Oracle Solaris 10 in both SPARC T4-4 servers have InfiniBand connections to the InfiniBand leaf switches and the shared storage on the Sun ZFS Storage 7320 appliance. However, the Application Domain running Oracle Solaris 10 will differ from the Database Domain in that it would not access the Exadata Storage Servers, which are used only for the Database Domain.

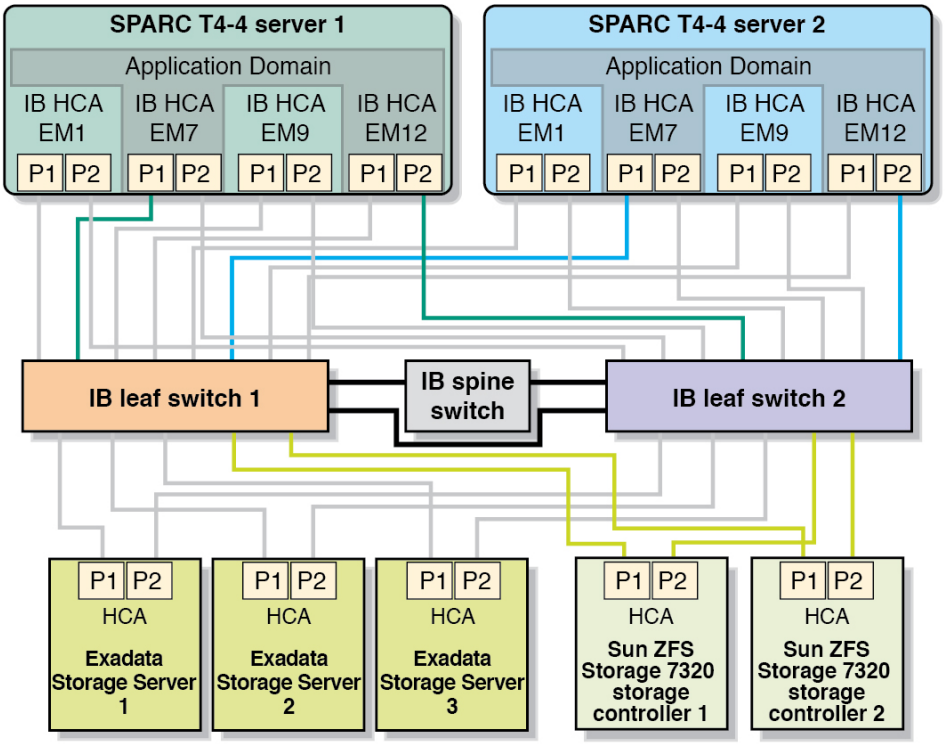

10 GbE Client Access Connections

Two of the four dual-ported Sun Dual 10 GbE SFP+ PCIe 2.0 Low Profile NICs in the SPARC T4-4 server are assigned to the Application Domain running Oracle Solaris 10, installed in these slots in the SPARC T4-4 server:

- Express module slot 6
- Express module slot 14

One port from each dual-ported 10 GbE NIC will be part of one IPMP group to provide redundancy. The IPMP grouping will be set during the initial configuration when the SPARC SuperCluster T4-4 is first installed at the customer site. In this example, the following ports are part of one IPMP group and are used for the physical connection to the 10 GbE network for the Application Domain running Oracle Solaris 10:

- 10 GbE NIC in express module slot 6 – Port 0
- 10 GbE NIC in express module slot 14 – Port 1

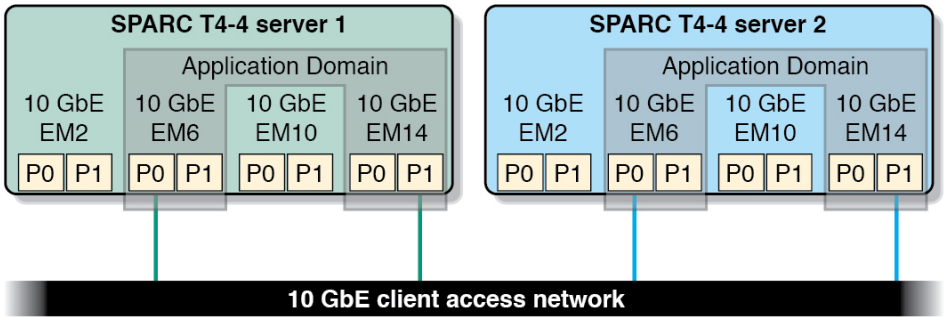

Understanding the Network Requirements

These topics describe the network requirements for the SPARC SuperCluster T4-4.

- "Network Requirements Overview" on page 46
- "Network Connection Requirements for SPARC SuperCluster T4-4" on page 50
- "Default IP Addresses" on page 51

Network Requirements Overview

The SPARC SuperCluster T4-4 includes SPARC T4-4 servers, Exadata Storage Servers, and the Sun ZFS Storage 7320 appliance, as well as equipment to connect the SPARC T4-4 servers

to your network. The network connections enable the servers to be administered remotely and enable clients to connect to the SPARC T4-4 servers.

Each SPARC T4-4 server consists of the following network components and interfaces:

- 4 embedded Gigabit Ethernet ports (NET0, NET1, NET2, and NET3) for connection to the host management network
- 1 Ethernet port (NET MGT) for Oracle Integrated Lights Out Manager (Oracle ILOM) remote management
- 4 dual-ported Sun QDR InfiniBand PCIe Low Profile host channel adapters (HCAs) for connection to the InfiniBand private network
- 4 dual-ported Sun Dual 10 GbE SFP+ PCIe 2.0 Low Profile network interface cards (NICs) for connection to the 10 GbE client access network

---

Note - The QSFP modules for the 10 GbE PCIe 2.0 network cards are purchased separately.

Each Exadata Storage Server consists of the following network components and interfaces:

- 1 embedded Gigabit Ethernet port (NET0) for connection to the host management network
- 1 dual-ported Sun QDR InfiniBand PCIe Low Profile host channel adapter (HCA) for connection to the InfiniBand private network
- 1 Ethernet port (NET MGT) for Oracle ILOM remote management

Each Sun ZFS Storage 7320 storage controller consists of the following network components and interfaces:

- 1 embedded Gigabit Ethernet port for connection to the host management network:
  - NET0 on the first storage controller (installed in slot 25 in the rack)
  - NET1 on the second storage controller (installed in slot 26 in the rack)
- 1 dual-port QDR InfiniBand Host Channel Adapter for connection to the InfiniBand private network
- 1 Ethernet port (NET0) for Oracle ILOM remote management using sideband management. The dedicate Oracle ILOM port is not used due to sideband.

The Cisco Catalyst 4948 Ethernet switch supplied with the SPARC SuperCluster T4-4 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, might 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 SPARC SuperCluster T4-4.

To deploy the SPARC SuperCluster T4-4, ensure that you meet the minimum network requirements. There are three networks for SPARC SuperCluster T4-4. 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 SPARC SuperCluster T4-4. It connects the servers, Oracle ILOM, and switches connected to the Ethernet switch in the rack. There is one uplink from the Ethernet switch in the rack to your existing management network.

Note - Network connectivity to the PDUs is only required if the electric current will be monitored remotely.

Each SPARC T4-4 server and Exadata Storage Server use two network interfaces for management. One provides management access to the operating system through the 1 GbE host management interface(s), and the other provides access to the Oracle Integrated Lights Out Manager through the Oracle ILOM Ethernet interface.

Note - The SPARC T4-4 servers have four 1 GbE host management interfaces (NET 0 - NET3). All four NET interfaces are physically connected, and use IPMP to provide redundancy. See “Understanding the SPARC SuperCluster T4-4 Configurations” on page 34 for more information.

The method used to connect the Sun ZFS Storage 7320 storage controllers to the management network varies depending on the controller:

- Storage controller 1: NET0 used to provide access to the Oracle ILOM network using sideband management, as well as access to the 1 GbE host management network.
- Storage controller 2: NET0 used to provide access to the Oracle ILOM network using sideband management, and NET1 used to provide access to the 1 GbE host management network.

The SPARC SuperCluster T4-4 is delivered with the 1 GbE host management and Oracle ILOM interfaces connected to the Ethernet switch on the rack. The 1 GbE host management

interfaces on the SPARC T4-4 servers should not be used for client or application network traffic. Cabling or configuration changes to these interfaces is not permitted.

- **Client access network –** This required 10 GbE network connects the SPARC T4-4 servers to your existing client network and is used for client access to the servers. Database applications access the database through this network using Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) addresses.
- **InfiniBand private network** – This network connects the SPARC T4-4 servers, Sun ZFS Storage 7320 appliance, and Exadata Storage Servers using the InfiniBand switches on the rack. For SPARC T4-4 servers configured with Database Domains, Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on Exadata Storage Servers and the Sun ZFS Storage 7320 appliance. For SPARC T4-4 servers configured with the Application Domain, Oracle Solaris Cluster uses this network for cluster interconnect traffic and to access data on the Sun ZFS Storage 7320 appliance. This non-routable network is fully contained in SPARC SuperCluster T4-4, and does not connect to your existing network. This network is automatically configured during installation.

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

The following figure shows the default network diagram.

FIGURE 5 Network Diagram for the SPARC SuperCluster T4-4

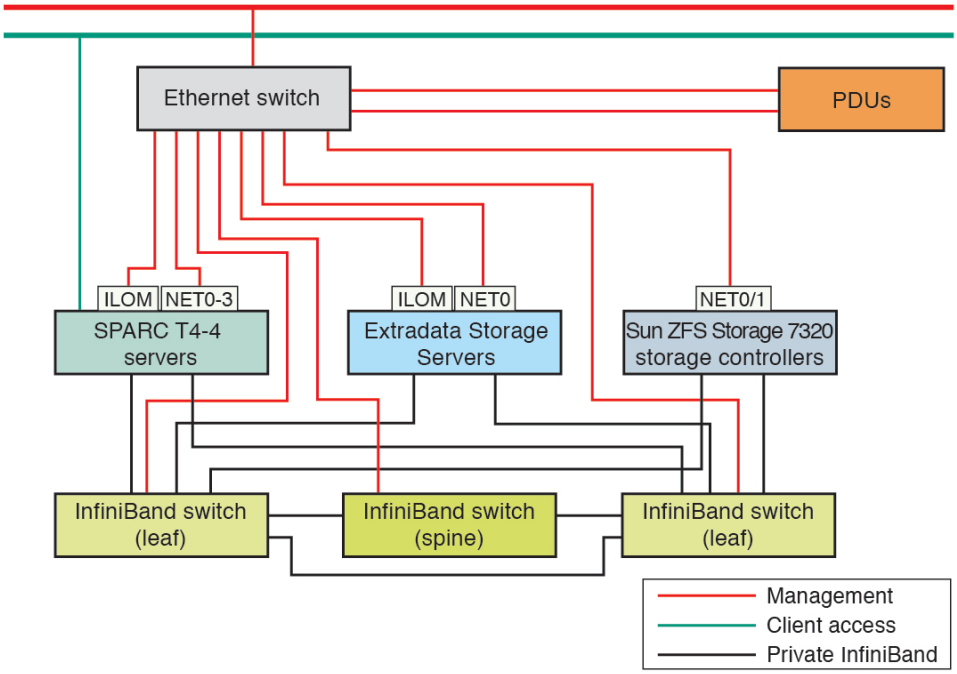

Network Connection Requirements for SPARC SuperCluster T4-4

The following connections are required for a SPARC SuperCluster T4-4 installation:

TABLE 1 New Network Connections Required for Installation

Connection Type	Number of connections	Comments
Management network	1 for Ethernet switch	Connect to the existing management network
Client access network	Typically 1–2 per SPARC T4-4 server: ■ Full rack: typically 4–8 connections ■ Half rack: typically 2–4 connections	Connect to the client access network. (You will not have redundancy through IPMP if there is only one connection per SPARC T4-4.)

Default IP Addresses (SPARC SuperCluster T4-4)

These topics list the default IP addresses assigned to SPARC SuperCluster T4-4 components during manufacturing.

Default IP Addresses

Four sets of default IP addresses are assigned at manufacturing:

- **Management IP addresses** – IP addresses used by Oracle ILOM for the SPARC T4-4 servers, Exadata Storage Servers, and the Sun ZFS Storage 7320 storage controllers.
- **Host IP addresses** – Host IP addresses used by the SPARC T4-4 servers, Exadata Storage Servers, and the Sun ZFS Storage 7320 storage controllers.
- **InfiniBand IP addresses** – InfiniBand interfaces are the default channel of communication among SPARC T4-4 servers, Exadata Storage Servers, and the Sun ZFS Storage 7320 storage controllers. If you are connecting a SPARC SuperCluster T4-4 to another SPARC SuperCluster T4-4 or to an Oracle Exadata or Exalogic machine on the same InfiniBand fabric, the InfiniBand interface enables communication between the SPARC T4-4 servers and storage server heads in one SPARC SuperCluster T4-4 and the other SPARC SuperCluster T4-4 or Oracle Exadata or Exalogic machine.
- **10 GbE IP addresses** – The IP addresses used by the 10 GbE client access network interfaces.

Tip - For more information about how these interfaces are used, see Figure 5.

Default Host Names and IP Addresses (Full Rack)

Refer to the following topics for the default IP addresses used in the SPARC SuperCluster T4-4 full rack:

Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Full Rack)

TABLE 2 Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Full Rack)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**Oracle ILOM Host Names**	**Oracle ILOM IP Addresses**	**Host Management Host Names**	**Host Management IP Addresses**
N/A	PDU-A (left from rear view)	sscpdua	192.168.1.210		N/A
	PCU-B (right from rear view)	sscpdub	192.168.1.211		N/A
42	Exadata Storage Server 6	ssces6-sp	192.168.1.106	cell06	192.168.1.6
41					
40	Exadata Storage Server 5	ssces5-sp	192.168.1.105	cell05	192.168.1.5
39					
38	Exadata Storage Server 4	ssces4-sp	192.168.1.104	cell04	192.168.1.4
37					
36	SPARC T4-4 Server 4	ssccn4-sp	192.168.1.111	ssccn4-m4	192.168.1.41
35				ssccn4-m3	192.168.1.31
34				ssccn4-m2	192.168.1.21
33				ssccn4-m1	192.168.1.11
32					
31	SPARC T4-4 Server 3	ssccn3-sp	192.168.1.110	ssccn3-m4	192.168.1.40
30				ssccn3-m3	192.168.1.30
29				ssccn3-m2	192.168.1.20
28				ssccn3-m1	192.168.1.10
27					
26	Sun ZFS Storage 7320 Storage Controller 2	sscsn2-sp	192.168.1.116	sscsn2-m1	192.168.1.16
25	Sun ZFS Storage 7320 Storage Controller 1	sscsn1-sp	192.168.1.115	sscsn1-m1	192.168.1.15
24	Sun Datacenter InfiniBand Switch 36 (Leaf 2)	sscnm1-m3	192.168.1.203	N/A	N/A
23	Sun Disk Shelf for the Sun ZFS Storage 7320 Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Cisco Catalyst 4948 Ethernet Management Switch	ssc4948-m1	192.168.1.200	N/A	N/A
18	Sun Datacenter InfiniBand Switch 36 (Leaf 1)	sscnm1-m2	192.168.1.202	N/A	N/A
17	SPARC T4-4 Server 2	ssccn2-sp	192.168.1.109	ssccn2-m4	192.168.1.49
16				ssccn2-m3	192.168.1.39

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**Oracle ILOM Host Names**	**Oracle ILOM IP Addresses**	**Host Management Host Names**	**Host Management IP Addresses**
15				ssccn2-m2	192.168.1.29
14				ssccn2-m1	192.168.1.9
13					
12	SPARC T4-4 Server 1	ssccn1-sp	192.168.1.108	ssccn1-m4	192.168.1.38
11				ssccn1-m3	192.168.1.28
10				ssccn1-m2	192.168.1.18
9				ssccn1-m1	192.168.1.8
8					
7	Exadata Storage Server 3	ssces3-sp	192.168.1.103	cell03	192.168.1.3
6					
5	Exadata Storage Server 2	ssces2-sp	192.168.1.102	cell02	192.168.1.2
4					
3	Exadata Storage Server 1	ssces1-sp	192.168.1.101	cell01	192.168.1.1
2					
1	Sun Datacenter InfiniBand Switch 36 (Spine)	sscnm1-m1	192.168.1.201	N/A	N/A

Default Host Names and IP Addresses for the InfiniBand and 10 GbE Client Access Networks (Full Rack)

TABLE 3 Default Host Names and IP Addresses for the InfiniBand and 10 GbE Client Access Networks (Full Rack)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**InfiniBand Host Names**	**InfiniBand IP Addresses**	**10 GbE Client Access IP Host names**	**10 GbE Client Access IP Addresses**
N/A	PDU-A (left from rear view)	N/A	N/A	N/A	N/A
	PCU-B (right from rear view)	N/A	N/A	N/A	N/A
42	Exadata Storage Server 6	ssces6-stor	192.168.10.6	N/A	N/A
41					
40	Exadata Storage Server 5	ssces5-stor	192.168.10.5	N/A	N/A
39					
38	Exadata Storage Server 4	ssces4-stor	192.168.10.4	N/A	N/A
37					
36	SPARC T4-4 Server 4	ssccn4-ib4	192.168.10.41	ssccn4-tg8 ssccn4-tg7	192.168.40.32 192.168.40.31
35		ssccn4-ib3	192.168.10.31	ssccn4-tg6 ssccn4-tg5	192.168.40.30 192.168.40.29

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**InfiniBand Host Names**	**InfiniBand IP Addresses**	**10 GbE Client Access IP Host names**	**10 GbE Client Access IP Addresses**
34		ssccn4-ib2	192.168.10.21	ssccn4-tg4 ssccn4-tg3	192.168.40.28 192.168.40.27
33		ssccn4-ib1	192.168.10.11	ssccn4-tg2	192.168.40.26
32				ssccn4-tg1	192.168.40.25
31	SPARC T4-4 Server 3	ssccn3-ib4	192.168.10.40	ssccn3-tg8 ssccn3-tg7	192.168.40.24 192.168.40.23
30		ssccn3-ib3	192.168.10.30	ssccn3-tg6 ssccn3-tg5	192.168.40.22 192.168.40.21
29		ssccn3-ib2	192.168.10.20	ssccn3-tg4 ssccn3-tg3	192.168.40.20 192.168.40.19
28		ssccn3-ib1	192.168.10.10	ssccn3-tg2	192.168.40.18
27				ssccn3-tg1	192.168.40.17
26	Sun ZFS Storage 7320 Storage Controller 2	sscsn1-stor2	N/A	N/A	N/A
25	Sun ZFS Storage 7320 Storage Controller 1	sscsn1-stor1	192.168.30.15	N/A	N/A
24	Sun Datacenter InfiniBand Switch 36 (Leaf 2)	N/A	N/A	N/A	N/A
23	Sun Disk Shelf for the Sun ZFS Storage 7320 Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Cisco Catalyst 4948 Ethernet Management Switch	N/A	N/A	N/A	N/A
18	Sun Datacenter InfiniBand Switch 36 (Leaf 1)	N/A	N/A	N/A	N/A
17	SPARC T4-4 Server 2	ssccn2-ib4	192.168.10.39	ssccn2-tg8 ssccn2-tg7	192.168.40.16 192.168.40.15
16		ssccn2-ib3	192.168.10.29	ssccn2-tg6 ssccn2-tg5	192.168.40.14 192.168.40.13
15		ssccn2-ib2	192.168.10.19	ssccn2-tg4 ssccn2-tg3	192.168.40.12 192.168.40.11
14		ssccn2-ib1	192.168.10.9	ssccn2-tg2	192.168.40.10
13				ssccn2-tg1	192.168.40.9
12	SPARC T4-4 Server 1	ssccn1-ib4	192.168.10.38	ssccn1-tg8	192.168.40.8

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**InfiniBand Host Names**	**InfiniBand IP Addresses**	**10 GbE Client Access IP Host names**	**10 GbE Client Access IP Addresses**
				ssccn1-tg7	192.168.40.7
11		ssccn1-ib3	192.168.10.28	ssccn1-tg6 ssccn1-tg5	192.168.40.6 192.168.40.5
10		ssccn1-ib2	192.168.10.18	ssccn1-tg4 ssccn1-tg3	192.168.40.4 192.168.40.3
9		ssccn1-ib1	192.168.10.8	ssccn1-tg2	192.168.40.2
8				ssccn1-tg1	192.168.40.1
7	Exadata Storage Server 3	ssces3-stor	192.168.10.3	N/A	N/A
6					
5	Exadata Storage Server 2	ssces2-stor	192.168.10.2	N/A	N/A
4					
3	Exadata Storage Server 1	ssces1-stor	192.168.10.1	N/A	N/A
2					
1	Sun Datacenter InfiniBand Switch 36 (Spine)	N/A	N/A	N/A	N/A

Default Host Names and IP Addresses (Half Rack)

Refer to the following topics for the default IP addresses used in the SPARC SuperCluster T4-4 half rack:

- “Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Half Rack)” on page 55
- “Default Host Names and IP Addresses for the InfiniBand and 10 GbE Client Access Networks (Half Rack)” on page 57

Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Half Rack)

TABLE 4 Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Half Rack)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**Oracle ILOM Host Names**	**Oracle ILOM IP Addresses**	**Host Management Host Names**	**Host Management IP Addresses**
N/A	PDU-A (left from rear view)	sscpdua	192.168.1.210		N/A
	PCU-B (right from rear view)	sscpdub	192.168.1.211		N/A
42	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
41					

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**Oracle ILOM Host Names**	**Oracle ILOM IP Addresses**	**Host Management Host Names**	**Host Management IP Addresses**
40					
39					
38	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
37					
36					
35					
34	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
33					
32					
31					
30	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
29					
28					
27					
26	Sun ZFS Storage 7320 Storage Controller 2	sscsn2-sp	192.168.1.116	sscsn2-m1	192.168.1.16
25	Sun ZFS Storage 7320 Storage Controller 1	sscsn1-sp	192.168.1.115	sscsn1-m1	192.168.1.15
24	Sun Datacenter InfiniBand Switch 36 (Leaf 2)	sscnm1-m3	192.168.1.203	N/A	N/A
23	Sun Disk Shelf for the Sun ZFS Storage 7320 Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Cisco Catalyst 4948 Ethernet Management Switch	ssc4948-m1	192.168.1.200	N/A	N/A
18	Sun Datacenter InfiniBand Switch 36 (Leaf 1)	sscnm1-m2	192.168.1.202	N/A	N/A
17	SPARC T4-4 Server 2	ssccn2-sp	192.168.1.109	ssccn2-m4	192.168.1.49
16				ssccn2-m3	192.168.1.39
15				ssccn2-m2	192.168.1.29
14				ssccn2-m1	192.168.1.9
13					
12	SPARC T4-4 Server 1	ssccn1-sp	192.168.1.108	ssccn1-m4	192.168.1.38
11				ssccn1-m3	192.168.1.28
10				ssccn1-m2	192.168.1.18
9				ssccn1-m1	192.168.1.8
8					

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**Oracle ILOM Host Names**	**Oracle ILOM IP Addresses**	**Host Management Host Names**	**Host Management IP Addresses**
7	Exadata Storage Server 3	ssces3-sp	192.168.1.103	cell03	192.168.1.3
6					
5	Exadata Storage Server 2	ssces2-sp	192.168.1.102	cell02	192.168.1.2
4					
3	Exadata Storage Server 1	ssces1-sp	192.168.1.101	cell01	192.168.1.1
2					
1	Sun Datacenter InfiniBand Switch 36 (Spine)	sscnm1-m1	192.168.1.201	N/A	N/A

Default Host Names and IP Addresses for the InfiniBand and 10 GbE Client Access Networks (Half Rack)

TABLE 5 Default Host Names and IP Addresses for the InfiniBand and 10 GbE Client Access Networks (Half Rack)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**InfiniBand Host Names**	**InfiniBand IP Addresses**	**10 GbE Client Access IP Host names**	**10 GbE Client Access IP Addresses**
N/A	PDU-A (left from rear view)	N/A	N/A	N/A	N/A
	PCU-B (right from rear view)	N/A	N/A	N/A	N/A
42	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
41					
40					
39					
38	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
37					
36					
35					
34	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
33					
32					
31					
30	4U Solid Filler Panel	Not applicable	Not applicable	Not applicable	Not applicable
29					
28					
27					
26	Sun ZFS Storage 7320 Storage Controller 2	sscsn1-stor2	N/A	N/A	N/A

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**InfiniBand Host Names**	**InfiniBand IP Addresses**	**10 GbE Client Access IP Host names**	**10 GbE Client Access IP Addresses**
25	Sun ZFS Storage 7320 Storage Controller 1	sscsn1-stor1	192.168.30.15	N/A	N/A
24	Sun Datacenter InfiniBand Switch 36 (Leaf 2)	N/A	N/A	N/A	N/A
23	Sun Disk Shelf for the Sun ZFS Storage 7320 Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Cisco Catalyst 4948 Ethernet Management Switch	N/A	N/A	N/A	N/A
18	Sun Datacenter InfiniBand Switch 36 (Leaf 1)	N/A	N/A	N/A	N/A
17	SPARC T4-4 Server 2	ssccn2-ib4	192.168.10.39	ssccn2-tg8 ssccn2-tg7	192.168.40.16 192.168.40.15
16		ssccn2-ib3	192.168.10.29	ssccn2-tg6 ssccn2-tg5	192.168.40.14 192.168.40.13
15		ssccn2-ib2	192.168.10.19	ssccn2-tg4 ssccn2-tg3	192.168.40.12 192.168.40.11
14		ssccn2-ib1	192.168.10.9	ssccn2-tg2 ssccn2-tg1	192.168.40.10 192.168.40.9
13					
12	SPARC T4-4 Server 1	ssccn1-ib4	192.168.10.38	ssccn1-tg8 ssccn1-tg7	192.168.40.8 192.168.40.7
11		ssccn1-ib3	192.168.10.28	ssccn1-tg6 ssccn1-tg5	192.168.40.6 192.168.40.5
10		ssccn1-ib2	192.168.10.18	ssccn1-tg4 ssccn1-tg3	192.168.40.4 192.168.40.3
9		ssccn1-ib1	192.168.10.8	ssccn1-tg2 ssccn1-tg1	192.168.40.2 192.168.40.1
8					
7	Exadata Storage Server 3	ssces3-stor	192.168.10.3	N/A	N/A
6					
5	Exadata Storage Server 2	ssces2-stor	192.168.10.2	N/A	N/A
4					
3	Exadata Storage Server 1	ssces1-stor	192.168.10.1	N/A	N/A
2					
1	Sun Datacenter InfiniBand Switch 36 (Spine)	N/A	N/A	N/A	N/A

Preparing the Site

This section describes the steps you should take to prepare the site for your system.

Cautions and Considerations

Consider the following when selecting a location for the new rack.

- Do not install the rack in a location that is exposed to:
  - Direct sunlight
  - Excessive dust
  - Corrosive gases
  - Air with high salt concentrations
  - Frequent vibrations
  - Sources of strong radio frequency interference
  - Static electricity
- Use power outlets that provide proper grounding.
  - A qualified electrical engineer must perform any grounding work.
  - Each grounding wire for the rack must be used only for the rack.
  - The grounding resistance must not be greater than 10 ohms.
  - Verify the grounding method for the building.
- Observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- (CHECKLIST) Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.
- (CHECKLIST) Clean and vacuum the area thoroughly in preparation for installation.

Reviewing System Specifications

Physical Specifications

Ensure that the installation site can properly accommodate the SPARC SuperCluster T4-4 by reviewing its physical specifications and space requirements.

Parameter	Metric	English
Height	1998 mm	78.66 in.
Width with side panels	600 mm	23.62 in.
Depth (with doors)	1200 mm	47.24 in.
Depth (without doors)	1112 mm	43.78 in.
Minimum ceiling height	2300 mm	90 in.
Minimum space between top of cabinet and ceiling	914 mm	36 in.
Weight (full rack)	862 kg	1,900 lbs
Weight (half rack)	594 kg	1,310 lbs

Related Information

- "Installation and Service Area" on page 60
- "Rack and Floor Cutout Dimensions" on page 61
- "Shipping Package Dimensions" on page 69

Installation and Service Area

Select an installation site that provides enough space to install and service the system.

Location	Maintenance Access	
Rear maintenance	914 mm	36 in.
Front maintenance	914 mm	36 in.
Top maintenance	914 mm	36 in.

Related Information

- “Physical Specifications” on page 60
- “Rack and Floor Cutout Dimensions” on page 61

Rack and Floor Cutout Dimensions

If you plan to route cables down through the bottom of the rack, cut a rectangular hole in the floor tile. Locate the hole below the rear portion of the rack, between the two rear casters and behind the rear inner rails. The suggested hole width is 280 mm (11 inches).

If you want a separate grounding cable, see “Install a Ground Cable (Optional)” on page 82.

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

FIGURE 6 Dimensions for Rack Stabilization

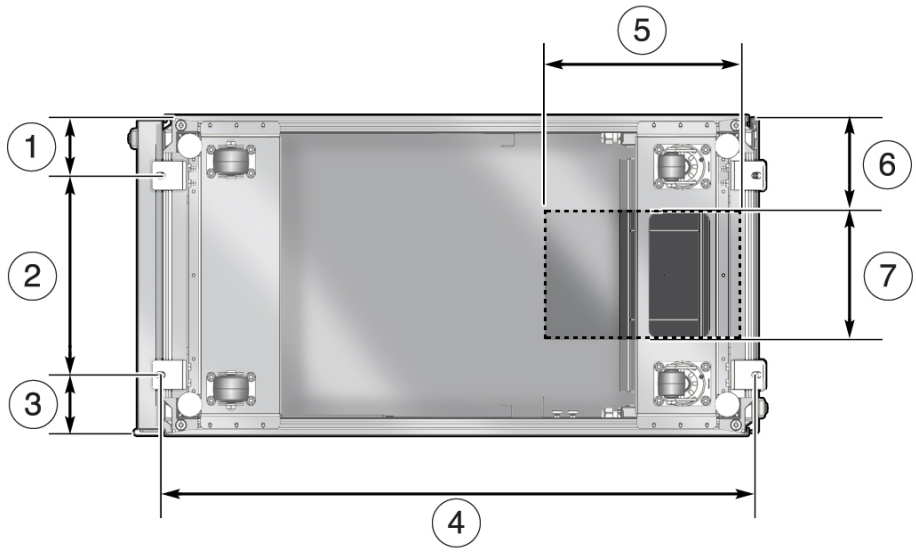

Figure Legend

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 floor cutout is 280 mm (11 inches)

Related Information

- “Perforated Floor Tiles” on page 69
- “Physical Specifications” on page 60
- “Installation and Service Area” on page 60

Reviewing Power Requirements

System Power Consumption

Comments	Full Rack	Half Rack
Maximum	15.81 kVA	8.97 kVA
Typical	15.02 kW	8.52 kW
	13.80 kVA	6.9 kVA
	13.11 kW	6.5 kW

Related Information

- “Facility Power Requirements” on page 62
- “Grounding Requirements” on page 63
- “PDU Power Requirements” on page 63

Facility Power Requirements

Provide a separate circuit breaker for each power cord.

Use dedicated AC breaker panels for all power circuits that supply power to the PDU. Breaker switches and breaker panels should not be shared with other high-powered equipment.

Balance the power load between AC supply branch circuits.

To protect the rack from electrical fluctuations and interruptions, you should have a dedicated power distribution system, an uninterruptible power supply (UPS), power-conditioning equipment, and lightning arresters.

Related Information

- "System Power Consumption" on page 62
- "Grounding Requirements" on page 63
- "PDU Power Requirements" on page 63

Grounding Requirements

Always connect the cords to grounded power outlets. Computer equipment requires electrical circuits to be grounded to the Earth.

Because different grounding methods vary by locality, 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.

Related Information

- "Facility Power Requirements" on page 62
- "System Power Consumption" on page 62
- "PDU Power Requirements" on page 63

PDU Power Requirements

When ordering SPARC SuperCluster T4-4 you selected one of these options:

- Low or high voltage
- Single or three phase power

Refer to the following tables for Oracle marketing and manufacturing part numbers.

Note - Each SPARC SuperCluster T4-4 has two power distribution units (PDUs). Both PDUs in a rack must be the same type.

TABLE 6 PDU Choices

Voltage	Phases	Reference
Low	1	Table 7
Low	3	Table 8
High	1	Table 9
High	3	Table 10

TABLE 7 Low Voltage 1 Phase PDUs for SPARC SuperCluster T4-4

Low Voltage	Single Phase	Comments
kVA Size	22 kVA	
Marketing part number	7100873	
Manufacturing part number	7018123	
Phase	1 ph	
Voltage	200-240 VAC	
Amps per PDU	110.4A (3 inputs x 36.8A)	
Outlets	42 C13	
	6 C19	
Number of inputs	3 inputs x50A, 1 ph	
Input current	36.8A max. current (per input)	
Data center receptacle	Hubbell CS8264C	
Outlet groups per PDU	6	
Usable PDU power cord length	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 8 Low Voltage 3 Phase PDUs for SPARC SuperCluster T4-4

Low Voltage	Three Phase	Comments
kVA Size	24 kVA	
Marketing part number	6444A	
Manufacturing part number	597-1162-01	
Phase	3 ph	
Voltage	200 to 240 VAC	
Amps per PDU	120A (6 x 20A)	
Outlets	42 C13	
	6 C19	
Number of inputs	2 inputs x 60A, 3 ph	
Input current	34.6 A max. per phase	
Data center receptacle	IEC 309-3P4W-IP67	(60A, 250V, AC, 3 ph)
Outlet groups per PDU	Hubbell HBL460R/C9W or equivalent	IEC309 60A 3ph 4 Wire

Low Voltage	Three Phase	Comments
Usable PDU power cord length	4	
kVA Size	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 9 High Voltage 1 Phase PDUs for SPARC SuperCluster T4-4

High Voltage	Single Phase	Comments
kVA Size	22 kVA	
Marketing part number	7100874	
Manufacturing part number	7018124	
Phase	1 ph	
Voltage	200 to 240 VAC	
Amps per PDU	96A (3 inputs x 32A)	
Outlets	42 C13 6 C19	
Number of inputs	3 inputs x 32A, 1 ph	
Input current	32A max. per input	
Data center receptacle	IEC 309-2P3W-IP44	(32A, 250V, AC, 3 ph)
Outlet groups per PDU	Hubbell HBL332R/C9W or equivalent	IEC309 32A 1 ph 3 Wire
Usable PDU power cord length	6	
kVA Size	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 10 High Voltage 3 Phase PDUs for SPARC SuperCluster T4-4

High Voltage	Three Phase	Comments
kVA Size	24 kVA	
Marketing part number	6445A	
Manufacturing part number	597-1163-01	
Phase	3 ph	
Voltage	220/380 to 240/415 VAC	
Amps per PDU	109A (6 inputs x 18.1A)	
Outlets	42 C13 6 C19	
Number of inputs	2 inputs x25A, 3 ph	
Input current	18 A max. per input	
Data center receptacle	IEC 309-4P5W-IP44	(32A, 400V, AC, 3 ph)
Outlet groups per PDU	Hubbell HBL532R/C9W or equivalent.	IEC309 32A 3ph 5 Wire
Usable PDU power cord length	4	

High Voltage	Three Phase	Comments
kVA Size	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

Related Information

- “Facility Power Requirements” on page 62
- “Grounding Requirements” on page 63
- “System Power Consumption” on page 62

Preparing for Cooling

Environmental Requirements

The following table lists temperature, humidity and altitude requirements.

Condition	Operating Requirement	Non-operating Requirement	Comments
Temperature	5 to 32°C (41 to 89.6°F)	-40 to 70°C (-40 to 158°F).	For optimal rack cooling, data center temperatures from 21 to 23°C (70 to 47°F)
Relative humidity	10 to 90% relative humidity, noncondensing	Up to 93% relative humidity.	For optimal data center rack cooling, 45 to 50%, noncondensing
Altitude	3048 m (10000 ft.) maximum[†]	12000 m (40000 ft.).	Ambient temperature is de-rated by 1 degree Celsius per 300 m above 900 m altitude above sea level

[†]Except in China where regulations may limit installations to a maximum altitude of 2000 m (6560 ft.).

Related Information

- “Heat Dissipation and Airflow Requirements” on page 67
- “Perforated Floor Tiles” on page 69

Heat Dissipation and Airflow Requirements

The following table lists the maximum rate of heat released from a system. In order to cool the system properly, ensure that adequate airflow travels through the system.

Comments	Full Rack	Half Rack
Maximum	53,966 BTU/hour	30,610 BTU/hour
	56,937 kJ/hour	32,295 kJ/hour
Typical	47,087 BTU/hour	23,543 BTU/hour
	49,679 kJ/hour	24,839 kJ/hour

The direction of airflow is front to back for the main and expansion racks.

Caution - Do not restrict the movement of cool air from the air conditioner to the cabinet, or the movement of hot air out of the rear of the cabinet.

Observe these additional requirements:

- Allow a minimum clearance of 914 mm (36 inches) at the front of the rack, and 914 mm (36 inches) at the rear of the rack for ventilation. There is no airflow requirement for the left and right sides, or the top of the rack.
- If the rack is not completely filled with components, cover the empty sections with filler panels.

FIGURE 7 Direction of Airflow Is Front to Back

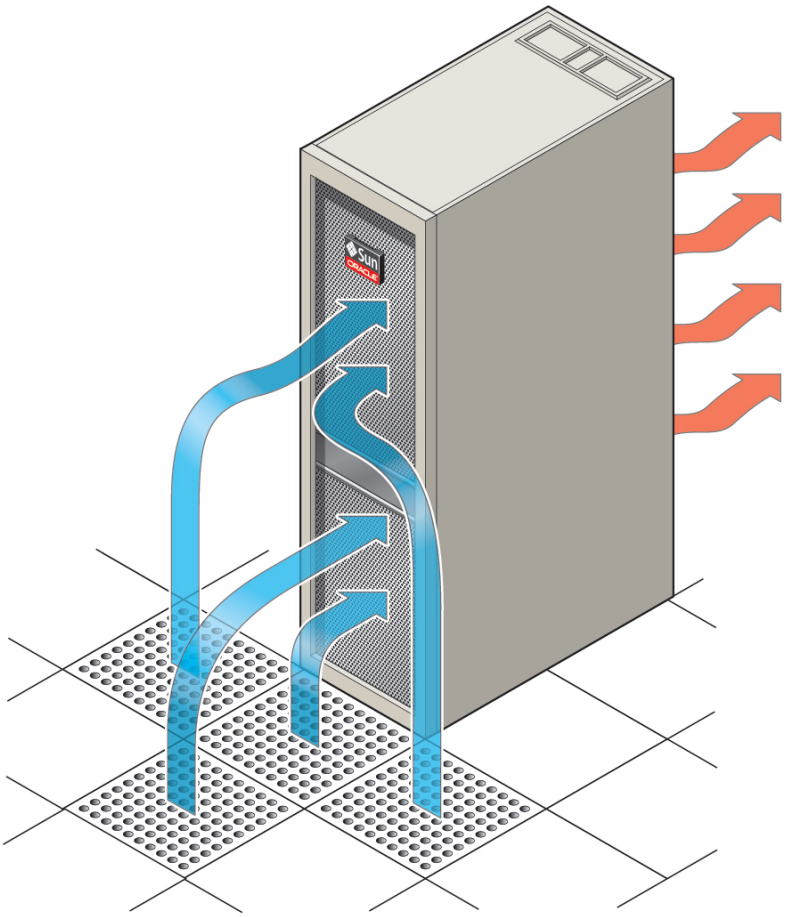

TABLE 11 Airflow (listed quantities are approximate)

Rack		Maximum	Typical
Full rack	CFM	2,498	2,180
Half rack	CFM	1,417	1,090

Related Information

- “Environmental Requirements” on page 66
- “Perforated Floor Tiles” on page 69

Perforated Floor Tiles

The default installation for the system assumes a slab floor, but if you install the server on a raised floor, use perforated tiles in front of the rack to supply cold air to the system. Each tile should support an airflow of approximately 400 CFM.

Perforated 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.

The following is the recommended number of floor tiles:

Rack	Number of Tiles
Full rack	6
Half rack	4

Related Information

- “Heat Dissipation and Airflow Requirements” on page 67
- “Environmental Requirements” on page 66
- “Rack and Floor Cutout Dimensions” on page 61

Preparing the Unloading Route and Unpacking Area

Ensure the system can get from the loading ramp to the installation site.

- “Shipping Package Dimensions” on page 69
- “Loading Dock and Receiving Area Requirements” on page 70
- “Access Route Guidelines” on page 70
- “Unpacking Area” on page 71

Shipping Package Dimensions

The following package dimensions apply to SPARC SuperCluster T4-4.

Parameter	Metric	English
Height	2159 mm	85 in.
Width	1219 mm	48 in.

Parameter	Metric	English
Depth	1575 mm	62 in.
Shipping weight (full rack)	1009.2 kg	2225 lbs
Shipping weight (half rack)	721 kg	1586 lbs

Related Information

- “Loading Dock and Receiving Area Requirements” on page 70
- “Access Route Guidelines” on page 70
- “Unpacking Area” on page 71
- “Physical Specifications” on page 60

Loading Dock and Receiving Area Requirements

Before the rack arrives, ensure that the receiving area is large enough for the shipping package.

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

When the rack arrives, leave the rack in its shipping packaging until it arrives in its installation site.

Note - Acclimatization time: If the shipping package is very cold or hot, allow the packaging to stand unopened in the computer room or a similar environment to come to the same temperature as the computer room. Acclimatization might require up to 24 hours.

Related Information

- “Shipping Package Dimensions” on page 69
- “Access Route Guidelines” on page 70
- “Unpacking Area” on page 71
- “Physical Specifications” on page 60

Access Route Guidelines

Leave the server in its shipping container until it reaches its final destination.

Plan the access route to enable smooth transport of the system. The entire access route to the installation site should be free of raised patterns that can cause vibration. Avoid obstacles such as doorways or elevator thresholds that can cause abrupt stops or shocks, and the route must meet the following requirements.

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 in.)	2040 mm (80.32 in.)
Minimum door width	1220 (48 in.)	600 mm (23.62 in.)
Minimum elevator depth	1575 mm (62 in.)	1200 mm (47.24 in.)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity (Shipping weights, see ??? on page 69)	1134 kg (2500 lbs)	1134 kg (2500 lbs)

Related Information

- “Shipping Package Dimensions” on page 69
- “Loading Dock and Receiving Area Requirements” on page 70
- “Unpacking Area” on page 71
- “Physical Specifications” on page 60

Unpacking Area

Use a conditioned space to remove the packaging material to reduce particles before entering the data center.

Allow enough space for unpacking it from its shipping cartons. Ensure that there is enough clearance and clear pathways for moving the rack from the unpacking location to the installation location.

Related Information

- “Shipping Package Dimensions” on page 69
- “Access Route Guidelines” on page 70
- “Loading Dock and Receiving Area Requirements” on page 70
- “Physical Specifications” on page 60

Preparing the Network

Prepare your network for the SPARC SuperCluster T4-4.

Network Connection Requirements

The minimum number of physical network connections are listed below.

Parameter	Comments	Full Rack	Half Rack
Network drops (minimum quantities)	1 Gb Ethernet	1 x 1 GB	1 x 1 GB
	10 Gb Ethernet	**4** x 10 GB if using one connection per SPARC T4-4†	**2** x 10 GB if using one connection per SPARC T4-4†
		8 x 10 GB if using two connections per SPARC T4-4	**4** x 10 GB if using two connections per SPARC T4-4
		32 x 10 GB if using four connections per SPARC T4-4	**16** x 10 GB if using four connections per SPARC T4-4

†You will not have redundancy using IPMP if there is only one connection per SPARC T4-4.

Related Information

Network IP Address Requirements

For the SPARC SuperCluster T4-4, the number of IP addresses that you will need for each network varies depending on the type of configuration you choose for your system. For more information on the number of IP addresses required for your configuration, refer to the appropriate SPARC SuperCluster T4-4 configuration worksheet.

Related Information

- “Prepare DNS for the System” on page 73
- “Understanding the Network Requirements” on page 46

▼ Prepare DNS for the System

Before You Begin You must finish the following tasks before the system arrives. Installation and the initial configuration cannot proceed until these tasks are finished.

Note - Grid Naming Service (GNS) is not configured on the rack until after initial configuration.

---

1. **Provide the necessary information in the following documents:**
   - SPARC SuperCluster T4-4 Site Checklists
   - The configuration worksheet document for your type of system
2. **Use the host names and IP addresses specified in the completed configuration worksheets document to create and register Domain Name System (DNS) addresses for the system.**

   All public addresses, single client access name (**SCAN**) addresses, and VIP addresses must also be registered in DNS prior to installation.

   Note - The configuration worksheets document defines the SCAN as a single name with three IP addresses on the client access network.

   ---

3. **Configure all addresses registered in DNS for both forward resolution and reverse resolution.**

   Reverse resolution must be forward confirmed (forward-confirmed reverse DNS) such that both the forward and reverse DNS entries match each other.

   The SCAN name to the three SCAN addresses must be configured in DNS for round robin resolution.

Related Information

- The configuration worksheets document for additional information about the worksheets
- *Oracle Grid Infrastructure Installation Guide for Linux* for additional information about SCAN addresses
- Your DNS vendor documentation for additional information about configuring round robin name resolution
- “Network IP Address Requirements” on page 72
- “Network Connection Requirements” on page 72

Installing the System

This chapter explains how to install the SPARC SuperCluster T4-4.

Installation Overview

This is a summary of the installation process.

1. Prepare for installation. 1. Review the safety precautions, guidelines, site checklists, and site requirements. 2. Ensure that the site is prepared for the installation. 3. Complete the worksheets.	**Reference** "Oracle Safety Information" on page 76 "Preparing the Site" *SPARC SuperCluster T4-4 Site Checklists* and *SPARC SuperCluster T4-4 Configuration Worksheets*
2. Begin the installation. 1. Unpack the SPARC SuperCluster T4-4. 2. Place the SPARC SuperCluster T4-4 in its allocated space. 3. Perform preliminary checks before connecting the power cords. 4. Verify the hardware by visually inspecting the system.	"Unpacking the System" on page 77.
3. Power on the PDUs in the rack. 1. Connect rack power. 2. Switch on the six PDU circuit breakers located on the rear of the main PDU (PDU A). 3. Wait 3 to 5 minutes for power-on self-test device checks to finish and for all Oracle ILOM service processors to boot.	"Powering on the System the First Time" on page 84

4. Verify that server standby power is on for each SPARC T4-4 server in the SPARC SuperCluster T4-4. 5. Verify that the main power is on for each database compute mode.	
4. The Sun ZFS Storage 7320 storage controllers should start up automatically. If the storage controllers do not start up: 1. Press the soft switches located on the front of the two Sun ZFS Storage 7320 storage controllers. 2. Wait 3 to 5 minutes for the Sun Storage 7320 System to initiate NFS services, daemons, and basic services. 3. Ping the IP address assigned to the Sun ZFS Storage 7320 storage controllers to verify whether the system is up and running.	
5. Start the Exadata Storage Cells. Wait for Oracle ILOM to boot, then either press the Power On buttons, or turn on the Exadata Storage Servers through Oracle ILOM.	
6. Start the SPARC T4-4 servers. 1. Press the soft switches located on the front of the SPARC T4-4 servers. 2. Verify that power is applied to the Ethernet switch. 3. Verify that power is applied to the Sun Datacenter Infiniband Switch 36 switches.	
7. Configure the SPARC SuperCluster T4-4 system.	
8. Finish the installation.	
(Optional) Install an express module to copy a database from another server.	"Using an Optional Fiber Channel Express Module" on page 96

Oracle Safety Information

Become familiar with Oracle's safety information before installing any Oracle server or equipment:

- Read the safety notices printed on the product packaging.
- Read the *Important Safety Information for Sun Hardware Systems* (821-1590) document that is included with the rack.
- Read all safety notices in the *Sun Rack II Safety and Compliance Guide* (820-4762) This guide is available at `http://www.oracle.com/documentation`.
- Read all safety notices in the *Sun Rack II Power Distribution Units Users Guide* (820-4760). This guide is also available at `http://www.oracle.com/documentation`.
- Read the safety labels that are on the equipment.

Unpacking the System

Tools for Installation

The following additional tools are required for installation:

- No. 2 Phillips screwdriver
- Antistatic wrist strap
- A tool to cut plastic strapping tape

▼ Find the Unpacking Instructions

- **Locate the unpacking instructions.**
  The unpacking instructions are attached to the outside of the shipping package.

This figure shows the major components of the shipping package.

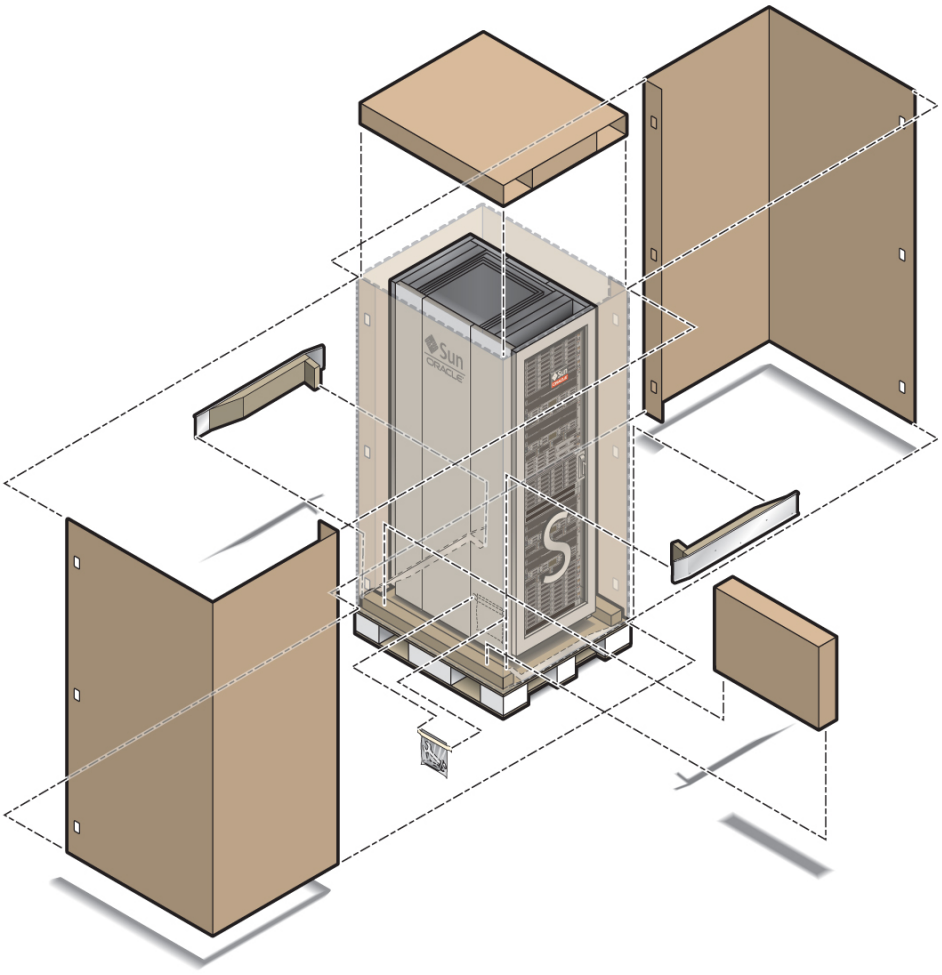

▼ Unpack and Inspect the Machine

Caution - Rocking or tilting the rack on the shipping pallet can cause it to fall over and cause serious injury or death.

1. **Use a No. 2 Phillips screwdriver to remove the shipping brackets from the lower front edges of each rack.**

2. **Check the rack for damage.**

3. **Check the rack for loose or missing screws.**

4. **Check the rack for the ordered configuration.**

   - Refer to the Customer Information Sheet (CIS) on the side of the packaging.
   - Two wrenches are included. The 12-mm wrench is for the cabinet leveling feet. The 17-mm wrench is for the bolts that hold the rack to the shipping pallet.
   - Two keys are included for the cabinet doors.
   - A box of other hardware and spare parts is included in the shipping carton. Note that some of these parts are not required for this installation.

5. **Check that all cable connections are secure and firmly in place:**

   a. **Check the power cords.**

      Ensure that the correct connectors have been supplied for the data center facilities power source.

   b. **Check the network data cables.**

6. **Check the floor tile preparations, if applicable.**

   See “Perforated Floor Tiles” on page 69.

7. **Check the data center airflow around the installation site.**

   See “Heat Dissipation and Airflow Requirements” on page 67 for more information.

Moving the Rack Into Place

▼ Move the SPARC SuperCluster T4-4

1. **Plan a route to the installation site.**

   See “Access Route Guidelines” on page 70.

2. **Ensure that the SPARC SuperCluster T4-4 doors are closed and secured.**

3. **Ensure that all four leveling feet on the rack are raised and out of the way.**

Caution - Use at least two people to move the rack: one person in front and one person in back to help guide the rack. Move the rack slowly at 0.65 meters per second (approximately two feet per second) or slower.

4. **Push the SPARC SuperCluster T4-4 from behind to the installation site.**

Note - The front casters do not swivel; you must steer the cabinet by turning the rear casters.

Caution - Never push on the side panels to move the rack. Pushing on the side panels can tip the rack over.

Caution - Never tip or rock the rack. It can fall over.

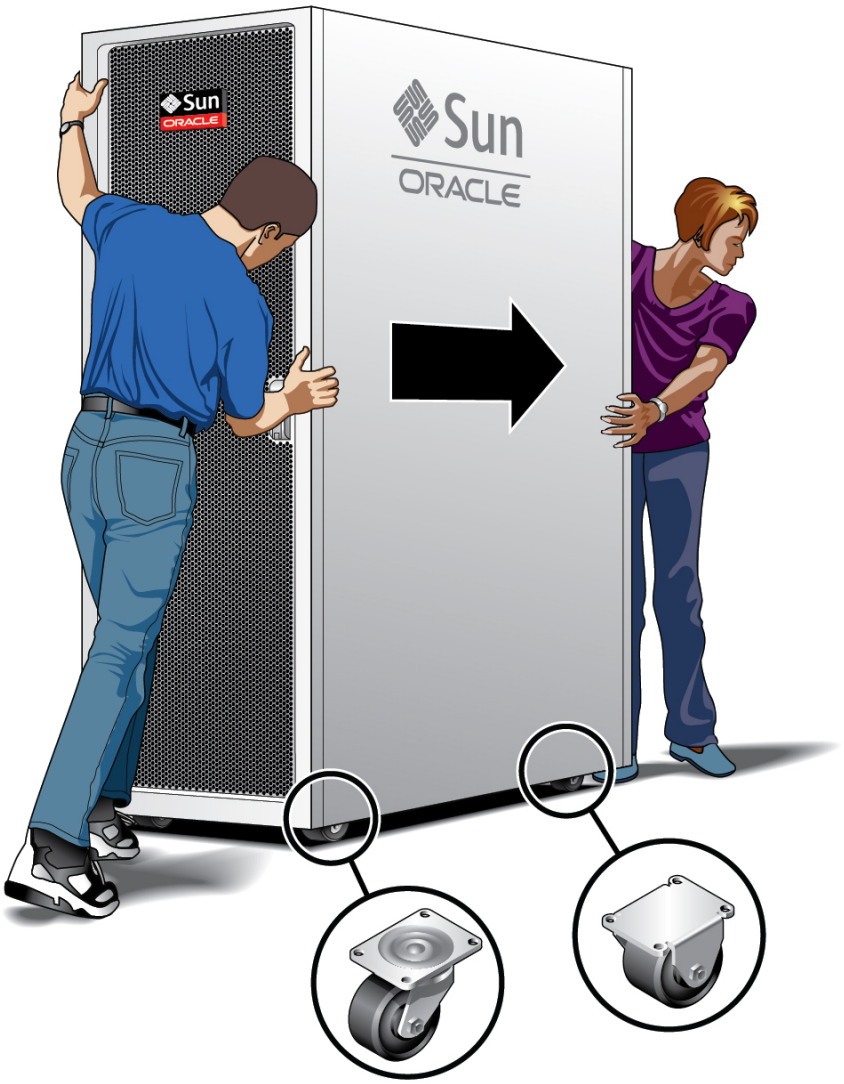

Install a Ground Cable (Optional)

The SPARC SuperCluster T4-4 power distribution units (PDUs) achieve earth ground through their power cords. For additional grounding, attach a chassis earth ground cable to the SPARC SuperCluster T4-4. The additional ground point enables electrical current leakage to dissipate more efficiently.

Caution - Do not install a ground cable until you confirm that there is proper PDU receptacle grounding. The PDU power input lead cords and the ground cable must reference a common earth ground.

Note - A grounding cable is not shipped with the system.

1. **Ensure that the installation site has properly grounded the power source in the data center.**

   An earth ground is required. See “Grounding Requirements” on page 63

2. **Ensure that all grounding points, such as raised floors and power receptacles, reference the facilities ground.**

3. **Ensure that direct, metal-to-metal contact is made for this installation.**

   The ground cable attachment area might have a painted or coated surface that must be removed to ensure solid contact.

4. **Attach the ground cable to one of the attachment points located at the bottom rear of the system frame.**

The attachment point is an adjustable bolt that is inside the rear of the rack on the right side.

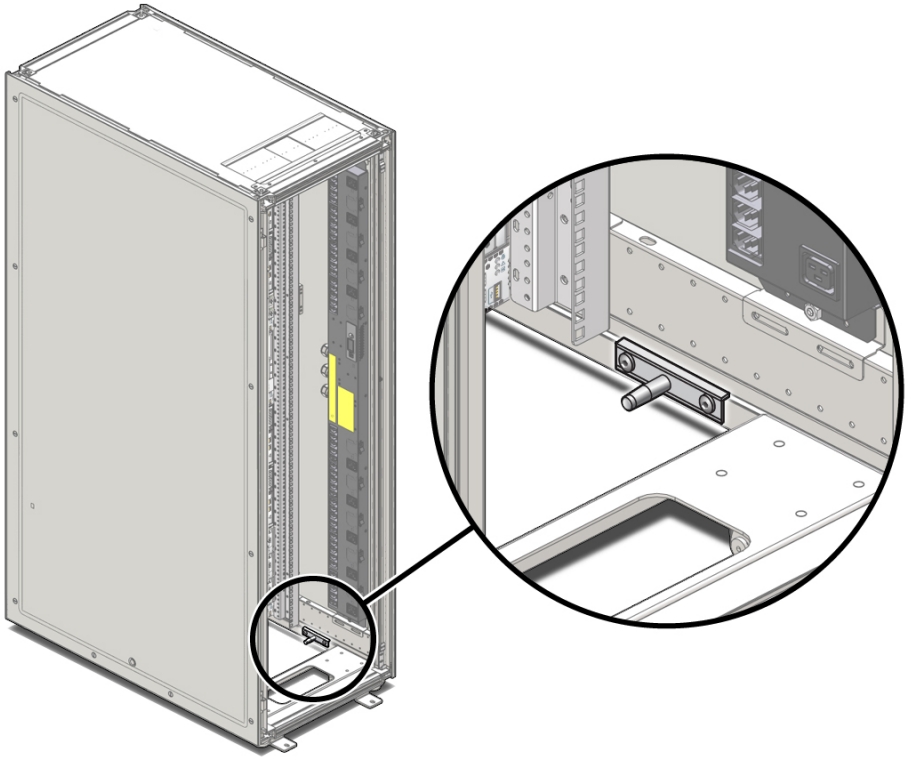

▼ Adjust the Leveling Feet

There are leveling feet at the four corners of the rack.

1. **Locate the 12-mm wrench inside the rack.**
2. **Use the wrench to lower the leveling feet to the floor.**

When lowered correctly, the four leveling feet should support the full weight of the rack.

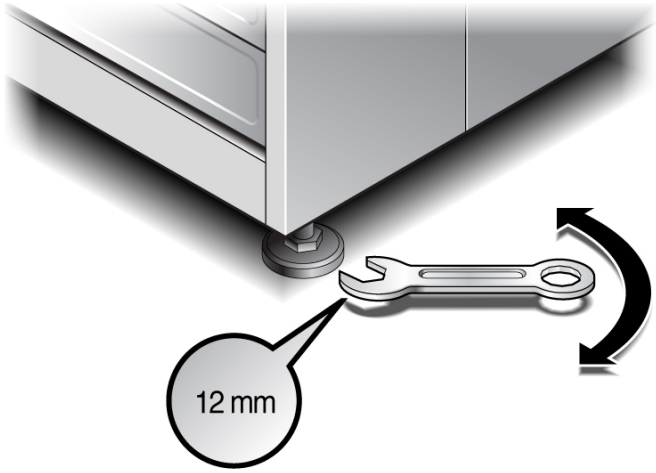

Powering on the System the First Time

▼ Connect Power Cords to the Rack

1. **At the facilities breaker panel, verify that all breakers for the rack are in the Off position.**

Caution - If the circuit breakers are in the On position, destructive sparking might occur when you attach the AC cables to the rack.

2. **Open the rear cabinet door.**

3. **Verify that the switches on the PDUs are in the Off position.**

   Ensure that both PDUs are turned completely off.

PDU-A is at the left side of the cabinet. PDU-B is at the right. Each PDU has six switches (circuit breakers), one for each socket group.

FIGURE 8 Power Switches on PDU

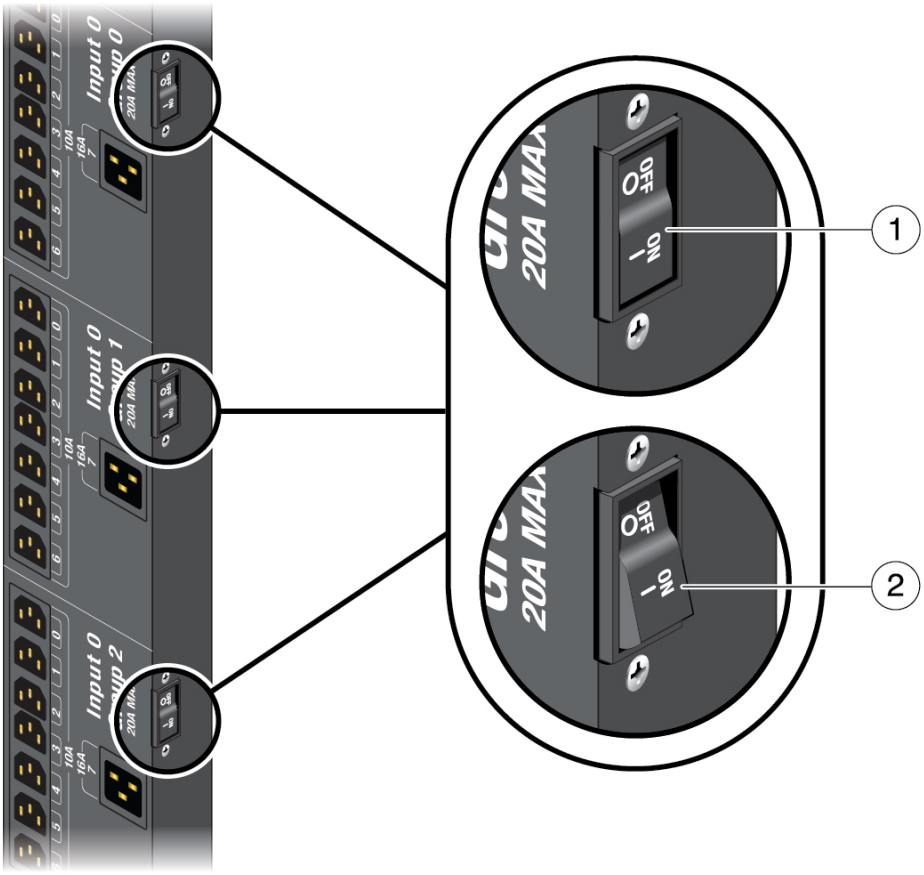

Figure Legend

1 Power switch lies flat in the On position.
2 Power switch is tilted in the Off position.

4. **Ensure that the correct power connectors have been supplied with the power cords.**

5. **Unfasten the power cord cable ties.**

The ties are for shipping only and are no longer needed.

6. **Route the power cords to the facility receptacles either above the rack or below the flooring.**

FIGURE 9 Routing Cables Below Flooring

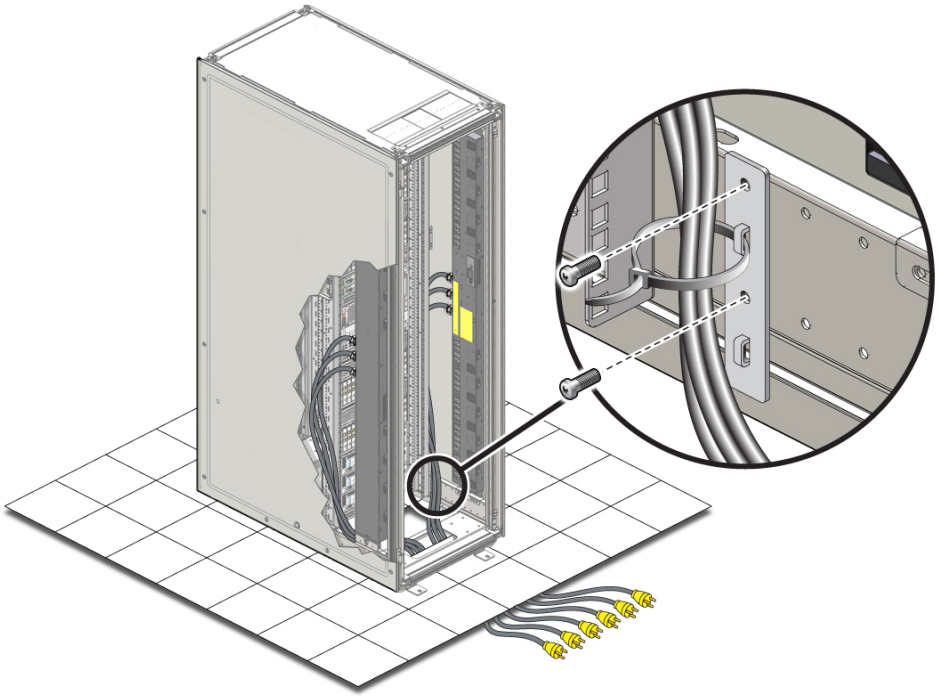

FIGURE 10 Routing Cables Out the Top of the Rack

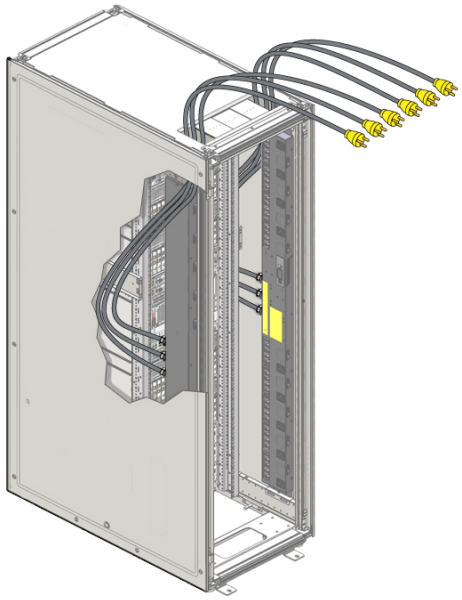

7. **Secure the power cords in bundles.**

8. **Connect the PDU power cord connectors into the facility receptacles.**

▼ Power On the SPARC SuperCluster T4-4

1. **Ensure that each of the main power cords is connected.**

2. **Turn on the facilities circuit breakers.**

3. **Switch on power distribution unit B (PDU B) only.**

Note - Do not turn on PDU A at this time.

PDU B is located on the right side of the rear of the rack. See below.

Press the ON (|) side of the toggle switches on PDU B.

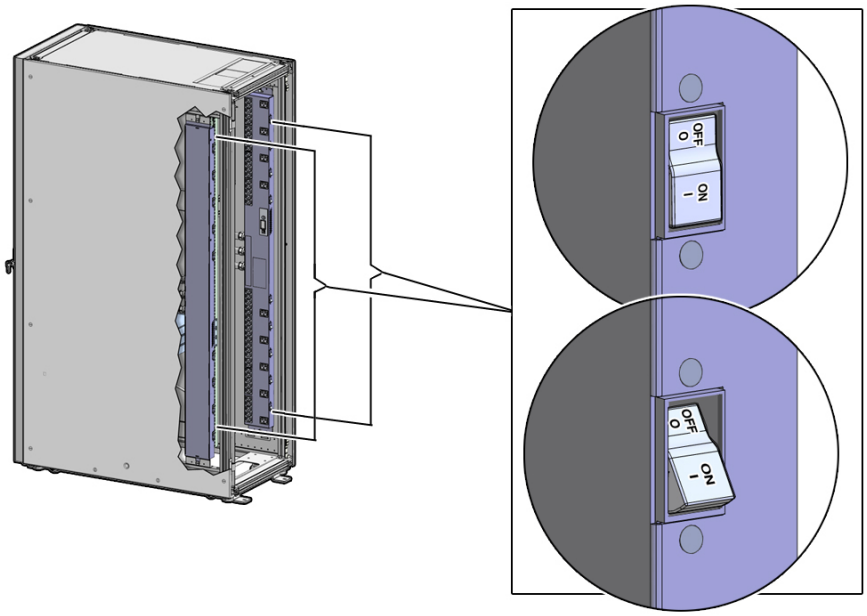

Powering on PDU B will power on only half of the power supplies in the rack. The remaining power supplies will be powered on in Step 4.

Note - For the location of each of the components, see "Identifying Hardware Components" on page 16.

The LEDs for the components should be in the following states when all of the PDU B circuit breakers have been turned on:

a. Check the SPARC T4-4 servers:

- Power OK green LED – Blinking
- Service Action Required amber LED – Off
- PS1 and PS3 Power LEDs – Green
- PS0 and PS2 Power LEDs – Amber

If the LEDs are not in these states, press the Power buttons at the fronts of the SPARC T4-4 servers.

b. Check the Sun ZFS Storage 7320 storage controllers:

- Power OK green LED (front panel) – Blinking while the operating system is booting up.

  The Power OK LED changes to Steady On after the operating system has booted up (this could take up to 10 minutes)
- Service Action Required amber LED (front panel) – Off
- Left power supply (rear panel) – Green
- Right power supply (rear panel) – Off

If the LEDs are not in these states, press the Power buttons at the fronts of the Sun ZFS Storage 7320 storage controllers.

c. At the rear of the Sun Disk Shelf, press both Power buttons to the On positions.

The LEDs should be:

- Service Action Required amber LED – Off
- Left power supply (rear panel) – Amber
- Right power supply (rear panel) – Green

d. Check the fronts of the Sun Datacenter InfiniBand Switch 36 switches:

- Left power supply LED (PS0 LED) – Red
- Right power supply LED (PS1 LED) – Green

e. Check the front of the Cisco Catalyst 4948 Ethernet management switch:

- PS1 LED – Red
- PS2 LED – Green

4. Switch on power distribution unit A (PDU A).

PDU A is located on the left side of the rack.

Press the ON (|) side of the toggle switches on PDU A.

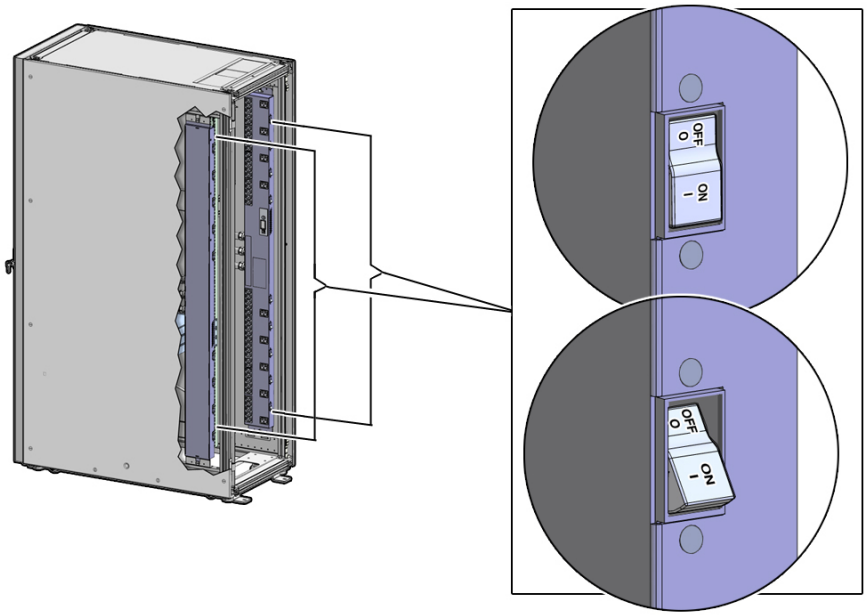

The LEDs for the components should be in the following states when all of the PDU A circuit breakers have been turned on.

a. Check the SPARC T4-4 servers:

- Power OK green LED – Blinking
- Service Action Required amber LED – Off
- PS1 and PS3 Power LEDs – Green
- PS0 and PS2 Power LEDs – Green

b. Check the Sun ZFS Storage 7320 storage controllers:

- Power OK green LED – On
- Service Action Required amber LED – Off
- Left power supply – Green
- Right power supply – Green

Tip - You can ping the IP address assigned to the Sun ZFS Storage 7320 storage controller to verify whether the system is up and running. For the default NET0 IP addresses, see "Default IP Addresses" on page 51. 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 SPARC SuperCluster T4-4" on page 91.

c. Check the Sun Disk Shelf:

- Service Action Required amber LED – Off
- Left power supply – Green
- Right power supply – Green

d. Check the Exadata Storage Server:

- Power OK LED – Off while Oracle ILOM is booting (about 3 minutes), then blinking
- Service Action Required amber LED – Off

e. Check the fronts of the Sun Datacenter InfiniBand Switch 36 switches:

- Left power supply LED (PS0 LED) – Green
- Right power supply LED (PS1 LED) – Green

f. Check the front of the Cisco Catalyst 4948 Ethernet management switch:

- PS1 LED – Green
- PS2 LED – Green

▼ Connect a Laptop to the SPARC SuperCluster T4-4

1. **Ensure that the laptop has 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 Catalyst 4948 Ethernet switch.**

   This switch is inside a vented filler panel in Unit 19 of your SPARC SuperCluster T4-4 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 SPARC SuperCluster T4-4 configurations.

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

   - **If you are using the Windows operating system, do the following:**

     a. **Go to Control Panel > Network Connections. Select your wired network adapter in the list of network connections, then 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 type a static IP address for your laptop.**

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 SPARC T4-4 servers assigned at the time of manufacturing or the custom IP address that you reconfigured using the Oracle SPARC SuperCluster T4-4 Configuration Utility tool. For the list of default NET0 IP addresses, see "Default IP Addresses" on page 51.

     d. **Although a default gateway is not necessary, enter the same IP address in the Default Gateway field.**

     e. **Click OK to exit the network connections screen.**

   - **If you are using a Linux operating system, do the following:**

     a. **Log in as a root user.**

     b. **At the command prompt, type the following command to display the network devices, such as ETH0, attached to the SPARC SuperCluster T4-4:**

        ```
        # ifconfig -a
        ```

The list of network devices or adapters attached to the SPARC SuperCluster T4-4 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 192.0.2.10 netmask 255.255.255.0 up
```

In this example, the `ifconfig` command assigns the IPv4 address *192.0.2.10*, 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 SPARC T4-4 server that is up and running.

Note - After you cable your laptop to the Cisco Catalyst 4948 Ethernet switch, you can use the *NET0* IP addresses of SPARC SuperCluster T4-4 components to communicate with them. For a list of default IP addresses assigned at the time of manufacturing, see "Default IP Addresses" on page 51.

Note - If you or the Oracle installer have not run the Oracle SPARC SuperCluster T4-4 Configuration Utility set of tools and scripts to reconfigure IP addresses for the SPARC SuperCluster T4-4, you can use a set of default IP addresses. If you or the Oracle installer have already run the Oracle SPARC SuperCluster T4-4 Configuration Utility set of tools and scripts, you can use the network IP address that you provided in the *SPARC SuperCluster T4-4. Configuration Worksheets and Site Checklists* document.

▼ Connect to a 10 GbE Client Access Network

Before You Begin

The IP addresses and host names for the client access network should be registered in the Domain Name System (DNS) prior to initial configuration. In addition, all public addresses, single client access name (SCAN) addresses, and VIP addresses should be registered in DNS prior to installation.

A 10 GbE client access network infrastructure is a required part of the installation process for the SPARC SuperCluster T4-4.

The following components have been included with the SPARC SuperCluster T4-4:

- Four dual-ported Sun Dual 10 GbE SFP+ PCIe NICs in each SPARC T4-4 server, preinstalled in PCIE slots 2, 6, 10, and 14 (see "Card and Port Locations (SPARC T4-4 Servers)" on page 25 for more information)
- Transceivers preinstalled in the 10 GbE NICs

- Four 10-meter SFP-QSFP optical splitter cables for a half rack or eight cables for a full rack

If you wish to use the supplied SFP-QSFP cables for the connection to your client access network, you must provide the following 10 GbE client access network infrastructure components:

- A 10 GbE switch with QSFP connections, such as the Sun Network 10GbE Switch 72p
- Transceivers for connections to your 10 GbE switch (four transceivers for a half rack or eight transceivers for a full rack)

If you do not want to use the supplied SFP-QSFP cables for the connection to your client access network, you must provide the following 10 GbE client access network infrastructure components:

- A 10 GbE switch
- Suitable optical cables with SFP+ connections for the SPARC T4-4 server side
- Suitable transceivers to connect all cables to your 10 GbE switch

If you do not have a 10 GbE client access network infrastructure set up at your site, you must have a 10 GbE network switch available at the time of installation that the SPARC SuperCluster T4-4 can be connected to, even if the network speed drops from 10 Gb to 1 Gb on the other side of the 10 GbE network switch. The SPARC SuperCluster T4-4 cannot be installed at the customer site without the 10 GbE client access network infrastructure in place.

1. **Locate the Sun Dual 10 GbE SFP+ PCIe NIC(s) that you will use to connect to the client access network.**

   The NIC connections will vary, depending on the type and number of domains that you have set up for your SPARC SuperCluster T4-4. See “10 GbE Client Access Network Physical Connections (SPARC T4-4 Servers)” on page 27 and “Understanding the SPARC SuperCluster T4-4 Configurations” on page 34 for more information.

2. **Locate the SFP cable that you will use for the connection to the 10 GbE client access network.**

   You can use the 10-meter SFP-QSFP optical splitter cables provided with the system or your own SFP cables.

3. **Connect the SFP side of the cable to the port on the 10 GbE NIC.**

4. **Connect the other end of the cable to the 10 GbE switch that you provided.**

   The following figure shows an example connection layout for the 10 GbE client access network, where:

   - The two QSFP ends of the two SFP-QSFP cables connect to the QSFP ports on the 10 GbE switch (in this example, the Sun Network 10GbE Switch 72p)

- All sixteen SFP+ ends of the two SFP-QSFP cables connect to the SFP+ ports on the 10 GbE NICs installed in the SPARC T4-4 server (some configurations would have fewer connections to the SFP+ ports on the 10 GbE NICs, depending on the configuration)

FIGURE 11 Example Connection to the 10 GbE Client Access Network

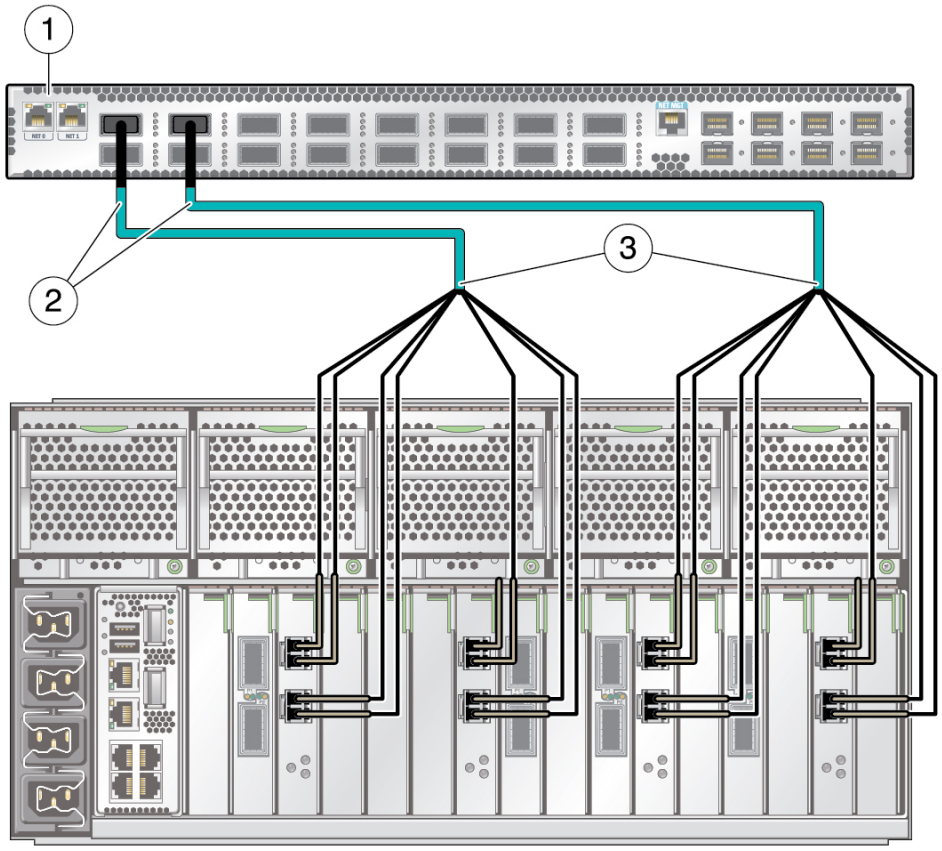

Figure Legend

1	10 GbE switch with QSFP connections (Sun Network 10GbE Switch 72p shown)
2	QSFP connector ends of SFP-QSFP cables, connecting to QSFP ports on 10 GbE switch
3	SFP+ connector ends of SFP-QSFP cables, connecting to SFP+ ports on 10 GbE NICs in SPARC T4-4 server

Using an Optional Fiber Channel Express Module

Optional Fiber Channel express modules are available to facilitate migration of data from legacy storage subsystems to the Exadata Storage Servers integrated with the SPARC SuperCluster T4-4. The optional Fiber Channel express modules are not included in standard SPARC SuperCluster T4-4 configurations and must be purchased separately.

The following options are supported:

- SG-XPCIEFCGBE-E8-N Sun StorageTek Dual 8Gb FC Dual 1 GbE HBA in ExpressModule form factor, RoHS-6 compliant, Emulex X-option
- SG-XPCIEFCGBE-Q8-N Sun StorageTek Dual 8Gb FC Dual 1 GbE HBA in ExpressModule form factor, RoHS-6 compliant, Qlogic. X-option

The optional Fiber Channel express modules can be installed in any or all of these remaining express module slots that are not populated by 10 GbE NICs or InfiniBand HCAs:

- Slot 0
- Slot 3
- Slot 4
- Slot 5
- Slot 8
- Slot 11
- Slot 13
- Slot 15

See “Card and Port Locations (SPARC T4-4 Servers)” on page 25 for more information on the slot locations on the SPARC T4-4 servers.

Slots for express modules in a SPARC T4-4 server are assigned to various domains. The selection of an express module slot depends on how domains are configured on your SPARC SuperCluster T4-4. See “Understanding the SPARC SuperCluster T4-4 Configurations” on page 34 for more information.

Note the following restrictions when using the optional Fiber Channel express modules:

- When installed in slots associated with Application Domains running either Oracle Solaris 10 or Oracle Solaris 11, the Fiber Channel express modules can be used for any purpose, including database file storage for supported databases other than Oracle Database 11gR2.
- When installed in slots associated with Database Domains, the Fiber Channel express modules can be used for data migration only, and not for storage of Oracle Database 11gR2 data.
- Oracle discourages the use of additional network interfaces based on the GbE ports on the Fiber Channel express modules. Oracle will not support questions or issues with networks based on these ports.

For instructions for installing an express module, refer to:

- The documentation provided with your express module
- The SPARC T4-4 server service manual at:

  http://docs.oracle.com/cd/E23411_01/html/E23412/

Maintaining the System

These topics describe maintenance information for the SPARC SuperCluster T4-4.

Cautions and Warnings

The following cautions and warnings apply to SPARC SuperCluster T4-4 systems.

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

Caution - Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to hardware components.

Caution - Keep the top, front, and back of the cabinets clear to allow proper airflow and prevent overheating of components.

Use only the supplied hardware.

Powering Off the System

- “Emergency Power Off” on page 102

Graceful Power Off Sequence

Use this procedure to shut down the system under normal circumstances.

1. Shut down Oracle Solaris Cluster:

```
# /usr/cluster/bin/cluster shutdown -g 0 -y
```

2. If Ops Center is running, shut down the enterprise controller:

```
# /opt/SUNWxvmoc/bin/ecadm ha-stop-no-relocate
```

3. Shut down the database using one of these methods:

   http://docs.oracle.com/cd/B28359_01/server.111/b28310/start003.htm

4. “Shut Down the Exadata Storage Servers” on page 100.
5. “Power Off the Exadata Storage Servers” on page 101
6. “Gracefully Shut Down LDoms” on page 101
7. Gracefully shut down each SPARC T4-4 server:

```
# init 0
```

8. Gracefully shut down the Sun ZFS Storage 7320 Storage Controller.

   Log in to the browser interface and click the power icon on the left side of the masthead.

9. Power off the switches, and the entire rack, by turning off the circuit breakers.

- To power the system back on, see “Power On the System” on page 102.

▼ Shut Down the Exadata Storage Servers

Perform this procedure for each Exadata Storage Server before you power them off. For more information on this task, see the Exadata documentation at:

http://wd0338.oracle.com/archive/cd_ns/E13877_01/doc/doc.112/e13874/maintenance.htm#CEGBHCJG

1. Run the following command to check if there are other offline disks:

```
CellCLI > LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes'
```

If any grid disks are returned, then it is not safe to take Exadata Storage Server offline, because proper Oracle ASM disk group redundancy will not be maintained. Taking Exadata Storage

Server offline when one or more grid disks are in this state will cause Oracle ASM to dismount the affected disk group, causing the databases to shut down abruptly.

2. **Inactivate all the grid disks when Exadata Storage Server is safe to take offline.**

```
CellCLI> ALTER GRIDDISK ALL INACTIVE
```

This command completes after all disks are inactive and offline.

3. **Verify that all grid disks are INACTIVE.**

```
CellCLI> LIST GRIDDISK WHERE STATUS != 'inactive'
```

If all grid disks are INACTIVE, then Exadata Storage Server can be shut down without affecting database availability.

4. **Shut down the cell.**
See "Power Off the Exadata Storage Servers" on page 101.

▼ Power Off the Exadata Storage Servers

Perform the following procedure for each Exadata Storage Server.

Note the following when powering off Exadata Storage Servers:

- All database and Oracle Clusterware processes should be shut down prior to shutting down more than one Exadata Storage Server.
- Powering off one Exadata Storage Server will not affect running database processes or Oracle ASM.
- Powering off or restarting Exadata Storage Servers can impact database availability.

- **The following command shuts down Exadata Storage Server immediately:**

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

▼ Gracefully Shut Down LDoms

The SPARC T4-4 server configurations vary based on the configuration chosen during installation. If running with one LDom you shutdown the machine just as you would any other server just by cleanly shutting down the OS. If running two Ldoms you will shutdown the guest domain first and then the primary (control). if running with three or more domains you will have to identify the domain(s) that is/are running off virtualized hardware and shut it/ them down first before moving on to shutting down the guest domain and finally the primary (control) domain.

1. **Shut down, stop, and unbind each of the non-I/O domains.**

```
# ldm stop domainname
LDom domainname stopped
# ldm unbind-domain domainname
```

2. **Shut down, stop, and unbind any active I/O domains.**

```
# ldm stop activedomainname
LDom activedomainname stopped
# ldm unbind-domain activedomainname
```

3. **Halt the primary domain.**

```
# shutdown -i5 -g0 -y
```

Because no other domains are bound, the firmware automatically powers off the system.

Emergency Power Off

If there is an emergency, such as earthquake or flood, an abnormal smell or smoke coming from the machine, or a threat to human safety, then power to the SPARC SuperCluster T4-4 should be halted immediately. In that case, use one of the following ways to power off the system.

- **Power off the system one of the following ways:**
  - Turn off power at the circuit breaker, or pull the emergency power-off switch in the computer room.
  - Turn off the site EPO switch to remove power from the SPARC SuperCluster T4-4.
  - Turn off the two PDUs in the rack.

After the emergency, contact Oracle Support Services to restore power to the machine.

Related Information

- “Graceful Power Off Sequence” on page 100

Power On the System

Power on the system in the reverse order of shutdown.

1. **Turn on both circuit breakers that provide power to the rack.**

The switches will power on, and the Exadata Storage Servers, SPARC T4-4 servers and the Sun ZFS Storage 7320 Storage Controllers will return to standby mode.

2. **Boot each Sun ZFS Storage 7320 Storage Controller.**
3. **Boot each SPARC T4-4 server.**
4. **Boot each Exadata Storage Server.**

Managing Oracle Solaris 11 Boot Environments

When the Oracle Solaris OS is first installed release on a system, a boot environment is created. You can use the `beadm(1M)` utility to create and administer additional boot environments on your system.

After your SPARC SuperCluster system is installed, create a backup of the original boot environment. If needed, you can then boot to the backup of the original boot environment. These topics describe how to create and manage boot environments on SPARC SuperCluster.

For more information about Oracle Solaris 11 boot environments, refer to:

`http://docs.oracle.com/cd/E23824_01/html/E21801/toc.html`

Advantages to Maintaining Multiple Boot Environments

Multiple boot environments reduce risk when updating or changing software because system administrators can create backup boot environments before making any updates to the system. If needed, they have the option of booting a backup boot environment.

The following examples show how having more than one Oracle Solaris boot environment and managing them with the `beadm` utility can be useful.

- You can maintain more than one boot environment on your system and perform various updates on each of them as needed. For example, you can clone a boot environment by using the `beadm create` command. The clone you create is a bootable copy of the original. Then, you can install, test, and update different software packages on the original boot environment and on its clone.

Although only one boot environment can be active at a time, you can mount an inactive boot environment by using the `beadm mount` command. Then, you could use the `pkg` command with the alternate root (`-R`) option to install or update specific packages on that environment.

- If you are modifying a boot environment, you can take a snapshot of that environment at any stage during modifications by using the `beadm create` command. For example, if you are doing monthly upgrades to your boot environment, you can capture snapshots for each monthly upgrade.

  Use the command as follows:

  ```
  # beadm create BeName@snapshotNamedescription
  ```

  The snapshot name must use the format, `BeName@snapshotdescription`, where BeName is the name of an existing boot environment that you want to make a snapshot from. Enter a custom snapshotdescription to identify the date or purpose of the snapshot.

  You can use the `beadm list -s` command to view the available snapshots for a boot environment.

  Although a snapshot is not bootable, you can create a boot environment based on that snapshot by using the `-e` option in the `beadm create` command. Then you can use the `beadm activate` command to specify that this boot environment will become the default boot environment on the next reboot.

For more information about the advantages of multiple Solaris 11 boot environments, see:

http://docs.oracle.com/cd/E23824_01/html/E21801/snap3.html#scrolltoc

▼ Create a Boot Environment

If you want to create a backup of an existing boot environment, for example, prior to modifying the original boot environment, you can use the `beadm` command to create and mount a new boot environment that is a clone of your active boot environment. This clone is listed as an alternate boot environment in the boot menu for SPARC systems.

1. Log in to the target system.

```
localsys% ssh systemname -l root
Password:
Last login: Wed Nov 13 20:27:29 2011 from dhcp-vpn-r
Oracle Corporation SunOS 5.11 solaris April 2011
root@sup46:~#
```

2. Manage ZFS boot environments with `beadm`.

```
root@sup46:~# beadm list

BE       Active    Mountpoint    Space    Policy     Created
------------------------------------------------------------
solaris    NR          /         2.17G    static   2011-07-13 12:01
```

Note - In the Active column, the first letter indicates the boot environment current status and the second letter indicates the status at next reboot. In the example above, N indicates the current (or Now) boot environment while the R indicates which boot environment will be active at next Reboot.

3. Create a new ZFS boot environment based on the current environment.

```
root@sup46:~# beadm create solaris_backup
root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris         NR        /       2.17G  static   2011-07-13 12:01
solaris_backup  -         -       35.0K  static   2011-07-17 21:01
```

4. Change to the next boot environment.

```
root@sup46:~# beadm activate solaris_backup
root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris_backup  R         -       2.17G  static   2011-07-17 21:01
solaris         N        /       1.86G  static   2011-07-13 12:01
```

5. Reboot to the new boot environment.

```
root@sup46:~# reboot
Connection to systemname closed by remote host.
Connection to systemname closed.
localsys% ssh systemname -l root
Password:
Last login: Thu Jul 14 14:37:34 2011 from dhcp-vpn-
Oracle Corporation SunOS 5.11 solaris April 2011

root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
```

```
solaris_backup  NR        -       2.19G  static   2011-07-17 21:01
solaris          -        /       4.12G  static   2011-07-13 12:01
```

Mount to a Different Build Environment

- **Use the following commands to mount to a different build environment and unmount the other build environment.**

```
root@sup46:~# beadm mount solaris_backup /mnt
root@sup46:~# df -k /mnt
Filesystem 1024-blocks Used Available Capacity Mounted on
rpool1/ROOT/solaris_backup
286949376 2195449 232785749 1% /mnt

root@sup46:~# df -k /
Filesystem 1024-blocks Used Available Capacity Mounted on
rpool1/ROOT/solaris_backup
286949376 2214203 232785749 1% /

root@sup46:~# ls /mnt
bin etc lib opt rpool1 system wwss
boot export media pkg sbin tmp
cdrom home micro platform scde usr
dev import mnt proc share var
devices java net re shared workspace
doe kernel nfs4 root src ws
root@sup46:~#

root@sup46:~# beadm umount solaris
root@sup46:~#
```

Reboot to the Original Boot Environment

- **Use the following commands to reboot to the original boot environment.**

```
root@sup46:~# beadm activate solaris
root@sup46:~# reboot
Connection to systemname closed by remote host.
Connection to systemname closed.
localsys%
ssh systemname -l root
Password: Last login: Thu Jul 14 14:37:34 2011 from dhcp-vpn-
Oracle Corporation SunOS 5.11 solaris April 2011
root@sup46:~#
```

▼ Remove Unwanted Boot Environments

- **Use the following commands to remove boot environments.**

```
root@sup46:~# beadm list

    BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris_backup  -         -      13.25G  static   2011-07-17 21:19
solaris         -         -       4.12G  static   2011-07-13 12:01

root@sup46:~# beadm destroy solaris_backup
Are you sure you want to destroy solaris_backup? This action cannot be undone(y/[n]): y
root@sup46:~# beadm list

   BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris       NR        /      4.12G  static   2011-07-13 12:01

root@sup46:~#
```

Using Dynamic Intimate Shared Memory

DISM Restrictions

Dynamic Intimate Shared Memory (DISM) is not supported for use on SPARC SuperCluster Solaris environments in instances other than the ASM instance. The use of DISM on the SPARC SuperCluster outside of the ASM instance can lead to several different issues ranging form excessive swap usage (even when memory is available) to kernel panics to performance problems. It has been determined that the ASM instance is typically such a small memory footprint that it should not cause an issue.

This behavior is typically occurs on instances created after installation, because Solaris 11 uses Automatic Memory Management by default. To prevent this DISM issue when creating Solaris 11 instances, disable DISM. For more information see: "Disable DISM" on page 108.

To decide if DISM is appropriate for your environment, and for more information about using DISM with an Oracle database, refer to the Oracle white paper "*Dynamic SGA Tuning of Oracle Database on Oracle Solaris with DISM*":

http://www.oracle.com/technetwork/articles/systems-hardware-architecture/using-dynamic-intimate-memory-sparc-168402.pdf

▼ Disable DISM

Dynamic Intimate Shared Memory (DISM) is not supported for use on SPARC SuperCluster Solaris environments in instances other than the ASM instance. For more information, see “ DISM Restrictions” on page 107.

Note - Do not disable the use of automatic shared memory management (ASMM) within the database, which is a very useful and desirable feature to reduce DBA management of the database.

- **Disable the use of DISM by the database on Solaris in one of two ways; either unsetting the SGA_MAX_SIZE / MEMORY_MAX_TARGET / MEMORY_TARGET parameters, or ensure SGA_MAX_SIZE is set to the same value as SGA_TARGET parameter or equal to the sum of all sga components in the instance.**

  - **For example, to set a 64 G SGA:**

```
alter system set SGA_TARGET=64G scope=spfile;
alter system set SGA_MAX_SIZE=64G scope=spfile;
alter system set MEMORY_MAX_TARGET=0 scope=spfile;
alter system set MEMORY_TARGET=0 scope=spfile;
```

Component-Specific Service Procedures

If you have a service contract for your SPARC SuperCluster T4-4, contact your service provider for maintenance.

If you do not have a service contract, refer to each component's documentation for general maintenance information.

TABLE 12 Component Service Documentation

Component Information	Maintenance Information
“SPARC T4-4 Servers” on page 20	http://download.oracle.com/docs/cd/E23411_01/index.html
“Sun ZFS Storage 7320 Appliance” on page 23.	http://download.oracle.com/docs/cd/E22471_01/html/820-4167/toc.html
“Sun Datacenter InfiniBand Switch 36 Switches” on page 23.	http://download.oracle.com/docs/cd/E19197-01/index.html
“Cisco Catalyst 4948 Ethernet Management Switch” on page 24	http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Component Information	Maintenance Information
"Exadata Storage Servers" on page 22	"Maintaining Exadata Storage Servers" on page 109
Oracle VM Server for SPARC	http://download.oracle.com/docs/cd/E23120_01/index.html
Oracle Solaris Cluster	http://download.oracle.com/docs/cd/E18728_01/index.html
Oracle Solaris 11	http://www.oracle.com/technetwork/documentation/solaris-11-192991.html
Oracle Solaris 10	http://download.oracle.com/docs/cd/E23823_01/index.html

Maintaining Exadata Storage Servers

The Exadata Storage Servers are highly optimized for use with the Oracle Database and employ a massively parallel architecture and Exadata Smart Flash Cache to dramatically accelerate Oracle Database processing and speed I/O operations. For more information, refer to: "Exadata Storage Servers" on page 22.

For general maintenance information, refer to the Exadata Storage Server documentation, located in the following directory on the Exadata Storage Servers installed in the SPARC SuperCluster T4-4:

```
/opt/oracle/cell/doc
```

These topics describe maintenance relevant to Exadata Storage Servers in the SPARC SuperCluster T4-4.

- "Monitor Write-through Caching Mode" on page 109
- "Shut Down an Exadata Storage Server" on page 111
- "Drop an Exadata Storage Server" on page 113

Related Information

- See the *Oracle Exadata Storage Server Software User's Guide* for additional information about the Oracle ASM disk repair timer.

▼ Monitor Write-through Caching Mode

The disk controller on each Exadata Storage Server periodically performs a discharge and charge of the controller battery. During the operation, the write cache policy changes from write-back caching to write-through caching. Write-through cache mode is slower than write-back cache mode. However, write-back cache mode has a risk of data loss if the Exadata Storage Server loses power or fails. For Exadata Storage Server releases earlier than release 11.2.1.3, the operation occurs every month. For Oracle Exadata Storage Server Software release

11.2.1.3 and later, the operation occurs every three months, for example, at 01:00 on the 17th day of January, April, July and October.

1. **To change the start time for when the learn cycle occurs, use a command similar to the following. The time reverts to the default learn cycle time after the cycle completes.**

```
CellCLI> ALTER CELL bbuLearnCycleTime="2011-01-22T02:00:00-08:00"
```

2. **To see the time for the next learn cycle, use the following command:**

```
CellCLI> LIST CELL ATTRIBUTES bbuLearnCycleTime
```

The Exadata Storage Server generates an informational alert about the status of the caching mode for logical drives on the cell, for example:

```
HDD disk controller battery on disk contoller at adapter 0 is going into a learn cycle. This
 is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12
 hours. The disk controller cache might go into WriteThrough caching mode during the learn
 cycle. Disk write throughput might be temporarily lower during this time. The message is
 informational only, no action is required.
```

3. **Use the following command to view the status of the battery:**

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -GetBbuStatus -a0
```

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

```
BBU status for Adapter: 0

BatteryType: iBBU08
Voltage: 3721 mV
Current: 541 mA
Temperature: 43 C

BBU Firmware Status:

Charging Status : Charging
Voltage : OK
Temperature : OK
Learn Cycle Requested : No
Learn Cycle Active : No
Learn Cycle Status : OK
Learn Cycle Timeout : No
I2c Errors Detected : No
Battery Pack Missing : No
Battery Replacement required : No
Remaining Capacity Low : Yes
```

```
Periodic Learn Required : No
Transparent Learn : No

Battery state:

GasGuageStatus:
Fully Discharged : No
Fully Charged : No
Discharging : No
Initialized : No
Remaining Time Alarm : Yes
Remaining Capacity Alarm: No
Discharge Terminated : No
Over Temperature : No
Charging Terminated : No
Over Charged : No

Relative State of Charge: 7 %
Charger System State: 1
Charger System Ctrl: 0
Charging current: 541 mA
Absolute State of Charge: 0%

Max Error: 0 %
Exit Code: 0x00
```

▼ Shut Down an Exadata Storage Server

When performing maintenance on Exadata Storage Servers, it may be necessary to power down or reboot the cell. If Exadata Storage Server is to be shut down when one or more databases are running, then verify that taking Exadata Storage Server offline will not impact Oracle ASM disk group and database availability. The ability to take Exadata Storage Server offline without affecting database availability depends on the level of Oracle ASM redundancy used on the affected disk groups, and the current status of disks in other Exadata Storage Servers that have mirror copies of data as Exadata Storage Server to be taken offline.

Use the following procedure describes to power down Exadata Storage Server.

1. **Run the following command to check if there are other offline disks:**

```
CellCLI> LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes'
```

If any grid disks are returned, then it is not safe to take Exadata Storage Server offline because proper Oracle ASM disk group redundancy will not be maintained. Taking Exadata Storage Server offline when one or more grid disks are in this state will cause Oracle ASM to dismount the affected disk group, causing the databases to shut down abruptly.

2. **Inactivate all the grid disks when Exadata Storage Server is safe to take offline:**

```
CellCLI> ALTER GRIDDISK ALL INACTIVE
```

The preceding command completes once all disks are inactive and offline.

3. **Verify that all grid disks are `INACTIVE` to allow safe shut down of Exadata Storage Server.**

```
CellCLI> LIST GRIDDISK WHERE STATUS != 'inactive'
```

If all grid disks are `INACTIVE`, then Exadata Storage Server can be shut down without affecting database availability.

4. **Shut down the cell.**

5. **After performing the maintenance, start the cell. The cell services start automatically.**

6. **Bring all grid disks online using the following command:**

```
CellCLI> ALTER GRIDDISK ALL ACTIVE
```

When the grid disks become active, Oracle ASM automatically synchronizes the grid disks to bring them back into the disk group.

7. **Verify that all grid disks have been successfully put online using the following command:**

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
```

Wait until `asmmodestatus` is `ONLINE` or `UNUSED` for all grid disks. For example:

```
DATA_CD_00_dm01cel01          ONLINE
DATA_CD_01_dm01cel01          SYNCING
DATA_CD_02_dm01cel01          OFFLINE
DATA_CD_02_dm02cel01          OFFLINE
DATA_CD_02_dm03cel01          OFFLINE
DATA_CD_02_dm04cel01          OFFLINE
DATA_CD_02_dm05cel01          OFFLINE
DATA_CD_02_dm06cel01          OFFLINE
DATA_CD_02_dm07cel01          OFFLINE
DATA_CD_02_dm08cel01          OFFLINE
DATA_CD_02_dm09cel01          OFFLINE
DATA_CD_02_dm10cel01          OFFLINE
DATA_CD_02_dm11cel01          OFFLINE
```

Oracle ASM synchronization is complete only when all grid disks show `asmmodestatus=ONLINE` or `asmmodestatus=UNUSED`. Before taking another Exadata Storage

Server offline, Oracle ASM synchronization must complete on the restarted Exadata Storage Server. If synchronization is not complete, the check performed on another Exadata Storage Server will fail. For example:

```
CellCLI> list griddisk attributes name where asmdeactivationoutcome != 'Yes'
 DATA_CD_00_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_01_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_02_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_03_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_04_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_05_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_06_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_07_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_08_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_09_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_10_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
DATA_CD_11_dm01cel02    "Cannot de-activate due to other offline disks in the diskgroup"
```

▼ Drop an Exadata Storage Server

1. **From Oracle ASM, drop the Oracle ASM disks on the physical disk using the following command:**

   ```
   ALTER DISKGROUP diskgroup-name DROP DISK asm-disk-name
   ```

   To ensure correct redundancy level in Oracle ASM, wait for the rebalance to complete before proceeding.

2. **Remove the IP address entry from the `cellip.ora` file on each database server that accesses the Exadata Storage Server.**

3. **From Exadata Storage Server, drop the grid disks, cell disks, and cell on the physical disk using the following command:**

   ```
   DROP CELLDISK celldisk-on-this-lun FORCE
   ```

4. **Shut down all services on the Exadata Storage Server.**

5. **Power down the cell.**

Tip - Refer to "Shut Down an Exadata Storage Server" on page 111 for additional information.

Tuning SPARC SuperCluster

The `ssctuner` utility is a small set of Perl/Korn shell scripts and configuration files designed to run on Oracle Solaris 10/11 global zones. It runs as an SMF service to monitor and tune, in real time:

- `/etc/system`
- `ndd parameters`
- `/kernel/drv/sd.conf`
- `/kernel/drv/ibd.conf`

and periodically checks for other things like use of DISM or suboptimal NFS mount options. The utility runs every two hours and modifies these parameters as needed.

Note - If you manually tune a parameter `ssctuner` requires to have a different value, it will set the value of that parameter back to what `ssctuner` requires and log the fact that it made changes at this interval check.

The utility also checks, every two minutes, to see if there are any virtual disk devices in the degraded state that have come back online and will clear that zpool if so.

If you must control one or more of the parameters `ssctuner` manages, consider turning off those specific components rather than disabling `ssctuner` completely.

These topic describe how to:

- "Monitor `ssctuner`" on page 114
- "Turn Off `ssctuner`" on page 116
- "Disable Individual `ssctuner` Components" on page 116

Related Information

- For more information about SMF services on the Oracle Solaris OS, see the *Oracle Solaris System Administration Guide: Common System Management Tasks* at:

  `http://docs.oracle.com/cd/E23824_01/html/821-1451/hbrunlevels-25516.html#scrolltoc`

▼ Monitor `ssctuner`

- **View `ssctuner` activity.**

```
# svcs ssctuner
```

In addition, ssctuner writes messages to syslog and to the ssctuner service log. Those messages are tagged as ssctuner and might point to other file locations for more information.

To view the ssctuner service log, type:

```
# svcs -x ssctuner
svc:/site/application/sysadmin/ssctuner:default (ssctuner for SPARC SuperCluster)
State: online since September 28, 2012 07:30:15 AM PDT
See: ssctuner(l)
See: /var/svc/log/site-application-sysadmin-ssctuner:default.log
Impact: None.

# more /var/svc/log/site-application-sysadmin-ssctuner\:default.log
[ Sep 28 07:30:00 Disabled. ]
[ Sep 28 07:30:00 Rereading configuration. ]
[ Sep 28 07:30:10 Enabled. ]
[ Sep 28 07:30:10 Executing start method ("/opt/oracle.supercluster/ssctuner.ksh start"). ]
ssctuner local0.notice success: Saved rollback for : /etc/system
ssctuner local0.notice success: Saved ndd rollback.
ssctuner local0.notice success: Saved rollback for : /kernel/drv/sd.conf
ssctuner local0.notice success: enabled, version 0.99e. daemon PID= 14599
[ Sep 28 07:30:15 Method "start" exited with status 0. ]
ssctuner local0.notice success: daemon executing
ssctuner local0.notice success: Changes made to /etc/system
ssctuner local0.notice success: Changes made to /kernel/drv/sd.conf
```

To view ssctuner messages in /var/adm, type:

```
# grep -i ssctuner /var/adm/messages
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved rollback for : /etc/system
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved ndd rollback.
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved rollback for : /kernel/drv/sd.conf
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: enabled, version 0.99e. daemon PID= 14599
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: daemon executing
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Changes made to /etc/system
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Changes made to /kernel/drv/sd.conf
```

Related Information

- “Disable Individual ssctuner Components” on page 116
- “Turn Off ssctuner” on page 116

Turn Off ssctuner

Caution - Leave ssctuner on, unless you need to specifically control your ndd settings. Note that turning off ssctuner can cause performance problems on your system. Consider "Disable Individual ssctuner Components" on page 116 before turning off ssctuner completely.

Use this procedure to turn off the repeated checks, but keep the tuning changes.

- **Disable ssctuner.**

```
# svcadm disable ssctuner
```

Note - Rebooting after disabling ssctuner will lose the ndd tunings, but keep any other changes made by the utility.

Related Information

- "Disable Individual ssctuner Components" on page 116
- "Monitor ssctuner" on page 114

Disable Individual ssctuner Components

While it is discouraged, you can disable individual components that ssctuner controls. You must restart the SMF service for any changes to take effect.

1. **In /opt/oracle.supercluster/ssctuner/etc/ssctuner.conf, set the portion you cannot have enabled to FALSE.**

```
NDD_TUNE=TRUE
SYSTEM_TUNE=TRUE
SDCONF_TUNE=TRUE
ZPOOL_FIX=TRUE
NFS_CHECK=TRUE
DISM_CHECK=TRUE
```

2. **Restart the SMF service for changes to take effect.**

Related Information

- "Monitor ssctuner" on page 114

- “Turn Off `ssctuner`” on page 116

Configuring CPU and Memory Allocation

CPU/Memory Tool Overview

The SPARC SuperCluster T4-4 compute node CPU and memory resources are initially allocated during installation and defined by your configuration. To change how those resources are allocated, use the CPU/Memory tool (`setcoremem`).This tool enables you to change the CPU and memory allocation for the current domains only (you cannot modify the number of domains with this tool). Supported configurations for each domain must contain a minimum of 4 cores and 32 Gbytes of memory, and must be configured in multiples of 4 and 32 respectively.

You can allocate less than 32 cores and less than 1 TB of memory for a compute node. Any unused cores are considered *parked* cores and are not counted for licensing purposes. However, parked cores are unusable and essentially wasted. A better alternative in most situations is to assign cores to a different domain.

The CPU/Memory tool package ships on the system, but you must install the package to use it. For package installation instructions, see the “Installing a Package” section of the *Adding and Updating Oracle Solaris 11.1 Software Packages* document:

`http://docs.oracle.com/cd/E26502_01/html/E28984/gihhp.html#gikvp`

Related Information

- “Change CPU/Memory Allocations” on page 118
- “Supported CPU/Memory Allocations” on page 117
- “Revert to a Previous CPU/Memory Configuration” on page 123

Supported CPU/Memory Allocations

The following table describes the initial CPU, memory and root complex allocations. You can change the CPU/Memory allocation using the CPU/Memory tool, however, this tool does not change the root complex configuration.

TABLE 13

	Primary Domain	Domain 1	Domain 2	Domain 3	Root Complex
	% of cores and memory allocated				
Layout 0	1 core, 32 Gb memory	remaining cores and memory			2 in primary, 2 in domain 1
Layout 1	25%	75%			2 in primary, 2 in domain 1
Layout 2	50%	50%			2 in primary, 2 in domain 1
Layout 3	75%	25%			2 in primary, 2 in domain 1
Layout 4	25%	25%	50%		2 in primary, 1 in domain 1, 1 in domain 2
Layout 5	25%	50%	25%		2 in primary, 1 in domain 1, 1 in domain 2
Layout 6	50%	25%	25%		2 in primary, 1 in domain 1, 1 in domain 2
Layout 7	25%	25%	25%	25%	Each domain with 1 root complex
Layout 128	100%				4 in primary

Related Information

- “Change CPU/Memory Allocations” on page 118
- “Revert to a Previous CPU/Memory Configuration” on page 123
- “CPU/Memory Tool Overview” on page 117

▼ Change CPU/Memory Allocations

Perform this procedure on a compute node to change its CPU and memory resource allocation. Perform this task on each compute node you want to adjust.

The tool makes the changes in the following order:

- modify domain resources
- stop non-primary domains
- reboot primary domains with new resources
- reboot system
- bring up non-primary domains with new resources

1. **Log in as `root` to the control domain on the compute node.**

2. **Activate any inactive domains using the `ldm bind` command.**

   The tool will not continue if any inactive domains are present.

3. **Run `setcoremem` to enter the new CPU and memory allocations.**

Running setcorememem is a two step process. When you issue `setcorememem -s 0`, the system reports back the current CPU and memory allocations and allows you to change these allocations. Issuing `setcorememem -s 1` makes the changes to the system.

If you implement changes using the `setcorememem -s 0` command (Step 3), you must run the `setcorememem -s 1` command (Step 5), otherwise the system is left in an unstable state.

```
$ ssh root@etc8cn06-app1
Password:
Last login: Mon Jul 11 1:03:50 2012
Oracle Corporation      SunOS 5.11      11.0    May 2012
root@etc8cn06-app1:~# cd /opt/oracle.supercluster/bin
root@etc8cn06-app1:/opt/oracle.supercluster/bin/setcorememem/lib/bin# ./setcorememem -s 0
Info: A log of the activities is being recorded in /opt/oracle.supercluster/setcorememem/log/modlocality.log

A previous configuration that has not been fully applied already exist,
Do you want to reconfigure again? (y,n default: n) y

Info: The current or Next Power On SP-Config is layout7_ML10212012031137
Info: Found 4 Domains including primary
         DOMAIN                      CORES             Memory(GB)
         ======                      =====             ==========
         primary                         20                    928
s11-sru2a-dump-1                          4                     32
s11-sru2a-dump-2                          4                     32
s11-sru2a-dump-3                          4                     32

With this configuration adjustment tool, existing domains can be re-configured
with new CPU-core count and memory size.

The minimum number of Cores for a domain is 4. Adjustments are allowed
in multiple of 4 cores.
The minimum amount of Memory for a domain is 32 GB, Adjustments are allowed
in multiple of 32 GB.

Info: Parking cores (for license purposes)
      ===================================
      'Parked' cores are cores that are not assigned to any domain and are
      therefore not counted for license purposes. Parked cores are unusable
      and essentially wasted. A better alternative in most situations is to
      avoid unwanted license charges by assigning cores to a different domain.

  This resource adjustment tool cannot be used to modify the number of domains.

Do you want to reconfigure cpu and memory for all 4 domains? (y,n default: n) y
         DOMAIN              CORES-Cur/New         Mem(GB)-Cur/New
         ======              =============         ===============
         primary                      20/0                  928/0
s11-sru2a-dump-1                       4/0                   32/0
s11-sru2a-dump-2                       4/0                   32/0
s11-sru2a-dump-3                       4/0                   32/0
```

Note - The previous table shows the current CPU and memory settings and the new settings, which display as 0 because they have not yet been entered.

```
Enter 0 as the value for Core or Memory if you wish to abort
Valid Core Options are 0, 4, 8, 12, 16, 20
For [primary                 ] Enter the number of Cores between 4-20: 4
Valid Memory Options are 0, 32, 64, 96, 128, 160, 192, 224, 256, 288,
320, 352, 384, 416, 448, 480, 512, 544, 576, 608, 640, 672, 704, 736,
768, 800, 832, 864, 896, 928
For [primary                 ] Enter Memory size in GB between 32-928: 32
         DOMAIN              CORES-Cur/New         Mem(GB)-Cur/New
         ======              =============         ===============
         primary                      20/4                   928/32
s11-sru2a-dump-1                       4/0                    32/0
s11-sru2a-dump-2                       4/0                    32/0
s11-sru2a-dump-3                       4/0                    32/0
```

Note - To quit the CPU/Memory tool during this step, enter 0 for either core or memory value.

```
Enter 0 as the value for Core or Memory if you wish to abort
Valid Core Options are 0, 4, 8, 12, 16, 20
For [s11-sru2a-dump-1        ] Enter the number of Cores between 4-20: 8
Valid Memory Options are 0, 32, 64, 96, 128, 160, 192, 224, 256, 288,
320, 352, 384, 416, 448, 480, 512, 544, 576, 608, 640, 672, 704, 736,
768, 800, 832, 864, 896, 928
For [s11-sru2a-dump-1        ] Enter Memory size in GB between 32-928: 288
         DOMAIN              CORES-Cur/New         Mem(GB)-Cur/New
         ======              =============         ===============
         primary                      20/4                   928/32
s11-sru2a-dump-1                       4/8                    32/288
s11-sru2a-dump-2                       4/0                    32/0
s11-sru2a-dump-3                       4/0                    32/0


Enter 0 as the value for Core or Memory if you wish to abort

Valid Core Options are 0, 4, 8, 12, 16
For [s11-sru2a-dump-2        ] Enter the number of Cores between 4-16: 8

Valid Memory Options are 0, 32, 64, 96, 128, 160, 192, 224, 256, 288,
320, 352, 384, 416, 448, 480, 512, 544, 576, 608, 640, 672
For [s11-sru2a-dump-2        ] Enter Memory size in GB between 32-672: 288
         DOMAIN              CORES-Cur/New         Mem(GB)-Cur/New
         ======              =============         ===============
         primary                      20/4                   928/32
s11-sru2a-dump-1                       4/8                    32/288
```

```
s11-sru2a-dump-2                        4/8                     32/288
s11-sru2a-dump-3                        4/0                     32/0
```

Note - The tool displays current and the proposed CPU and memory values for the each domain as they are entered.

```
Enter 0 as the value for Core or Memory if you wish to abort

Valid Core Options are 0, 4, 8, 12, 16
For [s11-sru2a-dump-3 ] Enter the number of Cores between 4-16: 8

Valid Memory Options are 0, 32, 64, 96, 128, 160, 192, 224, 256, 288,
320, 352, 384, 416, 448, 480, 512, 544, 576, 608, 640, 672
For [s11-sru2a-dump-3 ] Enter Memory size in GB between 32-672: 288
         DOMAIN               CORES-Cur/New         Mem(GB)-Cur/New
         ======               =============         ===============
         primary                       20/4                   928/32
s11-sru2a-dump-1                        4/8                     32/288
s11-sru2a-dump-2                        4/8                     32/288
s11-sru2a-dump-3                        4/0                     32/288

[4 Cores] and [128 GB] of Memory will be parked

Do you wish to continue with these changes? (y,n default: n) y

Please Execute ./setcorememm -s 1 and follow the instructions there

Execution of above command will cause all domains in this system
to be halted and will require a reboot of this domain.
Please enusre that you have stopped all applications in all domains
on this system, before proceeding with the above command
root@etc8cn06-app1:/opt/oracle.supercluster/bin#
```

Note - This example shows parked resources - those resources that are unused by the current configuration.

4. **Perform a clean shut down of all applications, zones and domains (other than the primary).**

5. **Run `setcorememm -s 1` to make the configuration changes to the primary domain.**

Note - During this step, the tool forcefully closes any active domains (except the primary).

```
root@etc8cn06-app1:/opt/oracle.supercluster/bin# ./setcorememm -s 1 -v

Info: A log of the activities is being recorded in
/opt/oracle.supercluster/log/modlocality.log
```

```
Info: Executing step 1 - Please do not interrupt until this step is complete
The system has 1024 GB memory
Info: The current ldoms config is layout7_ML10212012031137
Info: It will remain preserved, and a new config will be added at end of step 2
Info: The name of the new config will be layout7_ML10222012031925
LDom s11-sru2a-dump-3 stopped
Warning: Resulting memory size is zero.
Domain s11-sru2a-dump-3 uses physically bound core resources
Warning: setting required number of cores to zero
LDom s11-sru2a-dump-2 stopped
Warning: Resulting memory size is zero.
Domain s11-sru2a-dump-2 uses physically bound core resources
Warning: setting required number of cores to zero
LDom s11-sru2a-dump-1 stopped
Warning: Resulting memory size is zero.
Domain s11-sru2a-dump-1 uses physically bound core resources
Warning: setting required number of cores to zero
Initiating a delayed reconfiguration operation on the primary domain.
All configuration changes for other domains are disabled until the primary
domain reboots, at which time the new configuration for the primary domain
will also take effect.
--------------------------------------------------------------------------------
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
--------------------------------------------------------------------------------
Info: Adding 4 Cores to Primary Domain...
--------------------------------------------------------------------------------
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
--------------------------------------------------------------------------------
Info: Done with changes

Info: Done with step 1 - system reboot is required
Info: After reboot, Please run setcorememem -s 2
root@etc8cn06-app1:/opt/oracle.supercluster/bin#
```

6. **Reboot the system.**

7. **Log back in to the compute node and run `setcorememem -s 2` to make the configuration changes to the guest domain(s).**

```
$ ssh root@etc8cn04-app1
Password:
Last login: Mon Aug 23 10:44:50 2012
```

```
Oracle Corporation      SunOS 5.11      11.0    May 2012
You have new mail.
root@ssc6cn06:~# cd /opt/oracle.supercluster/bin
root@etc8cn06-app1:/opt/oracle.supercluster/bin# ./setcorememm -s 2 -v
Info: A log of the activities is being recorded in
/opt/oracle.supercluster/log/modlocality.log
Info: Executing step 2 - Please do not interrupt until this step is complete
Info: Setting CORE Constraints for Domain s11-sru2a-dump-1
Info: Setting COREs to 1,2,9,10,17,18,25,26 for Domain s11-sru2a-dump-1
Info: Setting Memory for Domain s11-sru2a-dump-1
Info: Binding Domain s11-sru2a-dump-1
Info: Starting Domain s11-sru2a-dump-1
LDom s11-sru2a-dump-1 started
Info: Setting CORE Constraints for Domain s11-sru2a-dump-2
Info: Setting COREs to 3,4,11,12,19,20,27,28 for Domain s11-sru2a-dump-2
Info: Setting Memory for Domain s11-sru2a-dump-2
Info: Binding Domain s11-sru2a-dump-2
Cannot validate NIC net9 on service domain s11-sru2a-dump-2
Cannot validate path /dev/rdsk/c0t5000C500436FEBB7d0s1 on service domain
s11-sru2a-dump-1
Info: Starting Domain s11-sru2a-dump-2
LDom s11-sru2a-dump-2 started
Info: Setting CORE Constraints for Domain s11-sru2a-dump-3
Info: Setting COREs to 5,6,13,14,21,22,29,30 for Domain s11-sru2a-dump-3
Info: Setting Memory for Domain s11-sru2a-dump-3
LDom s11-sru2a-dump-3 started
Info: The new SP Config is layout7_ML10222012031925
Info: Done with changes
Info: CPU and Memory configuration successfully changed
root@etc8cn06-app1:/opt/oracle.supercluster/bin#
```

Related Information

- “Revert to a Previous CPU/Memory Configuration” on page 123
- “Supported CPU/Memory Allocations” on page 117
- “CPU/Memory Tool Overview” on page 117

▼ Revert to a Previous CPU/Memory Configuration

Use this procedure to revert a compute node to a previous CPU/Memory configuration. You must perform this procedure on each member in a cluster; the tool does not automatically propagate changes to every cluster member.

Caution - All members in a cluster must use the same CPU/memory configuration or the SPARC SuperCluster T4-4 may become unstable.

1. **List previous configurations.**

You can also view previous configurations in the log files at this location:

```
/opt/oracle.supercluster/setcorememl/log
```

```
$ ssh root@hostname-app1
Password:
Last login: Mon Jul 11 1:53:50 2012
Oracle Corporation      SunOS 5.11      11.0    May 2012
# ldm list-confg
factory-default
layout6
layout6_ML10182012142558
layout6_ML10182012162049
layout6_ML10182012165059
layout6_ML10182012201150 [current]
```

In this example, `layout6` is the configuration initially installed. The remaining configurations are customer-created and can be identified by the naming format; ML stands for Modified Layout and the number string represents the local date and time that the configuration was created (MMDDYYYYHHMISS). So `layout6_ML10182012201150` is a customer-modified layout created on October 18, 2012 at 8:11:50 pm.

2. **Revert to a previous configuration.**

```
# ldm set-config layout6
```

3. **Halt all domains, then halt the primary domain.**

4. **Restart the system from the service processor.**

```
# #.

-> cd /SP
-> stop /SYS
Are you sure you want to stop /SYS (y/n) ? y
Stopping /SYS

-> start /SYS
Are you sure you want to start /SYS (y/n) ? y
Starting /SYS
```

5. **Boot all domains and zones.**

Related Information

- “CPU/Memory Tool Overview” on page 117
- “Change CPU/Memory Allocations” on page 118

- “Supported CPU/Memory Allocations” on page 117

Enabling and Disabling Capacity-on-Demand on SPARC SuperCluster

On a SPARC SuperCluster T4-4 that offers Capacity-on-Demand, a number of the cores on the compute servers can be enabled or disabled logically using the `ssccod` utility. Using the utility, a system can increase compute server CPU processing power when needed and then decrease it when the power is not needed. For example, a business may run 24 of the 32 cores (75%) in each SPARC SuperCluster T4-4 compute server the majority of time, and then enable the other 8 cores during the end of the quarter for running reports and queries. After the quarter-end processing is done, the 8 cores can be disabled. It is not necessary to restart the servers after running the script.

Capacity-on-Demand can be managed with the following script:

```
ssccod { -capacity_on_demand | -cod } { display | enable | disable }
```

This utility needs to be run in the control domain of each SPARC SuperCluster T4-4 compute node. To establish which nodes are your control domains, use the `virtinfo` shell command, which will specify "LDoms control" if you are in the control domain, e.g.:

```
$ /usr/sbin/virtinfo
Domain role: LDoms control I/O service root
```

The following table describes the parameters used with the `ssccod` utility.

Parameter	Description
`display`	Displays the number of CPUs in use and spare in each domain.
`enable`	Enables the spare CPUs reserved for additional on-demand capacity.
`disable`	Disables the spare CPUs.

Note - If the customer is using Solaris CPU resource controls (`pools` or `psets`) or processor binding inside any domain, these controls may need to be adjusted after Capacity-on-Demand is enabled, to allow existing workloads to utilize the newly-on-lined cores. This task is the responsibility of the customer. Similarly, the Capacity-on-Demand cores may need to be removed from resource pools or psets prior to Capacity-on-Demand being disabled. This task is also the responsibility of the customer. If existing CPU resource controls prevent Capacity-on-Demand from being disabled, the `ssccod` utility will inform you, and leave all spare cores enabled. This may incur additional Capacity-on-Demand operational charges. See the manual pages for `pooladm`(1M), `poolcfg`(1M), `psrset`(1M) and `pbind`(1M) for reference on Solaris CPU resource controls.

Monitoring the System

These topics describe the monitoring options for the SPARC SuperCluster T4-4.

Monitoring the System Using Auto Service Request

These topics describe how to monitor the SPARC SuperCluster T4-4 using Oracle Auto Service Request (ASR).

ASR Overview

Auto Service Request (ASR) is designed to automatically open service requests when specific SPARC SuperCluster T4-4 hardware faults occur. To enable this feature, the SPARC SuperCluster T4-4 components must be configured to send hardware fault telemetry to the ASR Manager software. The ASR Manager must be installed on a server that has connectivity to SPARC SuperCluster T4-4, and an outbound Internet connection using HTTPS or an HTTPS proxy.

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

Prior to using ASR, set up the following:

- Oracle Premier Support for Systems or Oracle/Sun Limited Warranty
- Technical contact responsible for SPARC SuperCluster T4-4
- Valid shipping address for SPARC SuperCluster T4-4 parts

An e-mail message is sent to both the My Oracle Support (MOS) e-mail account for Auto Service Request and the technical contact for the activated asset, notifying them of the creation of the service request.

Note - If a subscriber has not been set up, then the subsequent Auto Service Request activation will fail.

Consider the following information when using ASR:

- 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. For more information, see:

  `http://www.oracle.com/asr`

  Click the `Documentation` link on this page, then refer to the “ASR Fault Coverage Information” section at the bottom of the page.
- ASR is not a replacement for other monitoring mechanisms, such as SMTP, and SNMP alerts, within the customer data center. It is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.
- There are occasions when a service request may not be automatically filed. This can happen because of the unreliable nature of the SNMP protocol, or loss of connectivity to the ASR Manager. Oracle recommends that customers continue to monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been automatically filed.

Tip - Refer to the Oracle Auto Service Request web page at `http://www.oracle.com/asr` for more information on ASR.

Preparing to Configure ASR

Confirm your environment is supported and prepared before installing and configuring ASR on the SPARC SuperCluster T4-4:

▼ Prepare the ASR Environment

Before You Begin Before installing ASR, ensure the following conditions are met:

1. **Create a My Oracle Support (MOS) account at `http://support.oracle.com`.**

   Ensure the following are correctly set up:

   - Oracle Premier Support for Systems or Oracle/Sun Limited Warranty
   - Technical contact responsible for SPARC SuperCluster T4-4
   - Valid shipping address for SPARC SuperCluster T4-4 parts

2. **Identify and designate a system to serve as ASR Manager.**

   For more information, see:

   `http://www.oracle.com/asr`

   Click `additional details`, then click `Hardware and Network Configuration Recommendations`.

3. **Identify and verify ASR assets.**

4. **Ensure connectivity to the Internet using HTTPS.**

   You may need to open certain ports to your datacenter. For more information, see the *Oracle ASR Security Whitepaper*, located here:

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

   Click on the `Oracle ASR` user documentation link.

5. **Provide the necessary information in the following documents:**

   - *SPARC SuperCluster T4-4 Site Checklists*
   - *SPARC SuperCluster T4-4 Configuration Worksheets*

Software Requirements: ASR Manager

You will need root access to install the software to set up the ASR Manager.

- ASR Manager, version 3.6 or higher
- Oracle Services Tool Bundle (STB) for Solaris only

Software Requirements: SPARC SuperCluster T4-4

You will need root access to install the software to set up the ASR Manager.

- Operating System: Oracle Linux 5.3 and later, or Oracle Solaris 10 Update 10 (10u10) plus patches, and later
- Java Version: at least JRE/JDK 6.2.28
- Database Server: Exadata Software 11.2.3.1 and higher
- Exadata Storage Server: Exadata Storage Server Software 11.2.0.3 DB with Bundle Patch 1, and higher
- Sun ZFS Storage 7320 Storage Controller: Firmware version AK2011.1.0 and higher

Installing ASR Manager Components

Perform these procedures to install the ASR Manager components on the external system designated as the ASR Manager. You may use a pre-existing ASR Manager, as long as it conforms to the requirements listed in "Prepare the ASR Environment" on page 129.

▼ Install the Oracle Automated Service Manager (OASM) Package

1. **Verify that you have version 1.3.1 or later (if needed, download OASM).**

   As root:

   - Oracle Solaris: `pkginfo -l SUNWsasm`
   - Oracle Linux: `rpm -q SUNWsasm`

2. **Install the OASM package.**

   As root:

   - Oracle Solaris: `pkgadd -d SUNWsasm.`*version-number*`.pkg`
   - Oracle Linux: `rpm -i SUNWsasm.`*version-number*`.rpm`

▼ Install Service Tools Bundle (STB) for Solaris Only

1. **If needed, download Services Tools Bundle from:**

http://www.oracle.com/asr

and click on the Download link.

2. **Untar the STB bundle and run the installation script (`install_stb.sh`).**
   As part of the installation, select:

   - Type I for "install"
   - Type Y to replace existing SNEEP packages
   - Type Y to replace existing Service Tags packages

   Note - See Doc ID 1153444.1 to download the latest STB bundle from My Oracle Support (log in required): `https://support.oracle.com`

3. **Confirm that SNEEP is installed correctly:**

   ```
   sneep -a
   ```

4. **Verify that Service Tags is reporting your system attributes correctly:**

   ```
   stclient -E
   ```

   If the serial number does not display, then register the serial number manually:

   ```
   sneep -s serial-number
   ```

▼ Install the Oracle Auto Service Request (ASR) Package

1. **Download and unzip the ASR package.**
   As root:

   - Oracle Solaris: `pkgadd -d SUNWswasr.`*version-number*`.pkg`
   - Oracle Linux: `rpm -i SUNWswasr.`*version-number*`.rpm`

2. **Add the `asr` command to the `PATH` (update to the root's `.profile`, `.cshrc` , `.kshrc` or `.bashrc` as needed):**

   ```
   PATH=$PATH:/opt/SUNWswasr/bin
   export PATH
   ```

▼ Register ASR Manager

Before You Begin When registering ASR Manager, you will be asked to type your MOS single sign on information and any proxy servers, if needed.

1. **As root on the ASR Manager system, type:**

   ```
   # asr register
   ```

2. **Type "1" or "alternate URL for Managed OPS use only":**

   ```
   1) transport.oracle.com
   ```

3. **If you are using a proxy server to access the Internet, type the proxy server information.**

   Your screen output should look like this:

   ```
   Proxy server name: ? <proxy server name>
   Proxy port number: ? <proxy port number>
   Proxy authentication; if authentication is not required, enter -.
   Proxy user: <proxy user name>
   Proxy password: <proxy password>
   If this is an NTLM type proxy, enter the information below.
   Otherwise, enter -
   NTLM Domain: [?] <NTLM domain name>
   Enter the host the NTLM authentication request is originating
   from. Usually this is the hostname of the SASM server.
   NTLM Host: [?] <NTLM host name>
   NTLM support is not bundled with SASM but may be added now.

   1) Download jCIFS from http://jcifs.samba.org/
   2) Extract contents and locate the jcifs-*.jar file
   3) Enter full path to this file below

   jCIFS jar file: [?] <full path of jCIFS jar file>
   Note: These properties are stored in the
   /var/opt/SUNWsasm/configuration/config.ini file. You can update
   these properties if needed and then restart SASM.
   ```

4. **When prompted, type your My Oracle Support (MOS) username and password. ASR will validate the login. Once validated, the registration is complete. Note: Passwords are not stored.**

   Your MOS e-mail address receives output from notification of ASR problems and Service Request (SR) generation.

▼ Verify ASR Manager

1. **On the ASR Manager, verify that you have the correct version of ASR Manager:**

   ```
   # asr show_rules_version
   ```

   You should see that the version is 3.6 or later.

2. **Check the registration status:**

```
# asr show_reg_status
```

3. **Test the connection to ensure that ASR can send information to the transport server:**

```
# asr test_connection
```

Configure SNMP Trap Destinations for the Exadata Storage Servers

Note - Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

Complete the following steps on each Exadata Storage Server:

1. **Log in as `celladmin` on the Exadata Storage Server.**

2. **On the Exadata Storage Server, add SNMP trap destinations:**

```
# cellcli -e "alter cell snmpSubscriber=(host ='ASR-Manager-name-or-IP-address',port=162,
community=public,type=asr)"
```

Note that single quotes are required around the *ASR-Manager-name-or-IP-address* entry. Following are the element definitions for the command above:

- host=*'ASR-Manager-name-or-IP-address'* – The ASR Manager hostname can be used when DNS is enable for the site. If DNS is not running, the IP address is preferred, but the ASR Manager hostname can be used if the entry is added to the `/etc/hosts` file.
- type=asr – Shows the ASR Manager as being a special type of SNMP subscriber.
- community=public – The required value of the community string. This value can be modified to be a different string based on customer network requirements.
- port=162 – The SNMP port. This port value is customer dependant. It can be configured as a different port based on network requirements, or it may need to be changed for ASR to work correctly in a managed environment.

3. **Validate if Oracle ILOM auto-activation occurred (if the network and Oracle ILOM are set up correctly):**

```
# asr list_asset
```

Following is example output:

```
IP_ADDRESS     HOST_NAME       SERIAL_NUMBER   ASR       PROTOCOL   SOURCE
----------     ---------       -------------   ---       --------   -----
```

```
10.60.40.105   ssc1cel01       1234FMM0CA      Enabled   SNMP       ILOM
10.60.40.106   ssc1cel02       1235FMM0CA      Enabled   SNMP       ILOM
10.60.40.107   ssc1cel03       1236FMM0CA      Enabled   SNMP       ILOM
10.60.40.117   ssc1cel01-ilom  1234FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
10.60.40.118   ssc1cel02-ilom  1235FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
10.60.40.119   ssc1cel03-ilom  1236FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
```

- If all Oracle ILOMs for the Exadata Storage Servers are in the list, go to Step 5.
- If Oracle ILOMs are not in the list, go to Step 4.

4. **On the ASR Manager, activate the Oracle ILOMs of the Exadata Storage Servers:**

   # asr activate_asset -i *ILOM-IP-address*

   or

   # asr activate_asset -h *ILOM-hostname*

   Note - If the last step fails, verify that port 6481 on the Oracle ILOM is open. If port 6481 is open and the step still fails, contact ASR Support.

5. **Activate the Exadata OS side of the ASR support:**

   # asr activate_exadata -i *host-management-IP-address* **-h** *host-management-hostname* **-l** *ILOM-IP-address*

   or

   # asr activate_exadata -i *host-management-IP-address* **-h** *host-management-hostname* **-n** *ILOM-hostname*

6. **Validate all Exadata Storage Servers are visible on the ASR Manager:**

   # asr list_asset

   You should see both the Oracle ILOM and the host referenced in the list, with the same serial number, as shown in the following example output:

```
IP_ADDRESS     HOST_NAME       SERIAL_NUMBER   ASR       PROTOCOL   SOURCE
----------     ---------       -------------   ---       --------   -----
10.60.40.105   ssc1cel01       1234FMM0CA      Enabled   SNMP       ILOM
10.60.40.106   ssc1cel02       1235FMM0CA      Enabled   SNMP       ILOM
10.60.40.107   ssc1cel03       1236FMM0CA      Enabled   SNMP       ILOM
10.60.40.117   ssc1cel01-ilom  1234FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
10.60.40.118   ssc1cel02-ilom  1235FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
10.60.40.119   ssc1cel03-ilom  1236FMM0CA      Enabled   SNMP,HTTP  EXADATA-SW
```

7. **On the Exadata Storage Server, validate the configuration:**

   # cellcli -e “list cell attributes snmpsubscriber”

8. **On the Exadata Storage Server, validate the SNMP configuration:**

```
# cellcli -e “alter cell validate snmp type=asr”
```

The MOS contact will receive an email as confirmation.

9. **Repeat these instructions for every Exadata Storage Server in your SPARC SuperCluster T4-4.**

10. **When you have completed these instructions for every Exadata Storage Server in your SPARC SuperCluster T4-4, approve and verify contacts to the Exadata Storage Servers on MOS.**

    See “Approve and Verify ASR Activation for SPARC SuperCluster T4-4 Assets” on page 143 for those instructions.

    For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

▼ Configure ASR on the Sun ZFS Storage 7320 Storage Appliance

To activate the Sun ZFS Storage 7320 appliance included in your SPARC SuperCluster T4-4, complete the following steps on each Sun ZFS Storage 7320 storage controller:

1. **In a web browser, type the IP address or host name you assigned to the host management port of either Sun ZFS Storage 7320 storage controller as follows:**

   `https://`*storage-controller-ipaddress*`:215`

   or

   `https://`*storage-controller-hostname*`:215`

   The login screen appears.

2. **Type `root` into the `Username` field and the root password into this login screen, and press the `Enter` key.**

3. **Click the Configuration tab, and click SERVICES, and then on the left navigation pane, click Services to display the list of services.**

4. **Scroll down in the screen and click Phone Home, as shown in the following figure.**

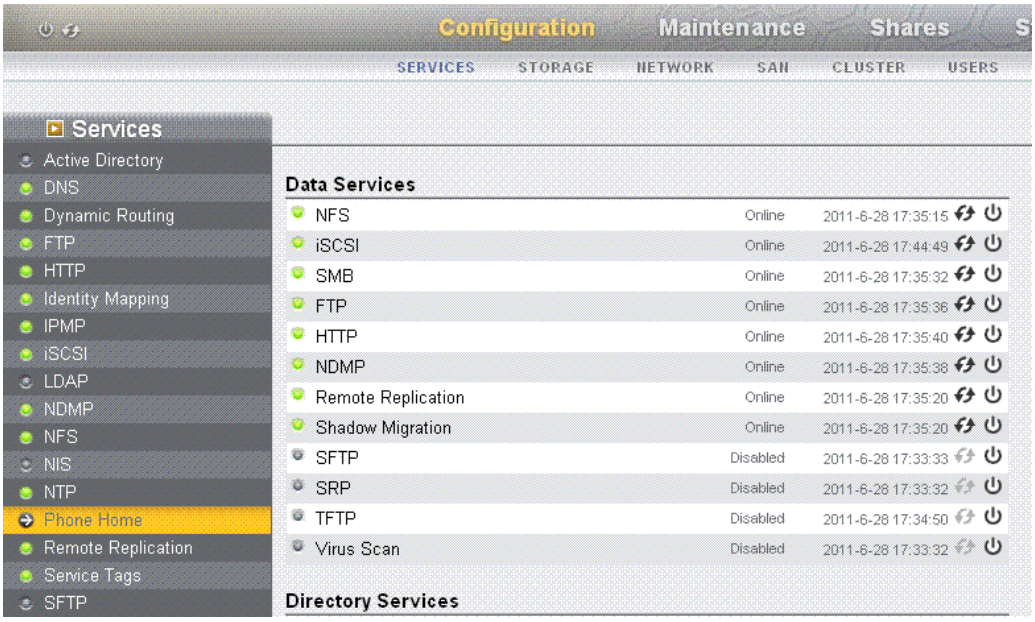

When you click Phone Home, the Phone Home page is displayed, as shown in the following figure.

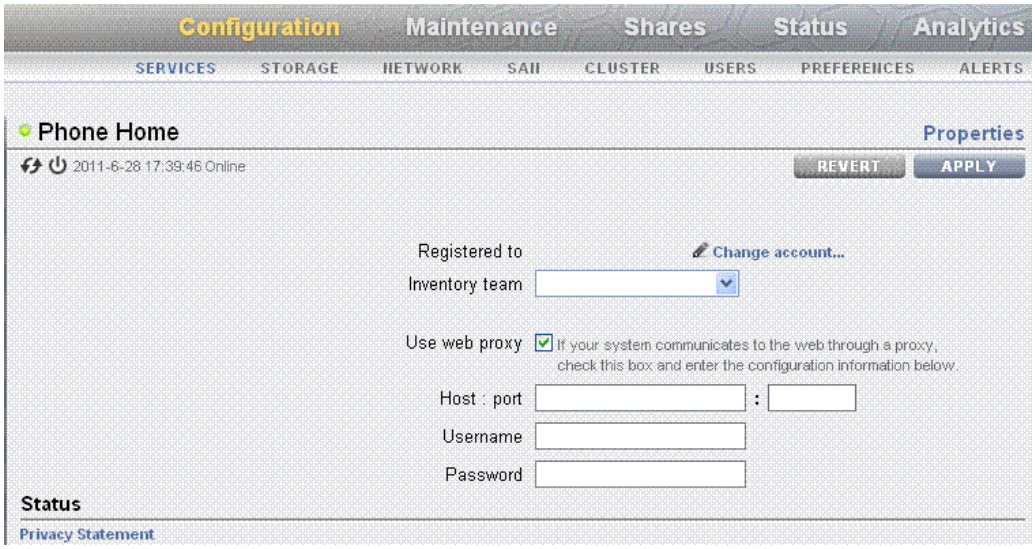

5. **If you are using a web proxy to connect to the Internet from the storage appliance, select the Use web proxy option, and type the following information:**

   - In the *Host:port* field, type the complete host name of your web proxy server and the port.
   - In the *Username* field, type your user name for the accessing the web proxy server.
   - In the *Password* field, type the password.

6. **Click the pencil icon in the registration section.**

   A Privacy Statement is displayed. Click OK, complete the section for My Oracle Support and password, and click OK.

7. **When the account is verified, select the Sun Inventory and Enable Phone Home options.**

8. **After typing the information, click APPLY.**

9. **When the Service Enable / Disable popup is presented, select the Enable option.**

10. **Repeat these instructions for every Sun ZFS Storage 7320 appliance in your SPARC SuperCluster T4-4.**

11. **When you have completed these instructions for every Sun ZFS Storage 7320 appliance in your SPARC SuperCluster T4-4, approve and verify contacts to the Sun ZFS Storage 7320 appliances on MOS.**

    See "Approve and Verify ASR Activation for SPARC SuperCluster T4-4 Assets" on page 143 for those instructions.

    For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

▼ Configure ASR on the SPARC T4-4 Servers: Oracle ILOM

Note - Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

To configure the Oracle ILOM for SPARC T4-4 servers, complete the following steps on each SPARC T4-4 server:

1. **Log in to the SPARC T4-4 server Oracle ILOM.**

2. **Display the available rules:**

   ```
   # show /SP/alertmgmt/rules
   ```

   This lists the rules available, similar to the following:

   ```
   1
   2
   3
   ...
   15
   ```

3. **Pick one of the rules and type the following command to determine if that rule is currently being used:**

   ```
   # show /SP/alertmgmt/rules/rule-number
   ```

   For example:

   ```
   # show /SP/alertmgmt/rules/1
   ```

   - If you see output similar to the following:

     ```
     Properties:
     ```

```
type = snmptrap
level = minor
destination = 10.60.10.243
destination_port = 0
community_or_username = public
snmp_version = 2c
testrule = (Cannot show property)
```

this rule is currently being used and should not be used for this exercise (the destination address shown would be the IP address of the ASR Manager in this case). If you see output similar to the preceding example, pick another rule and type the `show /SP/alertmgmt/rules/`*rule-number* command again, this time using another rule in the list.

- If you see output similar to the following:

```
Properties:
type = snmptrap
level = disable
destination = 0.0.0.0
destination_port = 0
community_or_username = public
snmp_version = 1
testrule = (Cannot show property)
```

this rule is currently unused and can be used for this exercise.

4. **Type this command using the unused rule:**

   # **`set /SP/alertmgmt/rules/`***unused-rule-number* **`type=snmptrap level=minor destination=`***IP-address-of-ASR-Manager* **`snmp_version=2c community_or_username=public`**

5. **Log in to the ASR Manager server.**

6. **Activate Oracle ILOM for the SPARC T4-4 server:**

   `asr>` **`activate_asset -i`** *ILOM-IP-address*

7. **Repeat these instructions on Oracle ILOM for every SPARC T4-4 server in your SPARC SuperCluster T4-4.**

8. **When you have completed these instructions for every SPARC T4-4 server in the SPARC SuperCluster T4-4, approve and verify contacts to the SPARC T4-4 servers on MOS.**

See “Approve and Verify ASR Activation for SPARC SuperCluster T4-4 Assets” on page 143 for those instructions.

For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

Configuring ASR on the SPARC T4-4 Servers: Oracle Solaris 11

Note - Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

Oracle Solaris 11 includes the ability to send ASR fault events and telemetry to Oracle using xml over HTTP to the ASR Manager.

To enable this capability use the `asr enable_http_receiver` command on the ASR Manager. Select a port for the HTTP receiver that is appropriate for your network environment and does not conflict with other network services.

Perform the following tasks:

- “Enable the HTTP Receiver on the ASR Manager” on page 140
- “Enable HTTPS on the ASR Manager (Optional)” on page 141
- “Register the SPARC T4-4 Server with Oracle Solaris 11 or Database Domains to ASR Manager” on page 142

▼ Enable the HTTP Receiver on the ASR Manager

Follow this procedure on the ASR Manager to enable the HTTP receiver for Oracle Solaris 11 ASR Assets.

1. **Log in to the ASR Manager system as `root`.**

2. **Verify the existing settings:**

   ```
   # asr show_http_receiver
   ```

3. **Enable the HTTP receiver:**

   ```
   # asr enable_http_receiver -p port-number
   ```

   where *port-number* is the port that you are designating for HTTP traffic.

Note - If you need to disable the HTTP receiver, run `asr disable_http_receiver`.

4. **Verify the updated configuration:**

```
# asr show_http_receiver
```

5. **Verify the HTTP receiver is up and running.**

In a browser, go to: `http://`*ASR-Manager-name*`:`*port-number*`/asr`

A message will display indicating that the HTTP receiver is up and running.

▼ Enable HTTPS on the ASR Manager (Optional)

If you need to use HTTPS for security purposes, you can set up HTTPS/SSL for the ASR Manager HTTP receiver.

1. **Once the SSL certificate from a trusted authority is loaded into keystore, then add the following SSL connector in `/var/opt/SUNWsasm/configuration/jetty/jetty.xml` below the `<Call name="addConnector">` sections:**

```
<Call name=”addConnector”>
  <Arg>
    <New class=”org.mortbay.jetty.security.SslSocketConnector”>
      <Set name=”Port”>443</Set>
      <Set name=”maxIdleTime”>30000</Set>
      <Set name=”keystore”>path-to-keystore</Set>
      <Set name=”password”>password</Set>
      <Set name=”keyPassword”key-password</Set>
      <Set name=”truststore”>path-to-keystore</Set>
      <Set name=”trustPassword”trust-password</Set>
    </New>
  </Arg>
</Call>
```

Passwords above can be plain text or obfuscated as follows:

`java -classpath lib/jetty-6.1.7.jar:lib/jetty-util-6.1.7.jar`

`org.mortbay.jetty.security.Password` *plaintext-password*

Then copy and paste the output line starting with `OBF:` (including the `OBF:` part) into this `jetty.xml` config file.

2. **Restart OASM.**

- On a system running Oracle Solaris, type:

```
# svcadm restart sasm
```

- On a system running Oracle Linux, type:

  # /opt/SUNWsasm/bin/sasm stop-instance

  # /opt/SUNWsasm/bin/sasm start-instance

3. **Verify the SSL setup by accessing the following URL from a browser:**

   `https://`*ASR-Manager-name*`/asr`

▼ Register the SPARC T4-4 Server with Oracle Solaris 11 or Database Domains to ASR Manager

Follow this procedure to register the SPARC T4-4 server with Oracle Solaris 11 or Database Domains to the ASR Manager.

1. **Log in to the SPARC T4-4 server as `root`.**

2. **Confirm that the `asr-notify` service is working:**

   # svcs asr-notify

   - If you see the following message:

     `svcs: Pattern ???asr-notify' doesn't match any instances`

     then confirm that the `asr-notify` service is installed:

     # pkg list asr-notify

     If you see the following message:

     `pkg list: no packages matching ???asr-modify' installed`

     then install the `asr-notify` service:

     # pkg install system/fault-management/asr-notify

     Enter the `svcs asr-notify` command again to confirm that the `asr-notify` service is working.
   - If you see the following message:

     ```
     # svcs asr-notify
     STATE     STIME       FMRI
     online    16:06:05    svc:/system/fm/asr-notify:default
     ```

     then the `asr-notify` service is installed and is working properly

3. **To register the ASR manager, run:**

   # asradm register -e http://*asr-manager-host***:***port-number***/asr**

For example:

```
# asradm register -e http://asrmanager1.mycompany.com:8777/asr
```

You should see screens asking for your Oracle Support account name and password. After entering your Oracle Support account name and password, you should see a notification, saying that your registration is complete:

```
Enter Oracle SSO User Name:
Enter password:

Registration complete.
```

4. **Run the following command:**

```
# asradm list
```

The screen output should be similar to the following:

```
PROPERTY VALUE
Status Successfully Registered with ASR manager
System Id system-identification-number
Asset Id asset-identification-number
User username
Endpoint URL http://asr-manager-host:port-number/asr
```

Upon successful results of the above commands, the registration of the ASR Manager is complete.

5. **Repeat these instructions for every SPARC T4-4 server with Oracle Solaris 11 or Database Domains in your SPARC SuperCluster T4-4.**

6. **When you have completed these instructions for every SPARC T4-4 server in your SPARC SuperCluster T4-4, approve and verify contacts to the SPARC T4-4 servers on MOS. See "Approve and Verify ASR Activation for SPARC SuperCluster T4-4 Assets" on page 143 for those instructions.**

For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

▼ Approve and Verify ASR Activation for SPARC SuperCluster T4-4 Assets

1. **On the standalone system where ASR Manager is running, run the following command to verify the status of your SPARC SuperCluster T4-4 assets:**

```
list_asset
```

This command should list ASR assets in your SPARC SuperCluster T4-4, including SPARC T4-4 servers, Exadata Storage Servers, and Sun ZFS Storage 7320 storage controllers.

2. **Log in to My Oracle Support (`https://support.oracle.com`).**

3. **In the My Oracle Support Dashboard, click the `More...` tab, then click `Settings` from the menu.**

4. **In the `Settings` pane on the left of the window, select `Pending ASR Activations` (located under the Administrative sub menu).**

   A complete list of all qualified ASR assets that are awaiting approval are displayed.

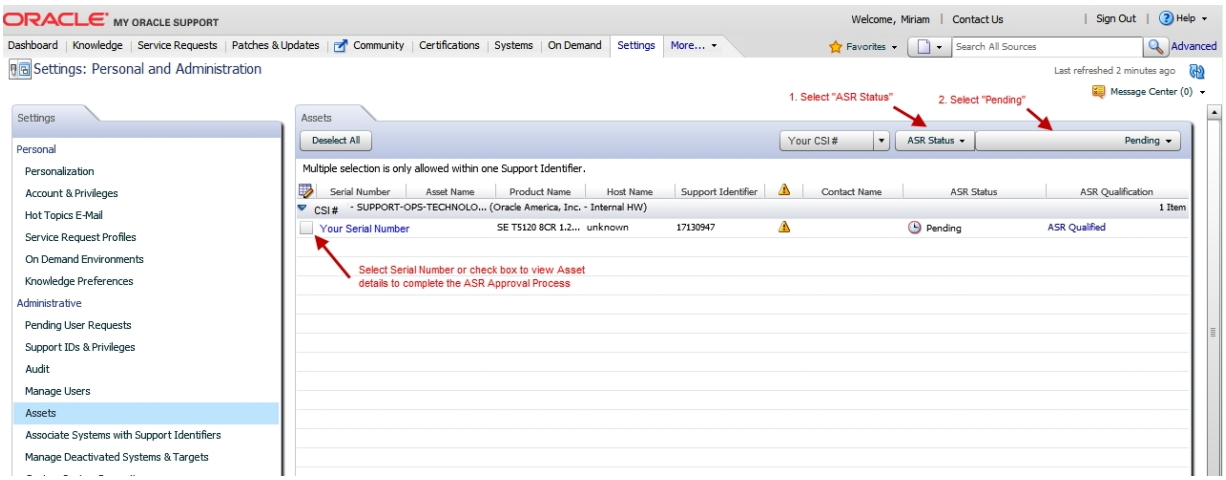

Note - By default, all support identifiers that you are associated with are displayed. If this list of assets is long, you can limit the display to show only assets associated to one support identifier. You can also search for an asset's serial number.

Note - For each component in the SPARC SuperCluster T4-4, you should see two host names associated with each serial number. If you see only the Oracle ILOM host name, that means that you did not activate ASR for that component. If you see more than two host names associated with each serial number, you might need to request help for ASR. To do this, open a hardware SR with “Problem Category” set to “My - Auto Service Request (ASR) Installation and Configuration Issues.”

---

5. **Click the asset's serial number.**

If any missing information about the asset is required, the information pop-up will indicate the needed information. The ASR Activation window will appear and look like the following figure.

Note - ASR Host name is updated when an activation request is sent to Oracle from the ASR software on the asset. (For example, from the `asr activate_asset` command on the ASR Manager.)

Required fields for ASR asset activation are:

- Contact Name: You can only select a name associated with the support identifier. Click the drop-down menu to see the list of available names.

  A contact must have the "Create SR" privilege for the asset's support identifier.
- Street Address 1: Type the street address for the asset.

  Note - By default, all support identifiers that you are associated with are displayed. If this list of assets is long, you can limit the display to show only assets associated to one support identifier. You can also search for an asset's serial number.
- Country: Select the asset's country location from the drop-down menu.
- ZIP/Postal Code: type the ZIP/postal code for the asset's location. If there is no postcode insert "-".
- Distribution Email List: You can add email addresses that will receive all ASR mail notifications. Separate multiple emails addresses with a comma. For example:

  asr-notifications-1@mycompany.com, asr-notifications-2@mycompany.com

ASR will send email to the Contact's email address and the Distribution Email List, if provided. This is a useful feature if your organization has a team that should be informed about Service Requests created by ASR.

6. **Click the “Approve” button to complete the ASR activation.**

Note - A SPARC SuperCluster T4-4 asset must be in an active ASR state in My Oracle Support in order for Service Request autocreate to work.

---

7. **To confirm that ASR can send information to the transport server, run:**

```
# asradm send test email-address@company.com
```

This command sends a test alert e-mail to the e-mail address.

Monitoring the System Using OCM

These topics describe how to monitor the system with OCM.

- “OCM Overview” on page 146
- “Install Oracle Configuration Manager on the SPARC T4-4 Servers” on page 147

OCM Overview

Oracle Configuration Manager (OCM) collects configuration information and uploads it to the Oracle repository. When the configuration information is uploaded daily, Oracle Support Services can analyze the data and provide better service. When a service request is logged, the configuration data is associated with the service request. The following are some of the benefits of Oracle Configuration Manager:

- Reduced time for problem resolution
- Proactive problem avoidance
- Improved access to best practices, and the Oracle knowledge base
- Improved understanding of the customer's business needs
- Consistent responses and services

The Oracle Configuration Manager software is installed and configured in each `ORACLE_HOME` directory on a host. For clustered databases, only one instance is configured for Oracle Configuration Manager. A configuration script is run on every database on the host. The Oracle Configuration Manager collects and then sends the data to a centralized Oracle repository.

For more information, see:

- *Oracle Configuration Manager Installation and Administration Guide*
- *Oracle Configuration Manager Collection Overview*

▼ Install Oracle Configuration Manager on the SPARC T4-4 Servers

Note - Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

1. Log in to the logical domain as root.

- For the Application Domain, log in through the 1 GbE host management network using the IP address or host name assigned to that domain. For example:

  ```
  ssh -l root ssc0101-mgmt
  ```

- For the Database Domain, log in through the 1 GbE host management network using the IP address or host name assigned to the SPARC T4-4 server. For example:

  ```
  ssh -l root ssc01db01
  ```

2. Locate the directory where the Oracle Configuration Manager is installed.

- For the Application Domain, Oracle Configuration Manager should be installed in the following directory:

  ```
  /opt/ocm
  ```

- For the Database Domain, Oracle Configuration Manager should be installed in the following directory:

  ```
  /usr/lib/ocm
  ```

- If the Oracle Configuration Manager is not installed in either of those directories, locate and change directories to the location where you installed Oracle Configuration Manager.

The directory where the Oracle Configuration Manager is installed will be referred to as *ocm-install-dir* for the remainder of these procedures.

3. Change directories to:

/ocm-install-dir`/ccr/bin`

4. Look for the file `emCCR`.

- If you see the `emCCR` file in this directory, then Oracle Configuration Manager has already been installed and configured on your SPARC T4-4 server. Do not proceed with these instructions in this case.

- If you do not see the `emCCR` file in this directory, then Oracle Configuration Manager has not been installed on your SPARC T4-4 server. Proceed to the next step.

5. **Type the following command to configure Oracle Configuration Manager on the SPARC T4-4 server:**

   /ocm-install-dir`/ccr/bin/configCCR`

6. **Type the following information in the appropriate fields in Oracle Configuration Manager:**

   - E-mail address/User Name
   - Password (optional)

   This installs Oracle Configuration Manager on your SPARC T4-4 server.

7. **If you are logged into the Database Domain, or if are running a database on the Application Domain, enable the collection of data from the database:**

   ```
   /ocm-install-dir/ccr/admin/scripts/installCCRSQL.sh collectconfig -s $SID -r SYS
   ```

   XML files should be generated in `/ocm-install-dir/ccr/hosts/`*SPARC-T4-4-server*`/state/review`.

   - If the XML files were generated in this directory, proceed to the next step.
   - If the XML files were not generated in this directory, contact Oracle for support with this issue.

8. **Perform a collection as follows:**

   ```
   ORACLE_HOME/ccr/bin/emCCR collect
   ```

9. **Log in to My Oracle Support.**

10. **On the home page, select the Customize page link at the top right of the Dashboard page.**

Drag the Targets button to the left and on to your dashboard.

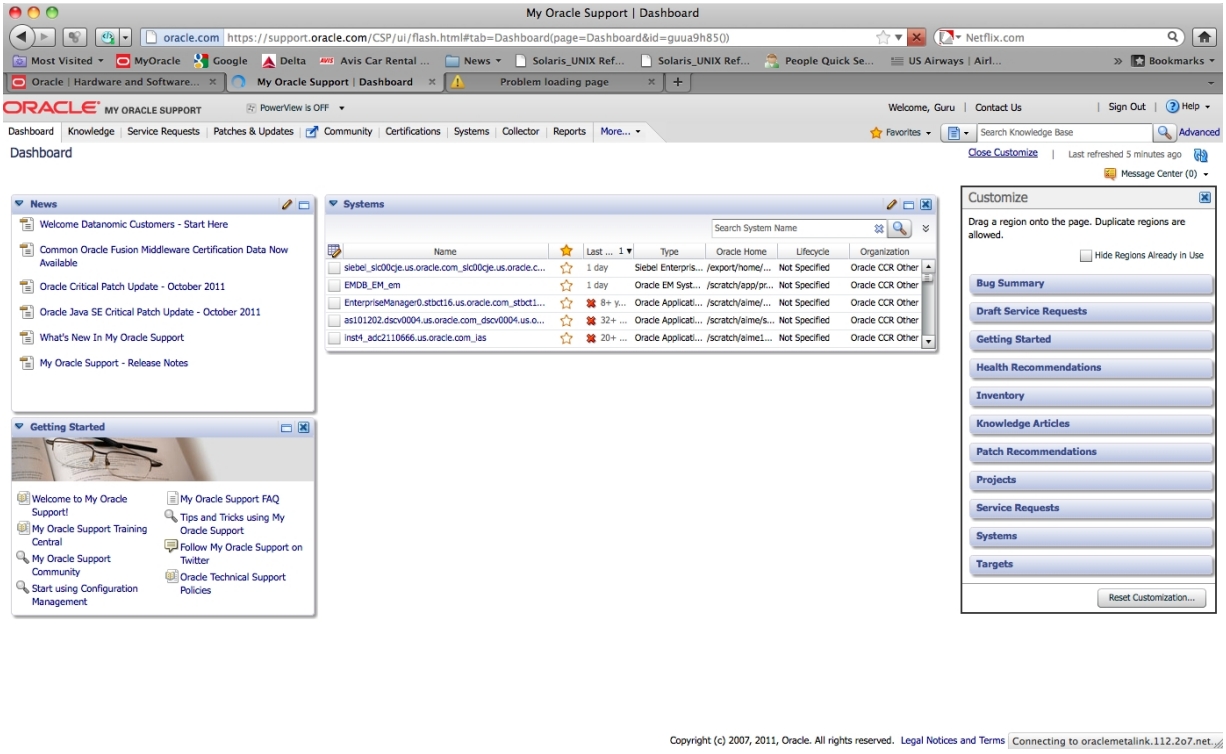

11. Search for your system in the targets search window at the top right of the screen.

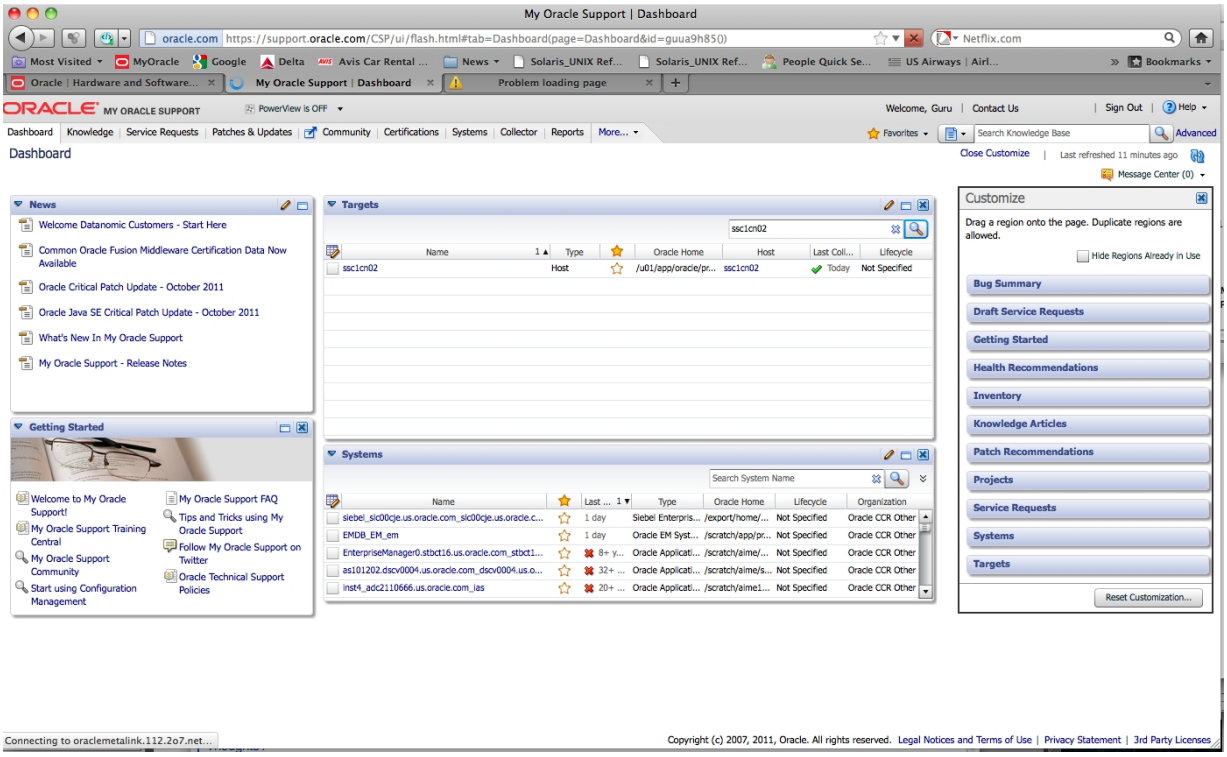

12. **Double-click on your system to get information for your system.**

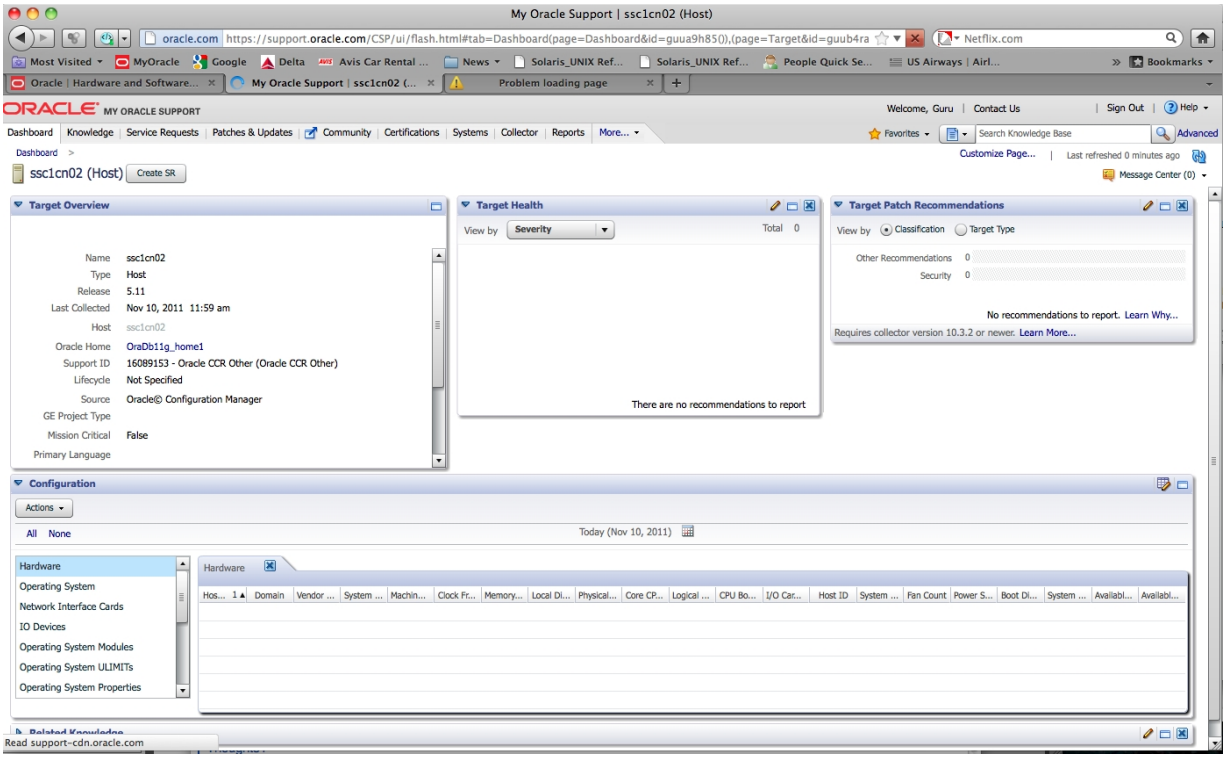

Monitoring the System Using EM Exadata Plug-in

Starting with SPARC SuperCluster 1.1 and later, you can monitor all Exadata-related software and hardware components in the cluster using Oracle Enterprise Manager Exadata 12.1.0.3 Plug-in only in the supported configuration described in the note below.

Note - Only Oracle SuperCluster with software Version 1.1 with Database Domain on Control LDOM-only environments are supported. Earlier versions of SPARC SuperCluster can be made compatible if you update to the October 2012 QMU release. You can confirm this requirement by looking at the version of the `compmon pkg` installed on the system (using either `pkg info compmon` or `pkg list compmon` commands to check). You must have the following minimum version of `compmon` installed: `pkg://exa-family/system/platform/exadata/compmon@0.5.11,5.11-0.1.0.11:20120726T024158Z`

Known issues:

- The pre-requisite check script `exadataDiscoveryPreCheck.pl` that is bundled in Exadata plug-in 12.1.0.3 does not support the `catalog.xml` file. Download the latest `exadataDiscoveryPreCheck.pl` file from My Oracle Support as described in the "Discovery Precheck Script" section of the

  Oracle Enterprise Manager Exadata Management Getting Started Guide (`docs.oracle.com/cd/E24628_01/doc.121/e27442/title.htm`)
- If multiple database clusters share the same Exadata Storage server, in one Enterprise Manager management server environment, you can discover and monitor the first DB Machine target and all its components. However, for additional DB Machine targets sharing the same Exadata Storage server, the Oracle Exadata Storage Server Grid system and the Oracle Database Exadata Storage Server System will have no Exadata Storage server members because they are already monitored.
- If the `perfquery` command installed in the Oracle SuperCluster has version 1.5.8 or later, you will encounter a bug (ID 15919339) where most columns in the HCA Port Errors metric in the host targets for the compute nodes will be blank. If there are errors occurring on the HCA ports, it will not be reported in Enterprise Manager.

  To check your version, run the following command:

  ```
  perfquery -V
  ```

Configuring Exalogic Software

This section provides information about using Exalogic software on the SPARC SuperCluster T4-4.

Overview of Exalogic-Specific Enhancements

Oracle Exalogic Elastic Cloud Software (EECS) includes performance optimizations for SPARC SuperCluster T4-4 to improve input/output, thread management, and request handling efficiency.

Additional optimizations include reduced buffer copies, which result in more efficient input/output. Finally, session replication performance and CPU utilization is improved through lazy deserialization, which avoids performing extra work on every session update that is only necessary when a server fails.

WebLogic Server clusters can be configured with cluster-wide optimizations that further improve server-to-server communication. The first optimization enables multiple replication channels, which improve network throughput among WebLogic Server cluster nodes. The second cluster optimization enables InfiniBand support for Sockets Direct Protocol, which reduces CPU utilization as network traffic bypasses the TCP stack.

Exalogic Software Prerequisites

The following are the prerequisites for configuring the Exalogic software products for the SPARC SuperCluster T4-4:

- Preconfiguring the environment, including database, storage, and network, as described in Chapter 3, “Network, Storage, and Database Preconfiguration” of the *Oracle Exalogic Enterprise Deployment Guide*, located at: http://docs.oracle.com/cd/E18476_01/doc.220/e18479/toc.htm
- Your Oracle Exalogic Domain is configured, as described in Chapter 5, “Configuration Oracle Fusion Middleware” in the *Oracle Exalogic Enterprise Deployment Guide*, located at: http://docs.oracle.com/cd/E18476_01/doc.220/e18479/toc.htm

▼ Enable Domain-Level Enhancements

To enable domain-level enhancements, complete the following steps:

1. **Log in to the Oracle WebLogic Server Administration Console.**
2. **Select** *Domainname* **in the left navigation pane. The Settings for** *Domainname* **screen is displayed. Click the** *General* **tab.**
3. **In your domain home page, select** *Enable Exalogic Optimizations***, and click** *Save***.**
4. **Activate changes.**
5. **Stop and start your domain.**

   The *Enable Exalogic Optimizations* setting collectively enables all of the individual features described in the following table. The Startup Option indicates how to independently enable and disable each feature.

Feature		
Scattered Reads		
	Description	Increased efficiency during I/O in environments with high network throughput
	Startup Option	*-Dweblogic.ScatteredReadsEnabled=true/false*
	MBean	*KernelMBean.setScatteredReadsEnabled*
Gathered Writes		
	Description	Increased efficiency during I/O in environments with high network throughput
	Startup Option	*-Dweblogic.GatheredWritesEnabled=true/false*
	MBean	*KernelMBean.setGatheredWritesEnabled*
Lazy Deserialization		
	Description	Increased efficiency with session replication
	Startup Option	*-Dweblogic.replication.enableLazyDeserialization=true/false*
	MBean	*ClusterMBean.setSessionLazyDeserializationEnabled*

Note - After enabling the optimizations, you may see the following message: `java.io.IOException: Broken pipe`. You may see the same message when storage failover occurs. In either case, you can ignore the error message.

▼ Enable Cluster-Level Session Replication Enhancements

1. **Ensure that Managed Servers in the** *Dept1_Cluster1* **cluster are up and running, as described in Section 5.16 "Starting Managed Servers on ComputeNode1 and ComputeNode2" of the *Oracle® Fusion Middleware Exalogic Enterprise Deployment Guide* at `http://docs.oracle.com/cd/E18476_01/doc.220/e18479/create_domain.htm#BABEGAFB`.**

2. **To set replication ports for a Managed Server, such as** *WLS1***, complete the following steps:**

   a. **Under** *Domain Structure***, click** *Environment* **and** *Servers***. The Summary of Servers page is displayed.**

   b. **Click** *WLS1* **on the list of servers. The Settings for WLS1 is displayed.**

   c. **Click the** *Cluster* **tab.**

   d. **In the** *Replication Ports* **field, enter a range of ports for configuring multiple replication channels. For example, replication channels for Managed Servers in** *Dept_1_Cluster1* **can listen on ports starting from** *7005* **to** *7015***. To specify this range of ports, enter** *7005-7015***.**

3. **Create a custom network channel for each Managed Server in the cluster (for example,** *WLS1***) as follows:**

   a. **Log in to the Oracle WebLogic Server Administration Console.**

   b. **If you have not already done so, click** *Lock & Edit* **in the Change Center.**

   c. **In the left pane of the Console, expand** *Environment* **and select** *Servers***.**
   The *Summary of Servers page* is displayed.

   d. **In the Servers table, click** *WLS1* **Managed Server instance.**

   e. **Select** *Protocols***, and then** *Channels***.**

f. Click *New***.**

g. Enter *ReplicationChannel* **as the name of the new network channel and select** *t3* **as the protocol, then click** *Next***.**

h. Enter the following information:

Listen address: *10.0.0.1*

Note - This is the floating IP assigned to *WLS1*.

Listen port: *700*5

i. Click *Next***, and in the Network Channel Properties page, select** *Enabled* **and** *Outbound Enabled***.**

j. Click *Finish***.**

k. Under the Network Channels table, select *ReplicationChannel***, the network channel you created for the** *WLS1* **Managed Server.**

l. Expand *Advanced***, and select** *Enable SDP Protocol***.**

m. Click *Save***.**

n. To activate these changes, in the Change Center of the Administration Console, click *Activate Changes***.**

You must repeat the above steps to create a network channel each for the remaining Managed Servers in the *Dept1_Cluster1* cluster. Enter the required properties, as described in the following table.

Managed Servers in Dept1_Cluster1	Name	Protocol	Listen Address	Listen Port	Additional Channel Ports
WLS2	*ReplicationChannel*	t3	10.0.0.2	7005	7006 to 7014
WLS3	*ReplicationChannel*	t3	10.0.0.3	7005	7006 to 7014
WLS4	*ReplicationChannel*	t3	10.0.0.4	7005	7006 to 7014
WLS5	*ReplicationChannel*	t3	10.0.0.5	7005	7006 to 7014
WLS6	*ReplicationChannel*	t3	10.0.0.6	7005	7006 to 7014
WLS7	*ReplicationChannel*	t3	10.0.0.7	7005	7006 to 7014
WLS8	*ReplicationChannel*	t3	10.0.0.8	7005	7006 to 7014

4. **After creating the network channel for each of the Managed Servers in your cluster, click** *Environment* **>** *Clusters***. The Summary of Clusters page is displayed.**

5. **Click** *Dept1_Cluster1* **(this is the example cluster to which you will deploy a web application at a later time). The Settings for Dept1_Cluster1 page is displayed.**

6. **Click the** *Replication* **tab.**

7. **In the** *Replication Channel* **field, ensure that** *ReplicationChannel* **is set as the name of the channel to be used for replication traffic.**

8. **In the** *Advanced* **section, select the** *Enable One Way RMI for Replication* **option.**

9. **Click** *Save***.**

10. **Activate changes, and restart the Managed Servers.**

11. **Manually add the system property** *-Djava.net.preferIPv4Stack=true* **to the** *startWebLogic.sh* **script, which is located in the** *bin* **directory of** *base_domain***, using a text editor as follows:**

    a. **Locate the following line in the startWebLogic.sh script:**
    *. ${DOMAIN_HOME}/bin/setDomainEnv.sh $**

    b. **Add the following property immediately after the above entry:**
    JAVA_OPTIONS="${JAVA_OPTIONS} -Djava.net.preferIPv4Stack=true"

    c. **Save the file and close.**

12. **Restart all Managed Servers as follows:**

    a. **In the administration console, click** *Environment* **>** *Servers***. The Summary of Servers page is displayed.**

    b. **Select a Managed Server, such as** *WLS1***, by clicking** *WLS1***. The Settings for WLS1 page is displayed.**

    c. **Click the** *Control* **tab. Select** *WLS1* **in the Server Status table. Click** *Start***.**

    d. **Repeat these steps for each of the Managed Servers in the WebLogic cluster.**

Note - To verify that multiple listening ports were opened, you can either run the `netstat -na` command on the command line or check the Managed Server logs.

Configuring Grid Link Data Source for Dept1_Cluster1

You must create a Grid Link Data Source for JDBC connectivity between Oracle WebLogic Server and a service targeted to an Oracle RAC cluster. It uses the Oracle Notification Service (ONS) to adaptively respond to state changes in an Oracle RAC instance. This section includes the following:

Grid Link Data Source

A Grid Link data source includes the features of generic data sources plus the following support for Oracle RAC:

Fast Connection Failover

A Grid Link data source uses Fast Connection Failover to:

- Provide rapid failure detection.
- Abort and remove invalid connections from the connection pool.
- Perform graceful shutdown for planned and unplanned Oracle RAC node outages. The data source allows in-progress transactions to complete before closing connections. New requests are load balanced to an active Oracle RAC node.
- Adapt to changes in topology, such as adding a new node.
- Distribute runtime work requests to all active Oracle RAC instances.

See Fast Connection Failover in the *Oracle Database JDBC Developer's Guide and Reference* at: `http://docs.oracle.com/cd/B19306_01/java.102/b14355/fstconfo.htm`.

Runtime Connection Load Balancing

Runtime Connection Load Balancing allows WebLogic Server to:

- Adjust the distribution of work based on back end node capacities such as CPU, availability, and response time.
- React to changes in Oracle RAC topology.

- Manage pooled connections for high performance and scalability.

If FAN (Fast Application Notification) is not enabled, Grid link data sources use a round-robin load balancing algorithm to allocate connections to Oracle RAC nodes.

XA Affinity

XA Affinity for global transactions ensures all the data base operations for a global transaction performed on an Oracle RAC cluster are directed to the same Oracle RAC instance. The first connection request for an XA transaction is load balanced using RCLB and is assigned an Affinity context. All subsequent connection requests are routed to the same Oracle RAC instance using the Affinity context of the first connection.

SCAN Addresses

Oracle Single Client Access Name (SCAN) addresses can be used to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. A Grid Link data source containing SCAN addresses does not need to change if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN URLs for your environment. For more information, see *http://www.oracle.com/technetwork/database/clustering/overview/scan-129069.pdf*.

Secure Communication using Oracle Wallet

Allows you to configure secure communication with the ONS listener using Oracle Wallet.

▼ Create a GridLink Data Source on Dept1_Cluster1

1. **Log in to the Oracle WebLogic Server Administration Console.**
2. **If you have not already done so, in the Change Center of the Administration Console, click** *Lock & Edit***.**
3. **In the** *Domain Structure* **tree, expand** *Services***, then select** *Data Sources***.**
4. **On the** *Summary of Data Sources* **page, click** *New* **and select** *GridLink Data Source***.**
   The *Create a New JDBC GridLink Data Source* page is displayed.
5. **Enter the following information:**
   - Enter a logical name for the datasource in the Name field. For example, *gridlink*.

- Enter a name for JNDI. For example, *jdbc/gridlink*.
- Click *Next.*

6. **In the** *Transaction Options* **page, de-select** *Supports Global Transactions***, and click** *Next***.**

7. **Select** *Enter individual listener information* **and click** *Next***.**

8. **Enter the following connection properties:**

   - *Service Name*: Enter the name of the Oracle RAC service in the **`Service Name`** field. For example, enter **`myService`** in *Service Name*.

     Note - The Oracle RAC Service name is defined on the database, and it is not a fixed name.

   - *Host Name* - Enter the DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
   - *Port* - Enter the port on which the database server listens for connections requests.
   - *Database User Name*: Enter the database user name. For example, *myDataBase*.
   - Password: Enter the password. For example, *myPassword1*.
   - Confirm Password and click *Next*.

   Tip - For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

   The console automatically generates the complete JDBC URL. For example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)
   (HOST=left)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(HOST=right)(PORT=1234))(ADDRESS=
   (PROTOCOL=TCP)(HOST=center)(PORT=1234)))(CONNECT_DATA=(SERVICE_NAME=myService)))
   ```

9. **On the** *Test GridLink Database Connection* **page, review the connection parameters and click** *Test All Listeners***.**

   Oracle WebLogic attempts to create a connection from the Administration Server to the database. Results from the connection test are displayed at the top of the page. If the test is unsuccessful, you should correct any configuration errors and retry the test.

   Click *Next*.

10. **In the** *ONS Client Configuration* **page, do the following:**

    - Select *Fan Enabled* to subscribe to and process Oracle FAN events.
    - In ONS host and port, enter a comma-separate list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.

- Click *Next*.

11. **On the** *Test ONS client configuration* **page, review the connection parameters and click** *Test All ONS Nodes***.**
    Click *Next*.

12. **In the** *Select Targets* **page, select** *Dept1_Cluster1* **as the target and** *All Servers in the cluster*

13. **Click** *Finish***.**

14. **Click** *Activate Changes***.**

15. **Now you must configure SDP-enabled JDBC drivers for the cluster. For instructions, go to "Configuring SDP-Enabled JDBC Drivers for Dept1_Cluster1" on page 161.**

Configuring SDP-Enabled JDBC Drivers for Dept1_Cluster1

You must configure SDP-enabled JDBC drivers for the *Dept1_Cluster1* cluster.

This section discusses the following topics:

Configure the Database to Support Infiniband

Before enabling SDP support for JDBC, you must configure the database to support InfiniBand, as described in the " Configuring SDP Protocol Support for Infinband Network Communication to the Database Server" section in the *Oracle Database Net Services Administrator's Guide*, located at:

`http://download.oracle.com/docs/cd/B28359_01/network.111/b28316/performance.htm#i1008413`

Ensure that you set the protocol to SDP.

▼ Enable SDP Support for JDBC

Complete the following steps:

1. **Ensure that you have created the Grid Link Data Sources for the JDBC connectivity on** *ComputeNode1* **and** *ComputeNode2***, as described in Section 7.6 "Configuring Grid Link Data Source for Dept1_Cluster1" of the *Oracle® Fusion Middleware Exalogic Enterprise Deployment Guide* at `http://docs.oracle.com/cd/E18476_01/doc.220/e18479/optimization.htm#BABHEDI`.**
   The console automatically generates the complete JDBC URL, as shown in the following example:
   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=192.x.x.x)(PORT=1522))
   (CONNECT_DATA=(SERVICE_NAME=myservice)))
   ```
2. **In the JDBC URL, replace** *TCP* **protocol with** *SDP* **protocol. For example:**
   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=sdp)(HOST=192.x.x.x)(PORT=1522))
   (CONNECT_DATA=(SERVICE_NAME=myservice)))
   ```
3. **Manually add the system property** *-Djava.net.preferIPv4Stack=true* **to the** *startWebLogic.sh* **script, which is located in the** *bin* **directory of** *base_domain***, using a text editor as follows:**
   a. **Locate the following line in the** *startWebLogic.sh* **script:**
      *. ${DOMAIN_HOME}/bin/setDomainEnv.sh $**
   b. **Add the following property immediately after the above entry:**
      ```
      JAVA_OPTIONS="${JAVA_OPTIONS} -Djava.net.preferIPv4Stack=true -Doracle.net.SDP=true"
      ```
   c. **Save the file and close.**
4. **Restart the Managed Server as follows:**
   a. **In the administration console, click** *Environment* **>** *Servers***. The Summary of Servers page is displayed.**
   b. **Select a Managed Server, such as** *WLS1***, by clicking** *WLS1***. The Settings for WLS1 page is displayed.**
   c. **Click the** *Control* **tab. Select** *WLS1* **in the Server Status table. Click** *Start***.**

▼ Monitor SDP Sockets Using `netstat` on Oracle Solaris

You can monitor SDP sockets by running the `netstat` command on the Application Domains running Oracle Solaris 11 containing the Oracle Exalogic Elastic Cloud Software in the SPARC

SuperCluster T4-4. Run the `netstat` command on these Application Domains running Oracle Solaris 11 and on the Database Domains to monitor SDP traffic between the Application Domains running Oracle Solaris 11 and the Database Domains.

- **Log in to the operating system as a *root*, and run the following command on the command line:**

```
# netstat -f sdp -s l
```

This command displays the status of all SDP sockets (established or not), as in the following sample output:

```
SDP    sdpActiveOpens      = 66357    sdpCurrEstab        =   748
       sdpPrFails          =    0    sdpRejects          =     0
       sdpOutSegs          =39985638793
       sdpInDataBytes      =9450383834191
       sdpOutDataBytes     =6228930927986

SDP    sdpActiveOpens      =     0    sdpCurrEstab        =     0
       sdpPrFails          =     0    sdpRejects          =     0
       sdpInSegs           = 14547
       sdpOutSegs          = 14525
       sdpInDataBytes      =3537194
       sdpOutDataBytes     =2470907
```

Configuring SDP InfiniBand Listener for Exalogic Connections

This section is intended for users who have the Oracle Exalogic Elastic Cloud Software connected to a Database Domain in the SPARC SuperCluster T4-4. It describes how to create an SDP listener on the InfiniBand network. The tasks described in this section should be run on each of the Database Domains in the SPARC SuperCluster T4-4 that uses SDP-enabled data sources.

■ “Create an SDP Listener on the InfiniBand Network” on page 163

▼ Create an SDP Listener on the InfiniBand Network

Oracle RAC 11*g* Release 2 supports client connections across multiple networks, and it provides load balancing and failover of client connections within the network they are connecting. To add a listener for the Oracle Exalogic Elastic Cloud Software connections coming in on the Infiniband network, first add a network resource for the Infiniband network with Virtual IP addresses.

Note - This example lists two Database Domains. If you have more than two Database Domains in your SPARC SuperCluster T4-4, you must repeat Database Domain-specific lines for each Database Domain in the cluster.

1. **Edit** */etc/hosts* **on each Database Domain in the cluster to add the virtual IP addresses you will use for the InfiniBand network. Make sure that these IP addresses are not used. The following is an example:**

   # Added for Listener over IB
   192.168.10.21 ssc01db01-ibvip.mycompany.com ssc01db01-ibvip
   192.168.10.22 ssc01db02-ibvip.mycompany.com ssc01db02-ibvip

2. **On one of the Database Domains, as the** *root* **user, create a network resource for the InfiniBand network, as in the following example:**

   ```
   # /u01/app/grid/product/11.2.0.2/bin/srvctl add network -k 2 -S
   192.168.10.0/255.255.255.0/bondib0
   ```

3. **Validate that the network was added correctly, by running one of the following commands:**

   ```
   # /u01/app/grid/product/11.2.0.2/bin/crsctl stat res -t | grep net
   ora.net1.network
   ```
   ora.net2.network -- Output indicating new Network resource
   or
   ```
   # /u01/app/grid/product/11.2.0.2/bin/srvctl config network -k 2
   ```
   Network exists: 2/192.168.10.0/255.255.255.0/bondib0, type static -- Output indicating Network resource on the 192.168.10.0 subnet

4. **Add the Virtual IP addresses on the network created in Step 2, for each node in the cluster.**

   ```
   srvctl add vip -n ssc01db01 -A ssc01db01-ibvip/255.255.255.0/bondib0 -k 2
   srvctl add vip -n ssc01db02 -A ssc01db02-ibvip/255.255.255.0/bondib0 -k 2
   ```

5. **As the "***oracle***" user (who owns the Grid Infrastructure Home), add a listener which will listen on the VIP addresses created in Step 3.**

   ```
   srvctl add listener -l LISTENER_IB -k 2 -p TCP:1522,/SDP:1522
   ```

6. **For each database that will accept connections from the middle tier, modify the** *listener_networks init* **parameter to allow load balancing and failover across multiple networks (Ethernet and InfiniBand). You can either enter the full tnsnames syntax in the initialization parameter or create entries in** *tnsnames.ora* **in** *$ORACLE_HOME/network/admin* **directory. The** *TNSNAMES.ORA* **entries must exist**

in the *GRID_HOME*. **The following example first updates** *tnsnames.ora*. **Complete this step on each Database Domain in the cluster with the correct IP addresses for that Database Domain.** *LISTENER_IBREMOTE* **should list all other Database Domains that are in the cluster.** *DBM_IB* **should list all Database Domains in the cluster.**

Note - The *TNSNAMES* entry is only read by the database instance on startup, if you modify the entry that is referred to by any *init.ora* parameter (*LISTENER_NETWORKS*), you must restart the instance or issue an *ALTER SYSTEM SET LISTENER_NETWORKS* command for the modifications to take affect by the instance.

```
(DESCRIPTION =
DBM =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01-scan)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)

(SERVICE_NAME = dbm)
))

DBM_IB =
(DESCRIPTION =
(LOAD_BALANCE=on)
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db01-ibvip)(PORT = 1522))
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db02-ibvip)(PORT = 1522))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dbm)
))

LISTENER_IBREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db02-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IBLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db01-ibvip.mycompany.com)(PORT = 1522))
(ADDRESS = (PROTOCOL = SDP)(HOST = ssc01db01-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IPLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc0101-vip.mycompany.com)(PORT = 1521))
))
```

```
LISTENER_IPREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01-scan.mycompany.com)(PORT = 1521))
))
```

7. **Modify the** *listener_networks init* **parameter. Connect to the database instance as** *sysdba***.**

```
SQLPLUS> alter system set listener_networks='((NAME=network2)
(LOCAL_LISTENER=LISTENER_IBLOCAL)(REMOTE_LISTENER=LISTENER_IBREMOTE))',
'((NAME=network1)(LOCAL_LISTENER=LISTENER_IPLOCAL)(REMOTE_LISTENER=LISTENER_IPREMOTE))'
scope=both;
```

8. **Stop and start** *LISTENER_IB* **for the modification in Step 7.**

```
srvctl stop listener -l LISTENER_IB

srvctl start listener -l LISTENER_IB
```

Understanding Internal Cabling

These topics show the cable layouts for SPARC SuperCluster T4-4.

Reference Drawings

The **InfiniBand (IB) fabric** connects together the IB switches and the servers listed below in the SPARC SuperCluster T4-4. For IB cable information, see "InfiniBand Fabric " on page 173.

The **Ethernet Management Switch** also connects to the servers listed below. For 1 Gb/s Cat 5E cabling information, see "Ethernet Management Switch" on page 178.

The following figures are in this section:

- **Sun Datacenter InfiniBand Switch 36**, Figure 12
  - IB Spine Switch (slot U1)
  - IB Leaf Switch No. 1 (slot U18)
  - IB Leaf Switch No. 2 (slot U24)
- **SPARC T4-4 servers**, Figure 13
  - SPARC T4-4 No. 1 (slot 8)
  - SPARC T4-4 No. 2 (slot 13)
  - SPARC T4-4 No. 3 (slot 27) (full rack only)
  - SPARC T4-4 No. 4 (slot 32) (full rack only)
- **Sun ZFS Storage 7320 storage controllers**, Figure 14
  - Sun ZFS Storage 7320 storage controller No. 1 (slot 25)

  - Sun ZFS Storage 7320 storage controller No. 2 (slot 26)
- **Exadata Storage Servers**, Figure 15
  - Exadata Storage Server No. 1 (slot 2)
  - Exadata Storage Server No. 2 (slot 4)
  - Exadata Storage Server No. 3 (slot 6)
  - Exadata Storage Server No. 4 (slot 37) (full rack only)
  - Exadata Storage Server No. 5 (slot 39) (full rack only)
  - Exadata Storage Server No. 6 (slot 41) (full rack only)
- The **Ethernet Management Switch**, Figure 16
- **PDU circuit breakers and AC sockets**, Figure 17

FIGURE 12 Sun Datacenter InfiniBand Switch 36

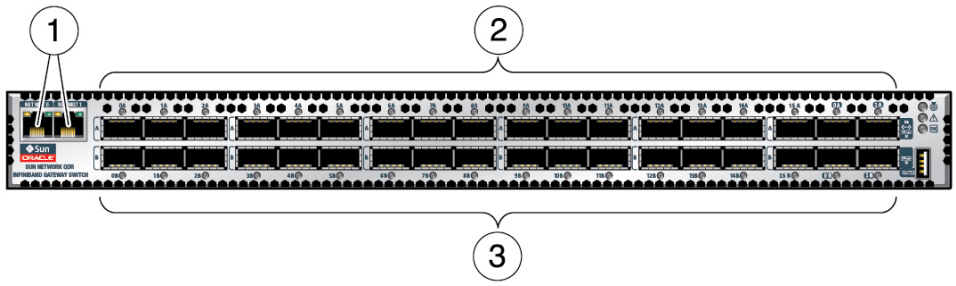

Figure Legend

1	NET MGT 0 and NET MGT 1 ports
2	InfiniBand ports 0A-17A (upper ports)
3	InfiniBand ports 0B-17B (lower ports)

FIGURE 13 SPARC T4-4 Server

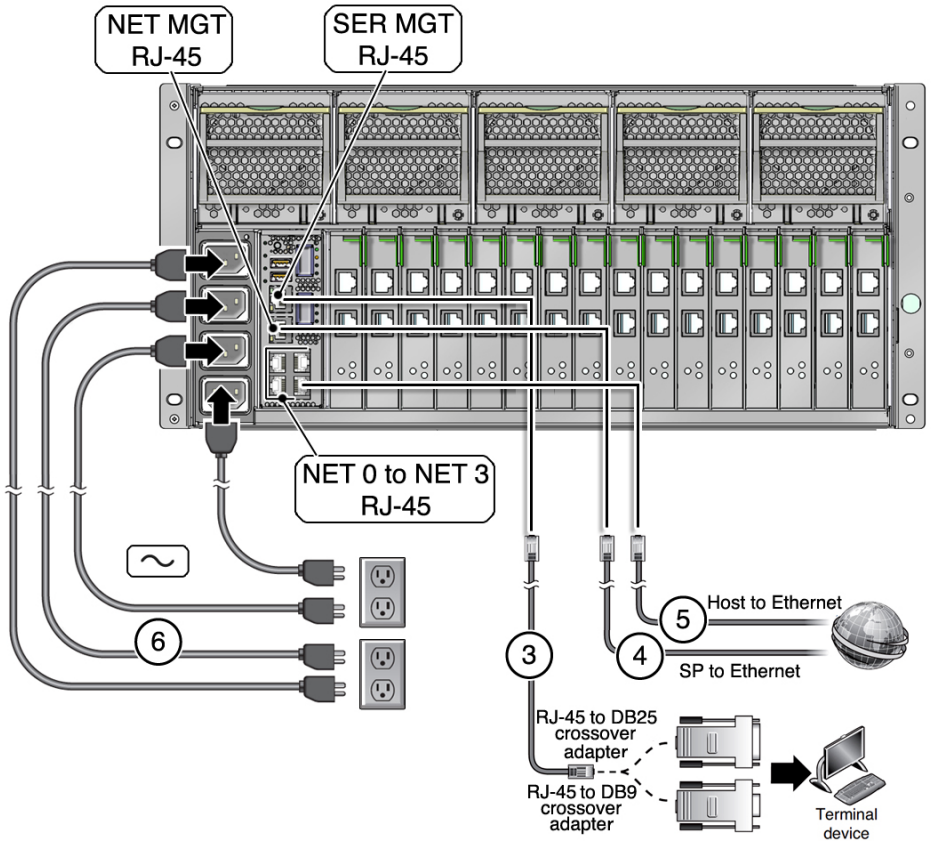

Figure Legend

1. NET MGT, network management port
2. SER MGT, serial management port
3. RJ-45 cable from SER MGT port to a terminal device, might require an adapter
4. RJ-45 cable, connects the SPARC T4-4 service processor to the Ethernet
5. RJ-45 cables, connect network ports NET 0-3 to the Ethernet
6. Power cords

FIGURE 14 Sun ZFS Storage 7320 Storage Controller

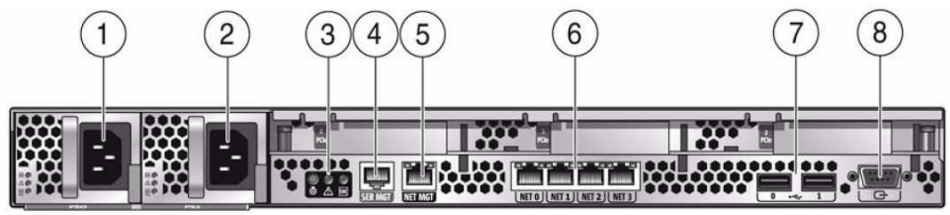

Figure Legend

1 Power supply 1
2 Power supply 2
3 System status LEDs
4 Serial management port
5 Service processor network management port
6 Gigabit Ethernet ports NET 0, 1, 2, 3
7 USB ports 0, 1
8 HD15 video connector

FIGURE 15 Exadata Storage Server

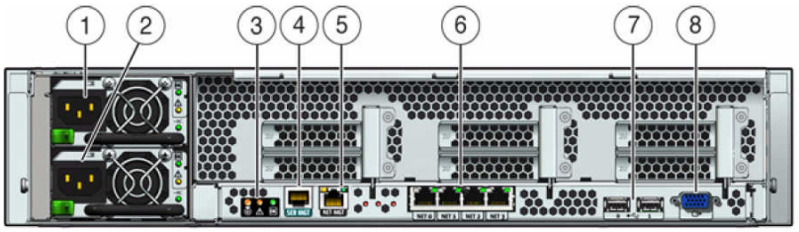

Figure Legend

1 Power supply 1
2 Power supply 2
3 System status LEDs
4 Serial management port
5 Service processor network management port
6 Gigabit Ethernet ports NET 0, 1, 2, 3
7 USB ports 0, 1
8 HD15 video connector

FIGURE 16 Ethernet Management Switch (Cisco Catalyst 4948 Ethernet Switch)

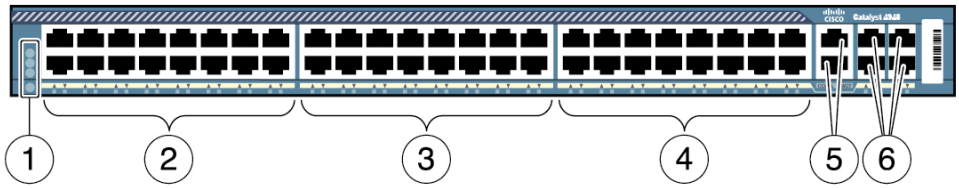

Figure Legend

1 Indicators and reset switch
2 Ports 1-16, 10/100/1000BASE-T Ethernet
3 Ports 17-32, 10/100/1000BASE-T Ethernet
4 Ports 33-48, 10/100/1000BASE-T Ethernet
5 CON (upper), MGT (lower)
6 Ports 45-48, 10-Gigabit Ethernet

FIGURE 17 Circuit Breakers and AC Sockets on the PDU

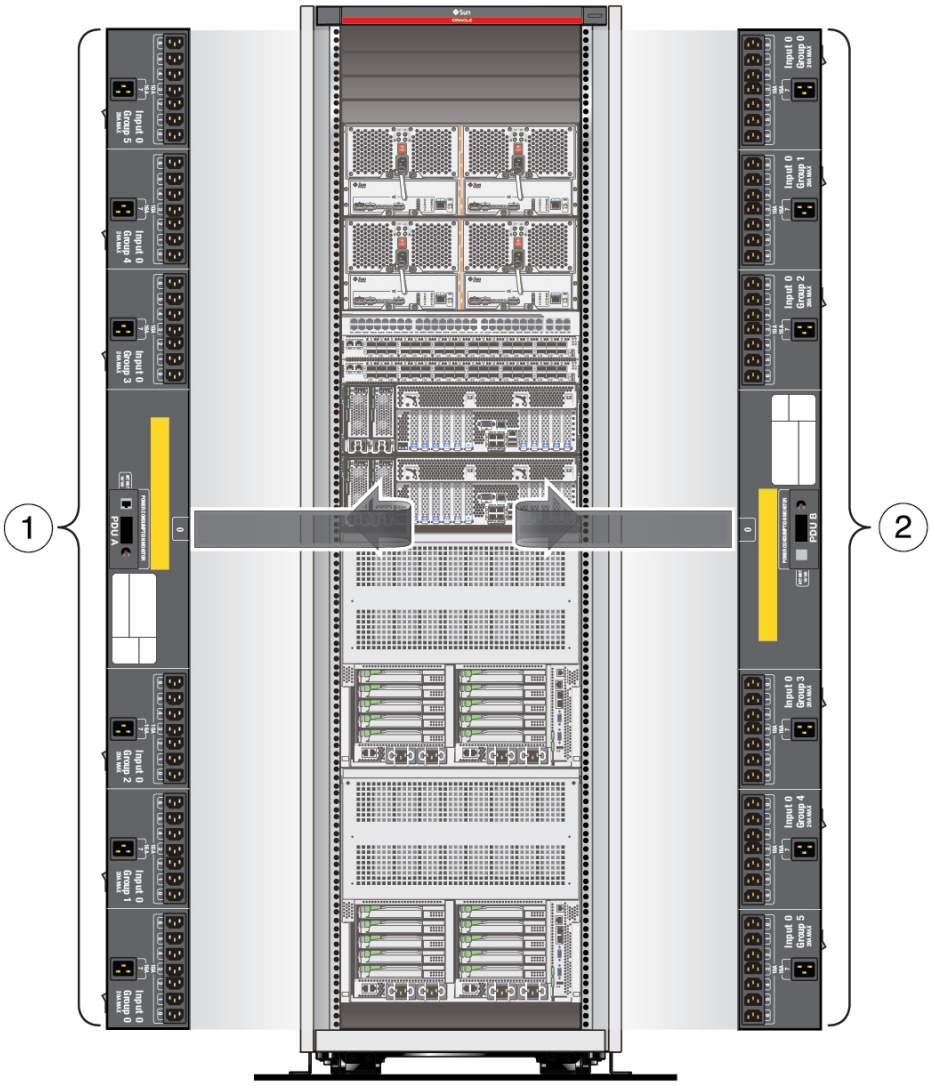

Figure Legend

1 PDU A
2 PDU B

InfiniBand Fabric

The following tables show InfiniBand (IB) connections for these switches.

- IB Spine switch, see Table 14
- IB leaf switch No. 1
  - Full rack, see Table 15
  - Half rack, see Table 16
- IB leaf switch No. 2
  - Full rack, see Table 17
  - Half rack, see Table 18

Note - For 1 Gb/s Ethernet, see "Ethernet Management Switch" on page 178.

IB Spine Switch

TABLE 14 **IB Spine Switch** (U1) Connections

Spine Switch	Port	To Device	Device Location	Device Port
U1	1B	Leaf switch 1	U18	8B
U1	0B	Leaf switch 2	U24	8B

IB Leaf Switch No. 1

TABLE 15 **Full Rack** Connections for IB Leaf Switch 1 (U18)

Leaf Switch 1	Port	To Device	Device Location	Device Port
U18	0A	T4-4 No. 4	U32	PCIE-12 P2
U18	0B	T4-4 No. 4	U32	PCIE-9 P2
U18	1A	T4-4 No. 4	U32	PCIE-7 P2
U18	1B	T4-4 No. 4	U32	PCIE-1 P2
U18	2A	—	—	—
U18	2B	—	—	—
U18	3A	Sun ZFS Storage 7320 storage controllers No. 2	U26	PCIE-0 P1
U18	3B	T4-4 No. 2	U13	PCIE-12 P1
U18	4A	T4-4 No. 2	U13	PCIE-9 P1
U18	4B	T4-4 No. 2	U13	PCIE-7 P1
U18	5A	T4-4 No. 2	U13	PCIE-1 P1

Leaf Switch 1	Port	To Device	Device Location	Device Port
U18	5B	T4-4 No. 3	U27	PCIE-12 P2
U18	6A	T4-4 No. 3	U27	PCIE-9 P2
U18	6B	T4-4 No. 3	U27	PCIE-7 P2
U18	7A	T4-4 No. 3	U27	PCIE-1 P2
U18	7B	—	—	—
U18	8A	IB leaf switch 2	U24	8A
U18	8B	IB spine switch	U1	1B
U18	9A	IB leaf switch 2	U24	9B
U18	9B	IB leaf switch 2	U24	9A
U18	10A	IB leaf switch 2	U24	10B
U18	10B	IB leaf switch 2	U24	10A
U18	11A	IB leaf switch 2	U24	11B
U18	11B	IB leaf switch 2	U24	11A
U18	12A	—	—	—
U18	12B	T4-4 No. 1	U8	PCIE-12 P1
U18	13A	T4-4 No. 1	U8	PCIE-9 P1
U18	13B	T4-4 No. 1	U8	PCIE-7 P1
U18	14A	T4-4 No. 1	U8	PCIE-1 P1
U18	14B	Sun ZFS Storage 7320 storage controller No. 1	U25	PCIE-0 P1
U18	15A	Exadata Storage Server No. 6	U41	PCIE-3 P2
U18	15B	Exadata Storage Server No. 5	U39	PCIE-3 P2
U18	16A	Exadata Storage Server No. 4	U37	PCIE-3 P2
U18	16B	Exadata Storage Server No. 3	U6	PCIE-3 P1
U18	17A	Exadata Storage Server No. 2	U4	PCIE-3 P1
U18	17B	Exadata Storage Server No. 1	U2	PCIE-3 P1

TABLE 16 **Half Rack** Connections for IB Leaf Switch No. 1 (U18)

Leaf Switch 1	Port	To Device	Device Location	Device Port
U18	0A	—	—	—
U18	0B	—	—	—
U18	1A	—	—	—
U18	1B	—	—	—
U18	2A	—	—	—
U18	2B	—	—	—
U18	3A	Sun ZFS Storage 7320 storage controllers No. 2	U26	PCIE-0 P1
U18	3B	T4-4 No. 2	U13	PCIE-12 P1

Leaf Switch 1	Port	To Device	Device Location	Device Port
U18	4A	T4-4 No. 2	U13	PCIE-9 P1
U18	4B	T4-4 No. 2	U13	PCIE-7 P1
U18	5A	T4-4 No. 2	U13	PCIE-1 P1
U18	5B	—	—	—
U18	6A	—	—	—
U18	6B	—	—	—
U18	7A	—	—	—
U18	7B	—	—	—
U18	8A	IB leaf switch No. 2	U24	8A
U18	8B	IB spine switch	U1	1B
U18	9A	IB leaf switch No. 2	U24	9B
U18	9B	IB leaf switch No. 2	U24	9A
U18	10A	IB leaf switch No. 2	U24	10B
U18	10B	IB leaf switch No. 2	U24	10A
U18	11A	IB leaf switch No. 2	U24	11B
U18	11B	IB leaf switch No. 2	U24	11A
U18	12A	—	—	—
U18	12B	T4-4 No. 1	U8	PCIE-12 P1
U18	13A	T4-4 No. 1	U8	PCIE-9 P1
U18	13B	T4-4 No. 1	U8	PCIE-7 P1
U18	14A	T4-4 No. 1	U8	PCIE-1 P1
U18	14B	Sun ZFS Storage 7320 storage controller No. 1	U25	PCIE-0 P1
U18	15A	—	—	—
U18	15B	—	—	—
U18	16A	—	—	—
U18	16B	Exadata Storage Server No. 3	U6	PCIE-3 P1
U18	17A	Exadata Storage Server No. 2	U4	PCIE-3 P1
U18	17B	Exadata Storage Server No. 1	U2	PCIE-3 P1

IB Leaf Switch No. 2

TABLE 17 **Full Rack** Connections for IB Leaf Switch No. 2 (U24)

Leaf Switch 2	Port	To Device	Device Location	Device Port
U24	0A	T4-4 No. 4	U32	PCIE-12 P1
U24	0B	T4-4 No. 4	U32	PCIE-9 P1
U24	1A	T4-4 No. 4	U32	PCIE-7 P1

Leaf Switch 2	Port	To Device	Device Location	Device Port
U24	1B	T4-4 No. 4	U32	PCIE-1 P1
U24	2A	—	—	—
U24	2B	—	—	—
U24	3A	Sun ZFS Storage 7320 storage controller No. 2	U26	PCIE-0 P2
U24	3B	T4-4 No. 2	U13	PCIE-12 P2
U24	4A	T4-4 No. 2	U13	PCIE-9 P2
U24	4B	T4-4 No. 2	U13	PCIE-7 P2
U24	5A	T4-4 No. 2	U13	PCIE-1 P2
U24	5B	T4-4 No. 3	U27	PCIE-12 P1
U24	6A	T4-4 No. 3	U27	PCIE-9 P1
U24	6B	T4-4 No. 3	U27	PCIE-7 P1
U24	7A	T4-4 No. 3	U27	PCIE-1 P1
U24	7B	—	—	—
U24	8A	IB leaf switch No. 1	U18	8A
U24	8B	IB spine switch	U1	0B
U24	9A	IB leaf switch No. 1	U18	9B
U24	9B	IB leaf switch No. 1	U18	9A
U24	10A	IB leaf switch No. 1	U18	10B
U24	10B	IB leaf switch No. 1	U18	10A
U24	11A	IB leaf switch No. 1	U18	11B
U24	11B	IB leaf switch No. 1	U18	11A
U24	12A	—	—	—
U24	12B	T4-4 No. 1	U8	PCIE-12 P2
U24	13A	T4-4 No. 1	U8	PCIE-9 P2
U24	13B	T4-4 No. 1	U8	PCIE-7 P2
U24	14A	T4-4 No. 1	U8	PCIE-1 P2
U24	14B	Sun ZFS Storage 7320 storage controller No. 1	U25	PCIE-0 P2
U24	15A	Exadata Storage Server No. 6	U41	PCIE 3 P1
U24	15B	Exadata Storage Server No. 5	U39	PCIE 3 P1
U24	16A	Exadata Storage Server No. 4	U37	PCIE 3 P1
U24	16B	Exadata Storage Server No. 3	U6	PCIE-3 P2
U24	17A	Exadata Storage Server No. 2	U4	PCIE-3 P2
U24	17B	Exadata Storage Server No. 1	U2	PCIE-3 P2

TABLE 18 **Half Rack** Connections for IB Leaf Switch No. 2 (U24)

Leaf Switch 2	Port	To Device	Device Location	Device Port
U24	0A	—	—	—
U24	0B	—	—	—
U24	1A	—	—	—
U24	1B	—	—	—
U24	2A	—	—	—
U24	2B	—	—	—
U24	3A	Sun ZFS Storage 7320 storage controller No. 2	U26	PCIE-0 P2
U24	3B	T4-4 No. 2	U13	PCIE-12 P2
U24	4A	T4-4 No. 2	U13	PCIE-9 P2
U24	4B	T4-4 No. 2	U13	PCIE-7 P2
U24	5A	T4-4 No. 2	U13	PCIE-1 P2
U24	5B	—	—	—
U24	6A	—	—	—
U24	6B	—	—	—
U24	7A	—	—	—
U24	7B	—	—	—
U24	8A	IB leaf switch No. 1	U18	8A
U24	8B	IB spine switch	U1	0B
U24	9A	IB leaf switch No. 1	U18	9B
U24	9B	IB leaf switch No. 1	U18	9A
U24	10A	IB leaf switch No. 1	U18	10B
U24	10B	IB leaf switch No. 1	U18	10A
U24	11A	IB leaf switch No. 1	U18	11B
U24	11B	IB leaf switch No. 1	U18	11A
U24	12A	—	—	—
U24	12B	T4-4 No. 1	U8	PCIE-12 P2
U24	13A	T4-4 No. 1	U8	PCIE-9 P2
U24	13B	T4-4 No. 1	U8	PCIE-7 P2
U24	14A	T4-4 No. 1	U8	PCIE-1 P2
U24	14B	Sun ZFS Storage 7320 storage controller No. 1	U25	PCIE-0 P2
U24	15A	—	—	—
U24	15B	—	—	—
U24	16A	—	—	—
U24	16B	Exadata Storage Server No. 3	U6	PCIE-3 P2
U24	17A	Exadata Storage Server No. 2	U4	PCIE-3 P2
U24	17B	Exadata Storage Server No. 1	U2	PCIE-3 P2

Ethernet Management Switch

The Ethernet management switch (Cisco Catalyst 4948 10 Gigabit Ethernet Switch, Figure 16) is located in the SPARC SuperCluster T4-4 at location U19.

The Ethernet management switch connects to the SPARC T4-4 servers, Sun ZFS Storage 7320 storage controllers, Exadata Storage Servers, and PDUs through the ports listed in the following tables.

Table 19 lists cables for the Ethernet Management switch in a full rack.

Table 20 lists cables for the Ethernet Management switch in a half rack.

Note - Ethernet management switch port numbers 45, 46, 47, and 48 are shared between item 4 (10BASE-T/100BASE-TX/1000BASE-T Ethernet) and item 6 (1000BASE-X Ethernet)

Note - Ethernet management switch ports 1, 2, 3, and 4 are not used at this time.

Ethernet Management Switch Connections in a Full Rack

TABLE 19 **Full Rack** Cat 5E Cables for Ethernet Management Switch (U19)

Ethernet Switch Port	To Device	Device Location	Device Port	Cable
1	—	—	—	—
2	—	—	—	—
3	—	—	—	—
4	—	—	—	—
5	T4-4 No. 4	U32	NET 3	Black, 3 m
6	T4-4 No. 3	U27	NET 3	Black, 3 m
7	T4-4 No. 4	U32	NET 2	Black, 3 m
8	T4-4 No. 3	U27	NET 2	Black, 3 m
9	T4-4 No. 4	U32	NET 1	Black, 3 m
10	T4-4 No. 3	U27	NET 1	Black, 3 m
11	T4-4 No. 4	U32	NET 0	Black, 3 m
12	T4-4 No. 4	U32	NET MGT	Red, 3 m
13	T4-4 No. 3	U27	NET-0	Black, 3 m
14	T4-4 No. 3	U27	NET MGT	Red, 3 m
15	PDU-A	PDU-A	NET MGT	White, 1 m
16	T4-4 No. 2	U13	NET-1	Black, 3 m

Ethernet Switch Port	To Device	Device Location	Device Port	Cable
17	T4-4 No. 2	U13	NET-0	Black, 3 m
18	T4-4 No. 2	U13	NET MGT	Red, 3 m
19	PDU-B	PDU-B	NET MGT	White, 1 m
20	T4-4 No. 2	U13	NET-2	Black, 3 m
21	T4-4 No. 1	U8	NET-0	Black, 3 m
22	T4-4 No. 1	U8	NET MGT	Red, 3 m
23	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-1	Blue, 3 m
24	T4-4 No. 2	U13	NET-3	Black, 3 m
25	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-0	Blue, 3 m
26	T4-4 No. 1	U8	NET-1	Black, 3 m
27	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-2	Blue, 3 m
28	T4-4 No. 1	U8	NET-2	Black, 3 m
29	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-1	Blue, 3 m
30	T4-4 No. 1	U8	NET-3	Black, 3 m
31	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-0	Blue, 3 m
32	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-2	Blue, 3 m
33	Exadata Storage Server No. 6	U41	NET-0	Black, 3 m
34	Exadata Storage Server No. 6	U41	NET MGT	Red, 3 m
35	Exadata Storage Server No. 5	U39	NET-0	Black, 3 m
36	Exadata Storage Server No. 5	U39	NET MGT	Red, 3 m
37	Exadata Storage Server No. 4	U37	NET-0	Black, 3 m
38	Exadata Storage Server No. 4	U37	NET MGT	Red, 3 m
39	Exadata Storage Server No. 3	U6	NET-0	Black, 3 m
40	Exadata Storage Server No. 3	U6	NET MGT	Red, 3 m
41	Exadata Storage Server No. 2	U4	NET-0	Black, 3 m
42	Exadata Storage Server No. 2	U4	NET MGT	Red, 3 m
43	Exadata Storage Server No. 1	U2	NET-0	Black, 3 m
44	Exadata Storage Server No. 1	U2	NET MGT	Red, 3 m
45	Sun Datacenter InfiniBand Switch 36 Leaf No. 2	U24	NET-0	Black, 3 m
46	Sun Datacenter InfiniBand Switch 36 Leaf No. 1	U18	NET-0	Black, 3 m
47	Sun Datacenter InfiniBand Switch 36 Spine	U1	NET-0	Black, 3 m
48	(Reserved for service access)	—	—	Blue, 3 m

Ethernet Management Switch Connections in a Half Rack

TABLE 20 **Half Rack** Cat 5E Cables for Ethernet Management Switch (U19)

Ethernet Switch Port	To Device	Device Location	Device Port	Cable
1	—	—	—	—
2	—	—	—	—
3	—	—	—	—
4	—	—	—	—
5	—	—	—	—
6	—	—	—	—
7	—	—	—	—
8	—	—	—	—
9	—	—	—	—
10	—	—	—	—
11	—	—	—	—
12	—	—	—	—
13	—	—	—	—
14	—	—	—	—
15	PDU-A	PDU-A	NET MGT	White, 1 m
16	T4-4 No. 2	U13	NET-1	Black, 3 m
17	T4-4 No. 2	U13	NET-0	Black, 3 m
18	T4-4 No. 2	U13	NET MGT	Red, 3 m
19	PDU-B	PDU-B	NET MGT	White, 1 m
20	T4-4 No. 2	U13	NET-2	Black, 3 m
21	T4-4 No. 1	U8	NET-0	Black, 3 m
22	T4-4 No. 1	U8	NET MGT	Red, 3 m
23	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-1	Blue, 3 m
24	T4-4 No. 2	U13	NET-3	Black, 3 m
25	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-0	Blue, 3 m
26	T4-4 No. 1	U8	NET-1	Black, 3 m
27	Sun ZFS Storage 7320 storage controller No. 2	U26	NET-2	Blue, 3 m
28	T4-4 No. 1	U8	NET-2	Black, 3 m
29	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-1	Blue, 3 m
30	T4-4 No. 1	U8	NET-3	Black, 3 m
31	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-0	Blue, 3 m
32	Sun ZFS Storage 7320 storage controller No. 1	U25	NET-2	Blue, 3 m
33	—	—	—	—

Ethernet Switch Port	To Device	Device Location	Device Port	Cable
34	—	—	—	—
35	—	—	—	—
36	—	—	—	—
37	—	—	—	—
38	—	—	—	—
39	Exadata Storage Server No. 3	U6	NET-0	Black, 3 m
40	Exadata Storage Server No. 3	U6	NET MGT	Red, 3 m
41	Exadata Storage Server No. 2	U4	NET-0	Black, 3 m
42	Exadata Storage Server No. 2	U4	NET MGT	Red, 3 m
43	Exadata Storage Server No. 1	U2	NET-0	Black, 3 m
44	Exadata Storage Server No. 1	U2	NET MGT	Red, 3 m
45	Sun Datacenter InfiniBand Switch 36 Leaf No. 2	U24	NET-0	Black, 3 m
46	Sun Datacenter InfiniBand Switch 36 Leaf No. 1	U18	NET-0	Black, 3 m
47	Sun Datacenter InfiniBand Switch 36 Spine	U1	NET-0	Black, 3 m
48	(Reserved for service access)	—	—	Blue, 3 m

Sun ZFS Storage 7320 Appliance

Table 21 shows the cabling for the Sun ZFS Storage 7320 appliance. The major components are:

- Sun ZFS Storage 7320 storage controller No. 1 (slot 25)
- Sun ZFS Storage 7320 storage controller No. 2 (slot 26)
- Sun Disk Shelf (slot 20)

TABLE 21 SAS Cable Connections for the Sun ZFS Storage 7320 appliance

From Storage Controller Slot	Storage Controller Port	To Slot	To Port	Cable
U25	PCIE-1 P0	U26	PCIE-1 P2	Yellow, 7 ft.
U25	PCIE-1 P1	U26	PCIE-1 P1	Black, 7 ft.
U25	PCIE-1 P2	U26	PCIE-1 P0	Green, 7 ft.
U26	PCIE-2 P1	U20	Left SAS Port 1	Black, 2 m
U25	PCIE-2 P1	U20	Left SAS Port 0	Black, 2 m
U26	PCIE-2 P0	U20	Right SAS Port 1	Black, 2 m
U25	PCIE-2 P0	U20	Right SAS Port 0	Black, 2 m

Single-Phase PDUs

If the rack uses single-phase AC power, connect the devices in the rack to the power distribution unit (PDU) as shown in Table 22. Each device requires two power cords for redundant power.

Caution - Turn off all PDU circuit breakers before connecting power cords to the PDU. Every socket group has a circuit breaker.

See Figure 17 for the locations of the PDU circuit breakers and for names of the PDU AC sockets.

TABLE 22 Single Phase PDU Power Cabling

Slot	**PDU-A/PSU-0** group-**socket**	**PDB-B/PDU-1** group-**socket**	**PDU-A/PSU-2** group-**socket**	**PDB-B/PSU-3** group-**socket**	**To Device**
U41	Group 5-4	Group 0-2	—	—	Exadata Storage Server No. 6
U39	Group 5-3	Group 0-3	—	—	Exadata Storage Server No. 5
U37	Group 5-2	Group 0-4	—	—	Exadata Storage Server No. 4
U32	Group 4-2	Group 1-4	Group 5-0	Group 0-6	SPARC T4-4 No. 4
U27	Group 3-6	Group 2-0	Group 4-1	Group 1-5	SPARC T4-4 No. 3
U26	Group 3-4	Group 2-2	—	—	Sun ZFS Storage 7320 storage controller No. 2
U25	Group 3-2	Group 2-4	—	—	Sun ZFS Storage 7320 storage controller No. 1
U24	Group 3-0	Group 2-6	—	—	Sun Datacenter InfiniBand Switch 36 Leaf Switch No. 2
U20	Group 2-6	Group 3-0	—	—	Sun Disk Shelf
U19	Group 2-5	Group 3-1	—	—	Ethernet management switch
U18	Group 2-3	Group 3-3	—	—	Sun Datacenter InfiniBand Switch 36 Leaf Switch No. 1
U13	Group 1-1	Group 4-5	Group 2-1	Group 3-5	SPARC T4-4 No. 2
U8	Group 0-6	Group 5-0	Group 1-0	Group 4-6	SPARC T4-4 No. 1
U6	Group 0-4	Group 5-2	—	—	Exadata Storage Server No. 3
U4	Group 0-3	Group 5-3	—	—	Exadata Storage Server No. 2
U2	Group 0-2	Group 5-4	—	—	Exadata Storage Server No. 1
U1	Group 0-1	Group 5-5	—	—	Sun Datacenter InfiniBand Switch 36 Spine Switch

Three-Phase PDUs

For three-phase AC power, connect the devices in the rack to the power distribution units (PDUs) as shown in Table 23 below. Each device requires two power cords for redundant power.

Caution - Turn off all PDU circuit breakers before connecting power cords to the PDU. Every socket group has a circuit breaker.

See Figure 17 (above) for the locations of the PDU circuit breakers and for names of the PDU AC sockets.

TABLE 23 Three Phase PDU Power Cabling

Slot	PDU-A/PSU-0 group-socket	PDB-B/PDU-1 group-socket	PDU-A/PSU-2 group-socket	PDB-B/PSU-3 group-socket	To Device
U41	Group 5-6	Group 2-0	—	—	Exadata Storage Server No. 6
U39	Group 5-4	Group 2-2	—	—	Exadata Storage Server No. 5
U37	Group 5-2	Group 2-4	—	—	Exadata Storage Server No. 4
U32	Group 4-2	Group 1-4	Group 5-0	Group 2-6	SPARC T4-4 No. 4
U27	Group 3-6	Group 0-0	Group 4-1	Group 1-5	SPARC T4-4 No. 3
U26	Group 3-4	Group 0-2	—	—	Sun ZFS Storage 7320 storage controller No. 2
U25	Group 3-2	Group 0-4	—	—	Sun ZFS Storage 7320 storage controller No. 1
U24	Group 3-0	Group 0-6	—	—	Sun Datacenter InfiniBand Switch 36 Leaf Switch No. 2
U20	Group 2-6	Group 5-0	—	—	Sun Disk Shelf
U19	Group 2-5	Group 5-1	—	—	Ethernet management switch
U18	Group 2-4	Group 5-2	—	—	Sun Datacenter InfiniBand Switch 36 Leaf Switch No. 1
U13	Group 1-1	Group 4-5	Group 2-0	Group 5-6	SPARC T4-4 No. 2
U8	Group 0-6	Group 3-0	Group 1-0	Group 4-6	SPARC T4-4 No. 1
U6	Group 0-4	Group 3-2	—	—	Exadata Storage Server No. 3
U4	Group 0-3	Group 3-3	—	—	Exadata Storage Server No. 2
U2	Group 0-2	Group 3-4	—	—	Exadata Storage Server No. 1
U1	Group 0-1	Group 3-5	—	—	Sun Datacenter InfiniBand Switch 36 Spine Switch

Connecting Multiple SPARC SuperCluster T4-4 Systems

These topics provide instructions for connecting one SPARC SuperCluster T4-4 to one or more SPARC SuperCluster T4-4 systems.

Multi-Rack Cabling Overview

Three Sun Datacenter InfiniBand Switch 36 switches are needed when cabling racks together. The SPARC SuperCluster T4-4 includes three Sun Datacenter InfiniBand Switch 36 switches. These switches attach to standard Quad Small Form-factor Pluggable (QSFP) connectors at the end of the InfiniBand cables. The procedures in this section assume the racks are adjacent to each other. If they are not, then longer cables may be required for the connections.

The switch at rack unit 1 (U1) is referred to as the spine switch. The switches at rack unit 18 (U18) and rack unit 24 (U24) are referred to as leaf switches. In a single rack, the two leaf switches are interconnected using seven connections. In addition, each leaf switch has one connection to the spine switch. The leaf switches connect to the spine switch as shown in the following graphic.

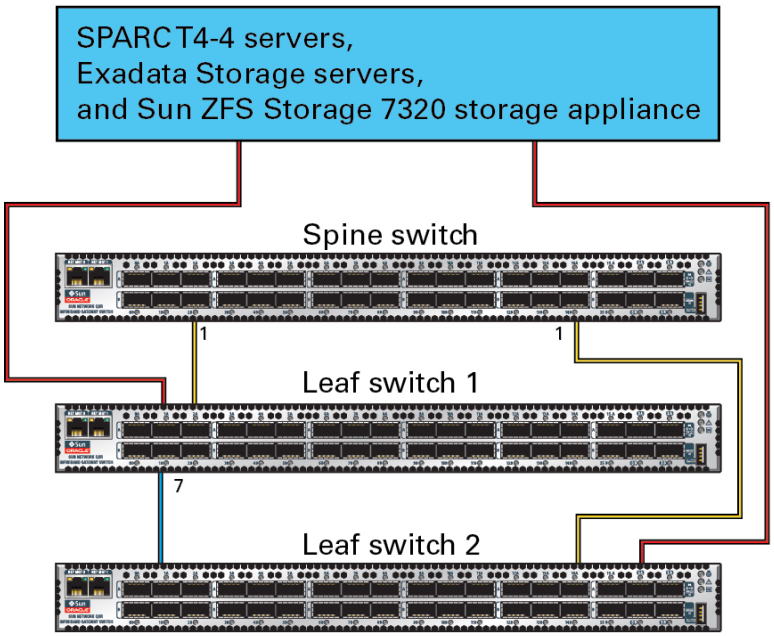

When connecting multiple racks together, remove the seven existing inter-switch connections between each leaf switches, as well as the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multi-rack environments, the leaf switches inside a rack are no longer directly interconnected, as shown in the following graphic.

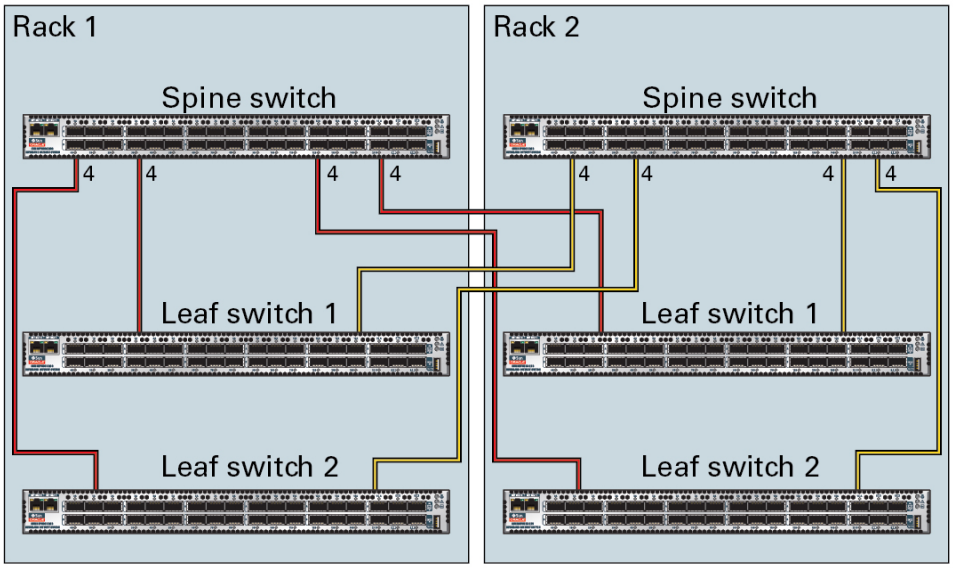

As shown in the preceding graphic, each leaf switch in rack 1 connects to the following switches:

- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2

The spine switch in rack 1 connects to the following switches:

- Eight connections to both internal leaf switches
- Eight connections to both leaf switches in rack 2

In the SPARC SuperCluster T4-4, the spine and leaf switches are installed in the following locations:

- Spine switch in U1
- Two leaf switches in U18 and U24

Two-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling two full racks together.

TABLE 24 Leaf Switch Connections for the First Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R1-U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	5 meters

Leaf Switch	Connection	Cable Length
	R1-U24-P8B to R1-U1-P4A	
	R1-U24-P9A to R1-U1-P5A	
	R1-U24-P9B to R1-U1-P6A	
R1-U24 to Rack 2	R1-U24-P10A to R2-U1-P7A	5 meters
	R1-U24-P10B to R2-U1-P8A	
	R1-U24-P11A to R2-U1-P9A	
	R1-U24-P11B to R2-U1-P10A	
R1-U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	5 meters
	R1-U18-P8B to R1-U1-P4B	
	R1-U18-P9A to R1-U1-P5B	
	R1-U18-P9B to R1-U1-P6B	
R1-U18 to Rack 2	R1-U18-P10A to R2-U1-P7B	5 meters
	R1-U18-P10B to R2-U1-P8B	
	R1-U18-P11A to R2-U1-P9B	
	R1-U18-P11B to R2-U1-P10B	

The following table shows the cable connections for the second spine switch (R2-U1) when cabling two full racks together.

TABLE 25 Leaf Switch Connections for the Second Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R2-U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	5 meters
	R2-U24-P8B to R2-U1-P4A	
	R2-U24-P9A to R2-U1-P5A	
	R2-U24-P9B to R2-U1-P6A	
R2-U24 to Rack 1	R2-U24-P10A to R1-U1-P7A	5 meters
	R2-U24-P10B to R1-U1-P8A	
	R2-U24-P11A to R1-U1-P9A	
	R2-U24-P11B to R1-U1-P10A	
R2-U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	5 meters
	R2-U18-P8B to R2-U1-P4B	
	R2-U18-P9A to R2-U1-P5B	
	R2-U18-P9B to R2-U1-P6B	
R2-U18 to Rack 1	R2-U18-P10A to R1-U1-P7B	5 meters

Leaf Switch	Connection	Cable Length
	R2-U18-P10B to R1-U1-P8B	
	R2-U18-P11A to R1-U1-P9B	
	R2-U18-P11B to R1-U1-P10B	

Three-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling three full racks together.

TABLE 26 Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1-U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	5 meters
	R1-U24-P8B to R1-U1-P4A	
	R1-U24-P9A to R1-U1-P5A	
R1-U24 to Rack 2	R1-U24-P9B to R2-U1-P6A	5 meters
	R1-U24-P10A to R2-U1-P7A	
	R1-U24-P10B to R2-U1-P8A	
R1-U24 to Rack 3	R1-U24-P11A to R3-U1-P9A	5 meters
	R1-U24-P11B to R3-U1-P10A	
R1-U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	5 meters
	R1-U18-P8B to R1-U1-P4B	
	R1-U18-P9A to R1-U1-P5B	
R1-U18 to Rack 2	R1-U18-P9B to R2-U1-P6B	5 meters
	R1-U18-P10A to R2-U1-P7B	
	R1-U18-P10B to R2-U1-P8B	
R1-U18 to Rack 3	R1-U18-P11A to R3-U1-P9B	5 meters
	R1-U18-P11B to R3-U1-P10B	

The following table shows the cable connections for the second spine switch (R2-U1) when cabling three racks together.

TABLE 27 Leaf Switch Connections for the Second Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2-U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	5 meters

Leaf Switch	Connection	Cable Length
	R2-U24-P8B to R2-U1-P4A	
	R2-U24-P9A to R2-U1-P5A	
R2-U24 to Rack 1	R2-U24-P11A to R1-U1-P9A	5 meters
	R2-U24-P11B to R1-U1-P10A	
R2-U24 to Rack 3	R2-U24-P9B to R3-U1-P6A	5 meters
	R2-U24-P10A to R3-U1-P7A	
	R2-U24-P10B to R3-U1-P8A	
R2-U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	5 meters
	R2-U18-P8B to R2-U1-P4B	
	R2-U18-P9A to R2-U1-P5B	
R2-U18 to Rack 1	R2-U18-P11A to R1-U1-P9B	5 meters
	R2-U18-P11B to R1-U1-P10B	
R2-U18 to Rack 3	R2-U18-P9B to R3-U1-P6B	5 meters
	R2-U18-P10A to R3-U1-P7B	
	R2-U18-P10B to R3-U1-P8B	

The following table shows the cable connections for the third spine switch (R3-U1) when cabling three full racks together.

TABLE 28 Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3-U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	5 meters
	R3-U24-P8B to R3-U1-P4A	
	R3-U24-P9A to R3-U1-P5A	
R3-U24 to Rack 1	R3-U24-P9B to R1-U1-P6A	5 meters
	R3-U24-P10A to R1-U1-P7A	
	R3-U24-P10B to R1-U1-P8A	
R3-U24 to Rack 2	R3-U24-P11A to R2-U1-P9A	5 meters
	R3-U24-P11B to R2-U1-P10A	
R3-U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	5 meters
	R3-U18-P8B to R3-U1-P4B	
	R3-U18-P9A to R3-U1-P5B	
R3-U18 to Rack 1	R3-U18-P9B to R1-U1-P6B	5 meters
	R3-U18-P10A to R1-U1-P7B	

Leaf Switch	Connection	Cable Length
R3-U18 to Rack 2	R3-U18-P10B to R1-U1-P8B R3-U18-P11A to R2-U1-P9B R3-U18-P11B to R2-U1-P10B	5 meters

Four-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling four full racks together.

TABLE 29 Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1-U24 within Rack 1	R1-U24-P8A to R1-U1-P3A R1-U24-P8B to R1-U1-P4A	5 meters
R1-U24 to Rack 2	R1-U24-P9A to R2-U1-P5A R1-U24-P9B to R2-U1-P6A	5 meters
R1-U24 to Rack 3	R1-U24-P10A to R3-U1-P7A R1-U24-P10B to R3-U1-P8A	5 meters
R1-U24 to Rack 4	R1-U24-P11A to R4-U1-P9A R1-U24-P11B to R4-U1-P10A	10 meters
R1-U18 within Rack 1	R1-U18-P8A to R1-U1-P3B R1-U18-P8B to R1-U1-P4B	5 meters
R1-U18 to Rack 2	R1-U18-P9A to R2-U1-P5B R1-U18-P9B to R2-U1-P6B	5 meters
R1-U18 to Rack 3	R1-U18-P10A to R3-U1-P7B R1-U18-P10B to R3-U1-P8B	5 meters
R1-U18 to Rack 4	R1-U18-P11A to R4-U1-P9B R1-U18-P11B to R4-U1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-U1) when cabling four full racks together.

TABLE 30 Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2-U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	5 meters

Leaf Switch	Connection	Cable Length
	R2-U24-P8B to R2-U1-P4A	
R2-U24 to Rack 1	R2-U24-P11A to R1-U1-P9A	5 meters
	R2-U24-P11B to R1-U1-P10A	
R2-U24 to Rack 3	R2-U24-P9A to R3-U1-P5A	5 meters
	R2-U24-P9B to R3-U1-P6A	
R2-U24 to Rack 4	R2-U24-P10A to R4-U1-P7A	5 meters
	R2-U24-P10B to R4-U1-P8A	
R2-U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	5 meters
	R2-U18-P8B to R2-U1-P4B	
R2-U18 to Rack 1	R2-U18-P11A to R1-U1-P9B	5 meters
	R2-U18-P11B to R1-U1-P10B	
R2-U18 to Rack 3	R2-U18-P9A to R3-U1-P5B	5 meters
	R2-U18-P9B to R3-U1-P6B	
R2-U18 to Rack 4	R2-U18-P10A to R4-U1-P7B	5 meters
	R2-U18-P10B to R4-U1-P8B	

The following table shows the cable connections for the third spine switch (R3-U1) when cabling four full racks together.

TABLE 31 Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R3-U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	5 meters
	R3-U24-P8B to R3-U1-P4A	
R3-U24 to Rack 1	R3-U24-P10A to R1-U1-P7A	5 meters
	R3-U24-P10B to R1-U1-P8A	
R3-U24 to Rack 2	R3-U24-P11A to R2-U1-P9A	5 meters
	R3-U24-P11B to R2-U1-P10A	
R3-U24 to Rack 4	R3-U24-P9A to R4-U1-P5A	5 meters
	R3-U24-P9B to R4-U1-P6A	
R3-U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	5 meters
	R3-U18-P8B to R3-U1-P4B	
R3-U18 to Rack 1	R3-U18-P10A to R1-U1-P7B	5 meters
	R3-U18-P10B to R1-U1-P8B	
R3-U18 to Rack 2	R3-U18-P11A to R2-U1-P9B	5 meters
	R3-U18-P11B to R2-U1-P10B	

Leaf Switch	Connection	Cable Length
R3-U18 to Rack 4	R3-U18-P9A to R4-U1-P5B	5 meters
	R3-U18-P9B to R4-U1-P6B	

The following table shows the cable connections for the fourth spine switch (R4-U1) when cabling four full racks together.

TABLE 32 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4-U24 within Rack 4	R4-U24-P8A to R4-U1-P3A	5 meters
	R4-U24-P8B to R4-U1-P4A	
R4-U24 to Rack 1	R4-U24-P9A to R1-U1-P5A	10 meters
	R4-U24-P9B to R1-U1-P6A	
R4-U24 to Rack 2	R4-U24-P10A to R2-U1-P7A	5 meters
	R4-U24-P10B to R2-U1-P8A	
R4-U24 to Rack 3	R4-U24-P11A to R3-U1-P9A	5 meters
	R4-U24-P11B to R3-U1-P10A	
R4-U18 within Rack 4	R4-U18-P8A to R4-U1-P3B	5 meters
	R4-U18-P8B to R4-U1-P4B	
R4-U18 to Rack 1	R4-U18-P9A to R1-U1-P5B	10 meters
	R4-U18-P9B to R1-U1-P6B	
R4-U18 to Rack 2	R4-U18-P10A to R2-U1-P7B	5 meters
	R4-U18-P10B to R2-U1-P8B	
R4-U18 to Rack 3	R4-U18-P11A to R3-U1-P9B	5 meters
	R4-U18-P11B to R3-U1-P10B	

Five-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling five full racks together.

TABLE 33 Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	3 meters
	R1-U24-P8B to R1-U1-P4A	

Leaf Switch	Connection	Cable Length
R1 U24 to Rack 2	R1-U24-P9A to R2-U1-P5A	5 meters
	R1-U24-P9B to R2-U1-P6A	
R1 U24 to Rack 3	R1-U24-P10A to R3-U1-P7A	5 meters
	R1-U24-P10B to R3-U1-P8A	
R1 U24 to Rack 4	R1-U24-P11A to R4-U1-P9A	10 meters
R1 U24 to Rack 5	R1-U24-P11B to R5-U1-P10A	10 meters
R1 U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	3 meters
	R1-U18-P8B to R1-U1-P4B	
R1 U18 to Rack 2	R1-U18-P9A to R2-U1-P5B	3 meters
	R1-U18-P9B to R2-U1-P6B	
R1 U18 to Rack 3	R1-U18-P10A to R3-U1-P7B	5 meters
	R1-U18-P10B to R3-U1-P8B	
R1 U18 to Rack 4	R1-U18-P11A to R4-U1-P9B	10 meters
R1 U18 to Rack 5	R1-U18-P11B to R5-U1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-U1) when cabling five full racks together.

TABLE 34 Leaf Switch Connections for the Second Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R2 U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	3 meters
	R2-U24-P8B to R2-U1-P4A	
R2 U24 to Rack 1	R2-U24-P11B to R1-U1-P10A	5 meters
R2 U24 to Rack 3	R2-U24-P9A to R3-U1-P5A	5 meters
	R2-U24-P9B to R3-U1-P6A	
R2 U24 to Rack 4	R2-U24-P10A to R4-U1-P7A	5 meters
	R2-U24-P10B to R4-U1-P8A	
R2 U24 to Rack 5	R2-U24-P11A to R5-U1-P9A	10 meters
R2 U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	3 meters
	R2-U18-P8B to R2-U1-P4B	
R2 U18 to Rack 1	R2-U18-P11B to R1-U1-P10B	5 meters
R2 U18 to Rack 3	R2-U18-P9A to R3-U1-P5B	5 meters
	R2-U18-P9B to R3-U1-P6B	
R2 U18 to Rack 4	R2-U18-P10A to R4-U1-P7B	5 meters
	R2-U18-P10B to R4-U1-P8B	
R2 U18 to Rack 5	R2-U18-P11A to R5-U1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-U1) when cabling five full racks together.

TABLE 35 Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	3 meters
	R3-U24-P8B to R3-U1-P4A	
R3 U24 to Rack 1	R3-U24-P11A to R1-U1-P9A	5 meters
R3 U24 to Rack 2	R3-U24-P11B to R2-U1-P10A	5 meters
R3 U24 to Rack 4	R3-U24-P9A to R4-U1-P5A	5 meters
	R3-U24-P9B to R4-U1-P6A	
R3 U24 to Rack 5	R3-U24-P10A to R5-U1-P7A	5 meters
	R3-U24-P10B to R5-U1-P8A	
R3 U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	3 meters
	R3-U18-P8B to R3-U1-P4B	
R3 U18 to Rack 1	R3-U18-P11A to R1-U1-P9B	5 meters
R3 U18 to Rack 2	R3-U18-P11B to R2-U1-P10B	5 meters
R3 U18 to Rack 4	R3-U18-P9A to R4-U1-P5B	5 meters
	R3-U18-P9B to R4-U1-P6B	
R3 U18 to Rack 5	R3-U18-P10A to R5-U1-P7B	5 meters
	R3-U18-P10B to R5-U1-P8B	

The following table shows the cable connections for the fourth spine switch (R4-U1) when cabling five full racks together.

TABLE 36 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 U24 within Rack 4	R4-U24-P8A to R4-U1-P3A	3 meters
	R4-U24-P8B to R4-U1-P4A	
R4 U24 to Rack 1	R4-U24-P10A to R1-U1-P7A	10 meters
	R4-U24-P10B to R1-U1-P8A	
R4 U24 to Rack 2	R4-U24-P11A to R2-U1-P9A	5 meters
R4 U24 to Rack 3	R4-U24-P11B to R3-U1-P10A	5 meters
R4 U24 to Rack 5	R4-U24-P9A to R5-U1-P5A	5 meters
	R4-U24-P9B to R5-U1-P6A	
R4 U18 within Rack 4	R4-U18-P8A to R4-U1-P3B	3 meters
	R4-U18-P8B to R4-U1-P4B	
R4 U18 to Rack 1	R4-U18-P10A to R1-U1-P7B	10 meters

Leaf Switch	Connection	Cable Length
	R4-U18-P10B to R1-U1-P8B	
R4 U18 to Rack 2	R4-U18-P11A to R2-U1-P9B	5 meters
R4 U18 to Rack 3	R4-U18-P11B to R3-U1-P10B	5 meters
R4 U18 to Rack 5	R4-U18-P9A to R5-U1-P5B	5 meters
	R4-U18-P9B to R5-U1-P6B	

The following table shows the cable connections for the fifth spine switch (R5-U1) when cabling five full racks together.

TABLE 37 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 U24 within Rack 5	R5-U24-P8A to R5-U1-P3A	3 meters
	R5-U24-P8B to R5-U1-P4A	
R5 U24 to Rack 1	R5-U24-P9A to R1-U1-P5A	10 meters
	R5-U24-P9B to R1-U1-P6A	
R5 U24 to Rack 2	R5-U24-P10A to R2-U1-P7A	10 meters
	R5-U24-P10B to R2-U1-P8A	
R5 U24 to Rack 3	R5-U24-P11A to R3-U1-P9A	5 meters
R5 U24 to Rack 4	R5-U24-P11B to R4-U1-P10A	5 meters
R5 U18 within Rack 5	R5-U18-P8A to R5-U1-P3B	3 meters
	R5-U18-P8B to R5-U1-P4B	
R5 U18 to Rack 1	R5-U18-P9A to R1-U1-P5B	10 meters
	R5-U18-P9B to R1-U1-P6B	
R5 U18 to Rack 2	R5-U18-P10A to R2-U1-P7B	10 meters
	R5-U24-P10B to R2-U1-P8B	
R5 U18 to Rack 3	R5-U18-P11A to R3-U1-P9B	5 meters
R5 U18 to Rack 4	R5-U18-P11B to R4-U1-P10B	5 meters

Six-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling six full racks together.

TABLE 38 Leaf Switch Connections for the First Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	3 meters

Leaf Switch	Connection	Cable Length
R1 U24 to Rack 2	R1-U24-P8B to R1-U1-P4A R1-U24-P9A to R2-U1-P5A R1-U24-P9B to R2-U1-P6A	5 meters
R1 U24 to Rack 3	R1-U24-P10A to R3-U1-P7A	5 meters
R1 U24 to Rack 4	R1-U24-P10B to R4-U1-P8A	10 meters
R1 U24 to Rack 5	R1-U24-P11A to R5-U1-P9A	10 meters
R1 U24 to Rack 6	R1-U24-P11B to R6-U1-P10A	10 meters
R1 U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	3 meters
R1 U18 to Rack 2	R1-U18-P8B to R1-U1-P4B R1-U18-P9A to R2-U1-P5B R1-U18-P9B to R2-U1-P6B	5 meters
R1 U18 to Rack 3	R1-U18-P10A to R3-U1-P7B	5 meters
R1 U18 to Rack 4	R1-U18-P10B to R4-U1-P8B	10 meters
R1 U18 to Rack 5	R1-U18-P11A to R5-U1-P9B	10 meters
R1 U18 to Rack 6	R1-U18-P11B to R6-U1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-U1) when cabling six full racks together.

TABLE 39 Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 U24 within Rack 2	R2-U24-P8A to R2-U1-P3A R2-U24-P8B to R2-U1-P4A	3 meters
R2 U24 to Rack 1	R2-U24-P11B to R1-U1-P10A	5 meters
R2 U24 to Rack 3	R2-U24-P9A to R3-U1-P5A R2-U24-P9B to R3-U1-P6A	5 meters
R2 U24 to Rack 4	R2-U24-P10A to R4-U1-P7A	5 meters
R2 U24 to Rack 5	R2-U24-P10B to R5-U1-P8A	10 meters
R2 U24 to Rack 6	R2-U24-P11Ato R6-U1-P9A	10 meters
R2 U18 within Rack 2	R2-U18-P8A to R2-U1-P3B R2-U18-P8B to R2-U1-P4B	3 meters
R2 U18 to Rack 1	R2-U18-P11B to R1-U1-P10B	5 meters
R2 U18 to Rack 3	R2-U18-P9A to R3-U1-P5B R2-U18-P9B to R3-U1-P6B	5 meters
R2 U18 to Rack 4	R2-U18-P10A to R4-U1-P7B	5 meters
R2 U18 to Rack 5	R2-U18-P10B to R5-U1-P8B	10 meters
R2 U18 to Rack 6	R2-U18-P11Ato R6-U1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-U1) when cabling six full racks together.

TABLE 40 Leaf Switch Connections for the Third Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R3 U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	3 meters
	R3-U24-P8B to R3-U1-P4A	
R3 U24 to Rack 1	R3-U24-P11A to R1-U1-P9A	5 meters
R3 U24 to Rack 2	R3-U24-P11B to R2-U1-P10A	5 meters
R3 U24 to Rack 4	R3-U24-P9A to R4-U1-P5A	5 meters
	R3-U24-P9B to R4-U1-P6A	
R3 U24 to Rack 5	R3-U24-P10A to R5-U1-P7A	5 meters
R3 U24 to Rack 6	R3-U24-P10B to R6-U1-P8A	5 meters
R3 U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	3 meters
	R3-U18-P8B to R3-U1-P4B	
R3 U18 to Rack 1	R3-U18-P11A to R1-U1-P9B	5 meters
R3 U18 to Rack 2	R3-U18-P11B to R2-U1-P10B	5 meters
R3 U18 to Rack 4	R3-U18-P9A to R4-U1-P5B	5 meters
	R3-U18-P9B to R4-U1-P6B	
R3 U18 to Rack 5	R3-U18-P10A to R5-U1-P7B	5 meters
R3 U18 to Rack 6	R3-U18-P10B to R6-U1-P8B	5 meters

The following table shows the cable connections for the fourth spine switch (R4-U1) when cabling six full racks together.

TABLE 41 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 U24 within Rack 4	R4-U24-P8A to R4-U1-P3A	3 meters
	R4-U24-P8B to R4-U1-P4A	
R4 U24 to Rack 1	R4-U24-P10B to R1-U1-P8A	10 meters
R4 U24 to Rack 2	R4-U24-P11A to R2-U1-P9A	5 meters
R4 U24 to Rack 3	R4-U24-P11B to R3-U1-P10A	5 meters
R4 U24 to Rack 5	R4-U24-P9A to R5-U1-P5A	5 meters
	R4-U24-P9B to R5-U1-P6A	
R4 U24 to Rack 6	R4-U24-P10A to R6-U1-P7A	5 meters
R4 U18 within Rack 4	R4-U18-P8A to R4-U1-P3B	3 meters
	R4-U18-P8B to R4-U1-P4B	

Leaf Switch	Connection	Cable Length
R4 U18 to Rack 1	R4-U18-P10B to R1-U1-P8B	10 meters
R4 U18 to Rack 2	R4-U18-P11A to R2-U1-P9B	5 meters
R4 U18 to Rack 3	R4-U18-P11B to R3-U1-P10B	5 meters
R4 U18 to Rack 5	R4-U18-P9A to R5-U1-P5B	5 meters
	R4-U18-P9B to R5-U1-P6B	
R4 U18 to Rack 6	R4-U18-P10A to R6-U1-P7B	5 meters

The following table shows the cable connections for the fifth spine switch (R5-U1) when cabling six full racks together.

TABLE 42 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 U24 within Rack 5	R5-U24-P8A to R5-U1-P3A	3 meters
	R5-U24-P8B to R5-U1-P4A	
R5 U24 to Rack 1	R5-U24-P10A to R1-U1-P7A	10 meters
R5 U24 to Rack 2	R5-U24-P10B to R2-U1-P8A	10 meters
R5 U24 to Rack 3	R5-U24-P11A to R3-U1-P9A	5 meters
R5 U24 to Rack 4	R5-U24-P11B to R4-U1-P10A	5 meters
R5 U24 to Rack 6	R5-U24-P9A to R6-U1-P5A	5 meters
	R5-U24-P9B to R6-U1-P6A	
R5 U18 within Rack 5	R5-U18-P8A to R5-U1-P3B	3 meters
	R5-U18-P8B to R5-U1-P4B	
R5 U18 to Rack 1	R5-U18-P10A to R1-U1-P7B	10 meters
R5 U18 to Rack 2	R5-U24-P10B to R2-U1-P8B	10 meters
R5 U18 to Rack 3	R5-U18-P11A to R3-U1-P9B	5 meters
R5 U18 to Rack 4	R5-U18-P11B to R4-U1-P10B	5 meters
R5 U18 to Rack 6	R5-U18-P9A to R6-U1-P5B	5 meters
	R5-U18-P9B to R6-U1-P6B	

The following table shows the cable connections for the sixth spine switch (R6-U1) when cabling six full racks together.

TABLE 43 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 U24 within Rack 6	R6-U24-P8A to R6-U1-P3A	3 meters
	R6-U24-P8B to R6-U1-P4A	

Leaf Switch	Connection	Cable Length
R6 U24 to Rack 1	R6-U24-P9A to R1-U1-P5A	10 meters
	R6-U24-P9B to R1-U1-P6A	
R6 U24 to Rack 2	R6-U24-P10A to R2-U1-P7A	10 meters
R6 U24 to Rack 3	R6-U24-P10B to R3-U1-P8A	5 meters
R6 U24 to Rack 4	R6-U24-P11A to R4-U1-P9A	5 meters
R6 U24 to Rack 5	R6-U24-P11B to R5-U1-P10A	5 meters
R6 U18 within Rack 6	R6-U18-P8A to R6-U1-P3B	3 meters
	R6-U18-P8B to R6-U1-P4B	
R6 U18 to Rack 2	R6-U24-P10A to R2-U1-P7B	10 meters
R6 U18 to Rack 1	R6-U18-P9A to R1-U1-P5B	10 meters
	R6-U18-P9B to R1-U1-P6B	
R6 U18 to Rack 3	R6-U18-P10B to R3-U1-P8B	5 meters
R6 U18 to Rack 4	R6-U18-P11A to R4-U1-P9B	5 meters
R6 U18 to Rack 5	R6-U18-P11B to R5-U1-P10B	5 meters

Seven-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling seven full racks together.

TABLE 44 Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	3 meters
	R1-U24-P8B to R1-U1-P4A	
R1 U24 to Rack 2	R1-U24-P9A to R2-U1-P5A	5 meters
R1 U24 to Rack 3	R1-U24-P9B to R3-U1-P6A	5 meters
R1 U24 to Rack 4	R1-U24-P10A to R4-U1-P7A	10 meters
R1 U24 to Rack 5	R1-U24-P10B to R5-U1-P8A	10 meters
R1 U24 to Rack 6	R1-U24-P11A to R6-U1-P9A	10 meters
R1 U24 to Rack 7	R1-U24-P11B to R7-U1-P10A	10 meters
R1 U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	3 meters
	R1-U18-P8B to R1-U1-P4B	
R1 U18 to Rack 2	R1-U18-P9A to R2-U1-P5B	5 meters
R1 U18 to Rack 3	R1-U18-P9B to R3-U1-P6B	5 meters
R1 U18 to Rack 4	R1-U18-P10A to R4-U1-P7B	10 meters
R1 U18 to Rack 5	R1-U18-P10B to R5-U1-P8B	10 meters
R1 U18 to Rack 6	R1-U18-P11A to R6-U1-P9B	10 meters

Leaf Switch	Connection	Cable Length
R1 U18 to Rack 7	R1-U18-P11B to R7-U1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-U1) when cabling seven full racks together.

TABLE 45 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	3 meters
	R2-U24-P8B to R2-U1-P4A	
R2 U24 to Rack 1	R2-U24-P11B to R1-U1-P10A	5 meters
R2 U24 to Rack 3	R2-U24-P9A to R3-U1-P5A	5 meters
R2 U24 to Rack 4	R2-U24-P9B to R4-U1-P6A	5 meters
R2 U24 to Rack 5	R2-U24-P10A to R5-U1-P7A	10 meters
R2 U24 to Rack 6	R2-U24-P10B to R6-U1-P8A	10 meters
R2 U24 to Rack 7	R2-U24-P11A to R7-U1-P9A	10 meters
R2 U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	3 meters
	R2-U18-P8B to R2-U1-P4B	
R2 U18 to Rack 1	R2-U18-P11B to R1-U1-P10B	5 meters
R2 U18 to Rack 3	R2-U18-P9A to R3-U1-P5B	5 meters
R2 U18 to Rack 4	R2-U18-P9B to R4-U1-P6B	5 meters
R2 U18 to Rack 5	R2-U18-P10A to R5-U1-P7B	10 meters
R2 U18 to Rack 6	R2-U18-P10Bto R6-U1-P8B	10 meters
R2 U18 to Rack 7	R2-U18-P11A to R7-U1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-U1) when cabling seven full racks together.

TABLE 46 Leaf Switch Connections for the Third Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	3 meters
	R3-U24-P8B to R3-U1-P4A	
R3 U24 to Rack 1	R3-U24-P11A to R1-U1-P9A	5 meters
R3 U24 to Rack 2	R3-U24-P11B to R2-U1-P10A	5 meters
R3 U24 to Rack 4	R3-U24-P9A to R4-U1-P5A	5 meters
R3 U24 to Rack 5	R3-U24-P9B to R5-U1-P6A	5 meters
R3 U24 to Rack 6	R3-U24-P10A to R6-U1-P7A	10 meters
R3 U24 to Rack 7	R3-U24-P10B to R7-U1-P8A	10 meters
R3 U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	3 meters

Leaf Switch	Connection	Cable Length
	R3-U18-P8B to R3-U1-P4B	
R3 U18 to Rack 1	R3-U18-P11A to R1-U1-P9B	5 meters
R3 U18 to Rack 2	R3-U18-P11B to R2-U1-P10B	5 meters
R3 U18 to Rack 4	R3-U18-P9A to R4-U1-P5B	5 meters
R3 U18 to Rack 5	R3-U18-P9B to R5-U1-P6B	5 meters
R3 U18 to Rack 6	R3-U18-P10A to R6-U1-P7B	10 meters
R3 U18 to Rack 7	R3-U18-P10B to R7-U1-P8B	10 meters

The following table shows the cable connections for the fourth spine switch (R4-U1) when cabling seven full racks together.

TABLE 47 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 U24 within Rack 4	R4-U24-P8A to R4-U1-P3A	3 meters
	R4-U24-P8B to R4-U1-P4A	
R4 U24 to Rack 1	R4-U24-P10B to R1-U1-P8A	10 meters
R4 U24 to Rack 2	R4-U24-P11A to R2-U1-P9A	5 meters
R4 U24 to Rack 3	R4-U24-P11B to R3-U1-P10A	5 meters
R4 U24 to Rack 5	R4-U24-P9A to R5-U1-P5A	5 meters
R4 U24 to Rack 6	R4-U24-P9B to R6-U1-P6A	5 meters
R4 U24 to Rack 7	R4-U24-P10A to R7-U1-P7A	10 meters
R4 U18 within Rack 4	R4-U18-P8A to R4-U1-P3B	3 meters
	R4-U18-P8B to R4-U1-P4B	
R4 U18 to Rack 1	R4-U18-P10B to R1-U1-P8B	10 meters
R4 U18 to Rack 2	R4-U18-P11A to R2-U1-P9B	5 meters
R4 U18 to Rack 3	R4-U18-P11B to R3-U1-P10B	5 meters
R4 U18 to Rack 5	R4-U18-P9A to R5-U1-P5B	5 meters
R4 U18 to Rack 6	R4-U18-P9B to R6-U1-P6B	5 meters
R4 U18 to Rack 7	R4-U18-P10A to R7-U1-P7B	10 meters

The following table shows the cable connections for the fifth spine switch (R5-U1) when cabling seven full racks together.

TABLE 48 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 U24 within Rack 5	R5-U24-P8A to R5-U1-P3A	3 meters
	R5-U24-P8B to R5-U1-P4A	
R5 U24 to Rack 1	R5-U24-P10A to R1-U1-P7A	10 meters

Leaf Switch	Connection	Cable Length
R5 U24 to Rack 2	R5-U24-P10B to R2-U1-P8A	10 meters
R5 U24 to Rack 3	R5-U24-P11A to R3-U1-P9A	5 meters
R5 U24 to Rack 4	R5-U24-P11B to R4-U1-P10A	5 meters
R5 U24 to Rack 6	R5-U24-P9A to R6-U1-P5A	5 meters
R5 U24 to Rack 7	R5-U24-P9B to R7-U1-P6A	5 meters
R5 U18 within Rack 5	R5-U18-P8A to R5-U1-P3B	3 meters
	R5-U18-P8B to R5-U1-P4B	
R5 U18 to Rack 1	R5-U18-P10A to R1-U1-P7B	10 meters
R5 U18 to Rack 2	R5-U24-P10B to R2-U1-P8B	10 meters
R5 U18 to Rack 3	R5-U18-P11A to R3-U1-P9B	5 meters
R5 U18 to Rack 4	R5-U18-P11B to R4-U1-P10B	5 meters
R5 U18 to Rack 6	R5-U18-P9A to R6-U1-P5B	5 meters
R5 U18 to Rack 7	R5-U18-P9B to R7-U1-P6B	5 meters

The following table shows the cable connections for the sixth spine switch (R6-U1) when cabling seven full racks together.

TABLE 49 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 U24 within Rack 6	R6-U24-P8A to R6-U1-P3A	3 meters
	R6-U24-P8B to R6-U1-P4A	
R6 U24 to Rack 1	R6-U24-P9B to R1-U1-P6A	10 meters
R6 U24 to Rack 2	R6-U24-P10A to R2-U1-P7A	10 meters
R6 U24 to Rack 3	R6-U24-P10B to R3-U1-P8A	5 meters
R6 U24 to Rack 4	R6-U24-P11A to R4-U1-P9A	5 meters
R6 U24 to Rack 5	R6-U24-P11B to R5-U1-P10A	5 meters
R6 U24 to Rack 7	R6-U24-P9A to R7-U1-P5A	5 meters
R6 U18 within Rack 6	R6-U18-P8A to R6-U1-P3B	3 meters
	R6-U18-P8B to R6-U1-P4B	
R6 U18 to Rack 1	R6-U18-P9B to R1-U1-P6B	10 meters
R6 U18 to Rack 2	R6-U24-P10A to R2-U1-P7B	10 meters
R6 U18 to Rack 3	R6-U18-P10B to R3-U1-P8B	5 meters
R6 U18 to Rack 4	R6-U18-P11A to R4-U1-P9B	5 meters
R6 U18 to Rack 5	R6-U18-P11B to R5-U1-P10B	5 meters
R6 U18 to Rack 7	R6-U24-P9A to R7-U1-P5B	5 meters

The following table shows the cable connections for the seventh spine switch (R7-U1) when cabling seven full racks together.

TABLE 50 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 U24 within Rack 7	R7-U24-P8A to R7-U1-P3A	3 meters
	R7-U24-P8B to R7-U1-P4A	
R7 U24 to Rack 1	R7-U24-P9A to R1-U1-P5A	10 meters
R7 U24 to Rack 2	R7-U24-P9B to R2-U1-P6A	10 meters
R7 U24 to Rack 3	R7-U24-P10A to R3-U1-P7A	10 meters
R7 U24 to Rack 4	R7-U24-P10B to R4-U1-P8A	10 meters
R7 U24 to Rack 5	R7-U24-P11A to R5-U1-P9A	5 meters
R7 U24 to Rack 6	R7-U24-P11B to R6-U1-P10A	5 meters
R7 U18 within Rack 7	R7-U18-P8A to R7-U1-P3B	3 meters
	R7-U18-P8B to R7-U1-P4B	
R7 U18 to Rack 1	R7-U18-P9A to R1-U1-P5B	10 meters
R7 U18 to Rack 2	R7-U24-P9B to R2-U1-P6B	10 meters
R7 U18 to Rack 3	R7-U18-P10A to R3-U1-P7B	10 meters
R7 U18 to Rack 4	R7-U18-P10B to R4-U1-P8B	10 meters
R7 U18 to Rack 5	R7-U18-P11A to R5-U1-P9B	5 meters
R7 U18 to Rack 6	R7-U24-P11B to R6-U1-P10B	5 meters

Eight-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-U1) when cabling eight full racks together.

TABLE 51 Leaf Switch Connections for the First Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 U24 within Rack 1	R1-U24-P8A to R1-U1-P3A	3 meters
R1 U24 to Rack 2	R1-U24-P8B to R2-U1-P4A	5 meters
R1 U24 to Rack 3	R1-U24-P9A to R3-U1-P5A	5 meters
R1 U24 to Rack 4	R1-U24-P9B to R4-U1-P6A	10 meters
R1 U24 to Rack 5	R1-U24-P10A to R5-U1-P7A	10 meters
R1 U24 to Rack 6	R1-U24-P10B to R6-U1-P8A	10 meters
R1 U24 to Rack 7	R1-U24-P11A to R7-U1-P9A	10 meters
R1 U24 to Rack 8	R1-U24-P11B to R8-U1-P10A	10 meters
R1 U18 within Rack 1	R1-U18-P8A to R1-U1-P3B	3 meters
R1 U18 to Rack 2	R1-U18-P8B to R2-U1-P4B	5 meters
R1 U18 to Rack 3	R1-U18-P9A to R3-U1-P5B	5 meters
R1 U18 to Rack 4	R1-U18-P9B to R4-U1-P6B	10 meters

Leaf Switch	Connection	Cable Length
R1 U18 to Rack 5	R1-U18-P10A to R5-U1-P7B	10 meters
R1 U18 to Rack 6	R1-U18-P10B to R6-U1-P8B	10 meters
R1 U18 to Rack 7	R1-U18-P11A to R7-U1-P8B	10 meters
R1 U18 to Rack 8	R1-U18-P11B to R8-U1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-U1) when cabling eight full racks together.

TABLE 52 Leaf Switch Connections for the Second Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R2 U24 within Rack 2	R2-U24-P8A to R2-U1-P3A	3 meters
R2 U24 to Rack 1	R2-U24-P11B to R1-U1-P10A	5 meters
R2 U24 to Rack 3	R2-U24-P8B to R3-U1-P4A	5 meters
R2 U24 to Rack 4	R2-U24-P9A to R4-U1-P5A	5 meters
R2 U24 to Rack 5	R2-U24-P9B to R5-U1-P6A	10 meters
R2 U24 to Rack 6	R2-U24-P10A to R6-U1-P7A	10 meters
R2 U24 to Rack 7	R2-U24-P10B to R7-U1-P8A	10 meters
R2 U24 to Rack 8	R2-U24-P11A to R8-U1-P9A	10 meters
R2 U18 within Rack 2	R2-U18-P8A to R2-U1-P3B	3 meters
R2 U18 to Rack 1	R2-U18-P11B to R1-U1-P10B	5 meters
R2 U18 to Rack 3	R2-U18-P8B to R3-U1-P4B	5 meters
R2 U18 to Rack 4	R2-U18-P9A to R4-U1-P5B	5 meters
R2 U18 to Rack 5	R2-U18-P9B to R5-U1-P6B	10 meters
R2 U18 to Rack 6	R2-U18-P10A to R6-U1-P7B	10 meters
R2 U18 to Rack 7	R2-U18-P10B to R7-U1-P8B	10 meters
R2 U18 to Rack 8	R2-U18-P11A to R8-U1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-U1) when cabling eight full racks together.

TABLE 53 Leaf Switch Connections for the Third Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 U24 within Rack 3	R3-U24-P8A to R3-U1-P3A	3 meters
R3 U24 to Rack 1	R3-U24-P11A to R1-U1-P9A	5 meters
R3 U24 to Rack 2	R3-U24-P11B to R2-U1-P10A	5 meters
R3 U24 to Rack 4	R3-U24-P8B to R4-U1-P4A	5 meters
R3 U24 to Rack 5	R3-U24-P9A to R5-U1-P5A	5 meters
R3 U24 to Rack 6	R3-U24-P9B to R6-U1-P6A	5 meters
R3 U24 to Rack 7	R3-U24-P10A to R7-U1-P7A	10 meters

Leaf Switch	Connection	Cable Length
R3 U24 to Rack 8	R3-U24-P10B to R8-U1-P8A	10 meters
R3 U18 within Rack 3	R3-U18-P8A to R3-U1-P3B	3 meters
R3 U18 to Rack 1	R3-U18-P11A to R1-U1-P9B	5 meters
R3 U18 to Rack 2	R3-U18-P11B to R2-U1-P10B	5 meters
R3 U18 to Rack 4	R3-U18-P8B to R4-U1-P4B	5 meters
R3 U18 to Rack 5	R3-U18-P9A to R5-U1-P5B	5 meters
R3 U18 to Rack 6	R3-U18-P9B to R6-U1-P6B	5 meters
R3 U18 to Rack 7	R3-U18-P10A to R7-U1-P7B	10 meters
R3 U18 to Rack 8	R3-U18-P10B to R8-U1-P8B	10 meters

The following table shows the cable connections for the fourth spine switch (R4-U1) when cabling eight full racks together.

TABLE 54 Leaf Switch Connections for the Fourth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 U24 within Rack 4	R4-U24-P8A to R4-U1-P3A	3 meters
R4 U24 to Rack 1	R4-U24-P10B to R1-U1-P8A	10 meters
R4 U24 to Rack 2	R4-U24-P11A to R2-U1-P9A	5 meters
R4 U24 to Rack 3	R4-U24-P11B to R3-U1-P10A	5 meters
R4 U24 to Rack 5	R4-U24-P8B to R5-U1-P4A	5 meters
R4 U24 to Rack 6	R4-U24-P9A to R6-U1-P5A	5 meters
R4 U24 to Rack 7	R4-U24-P9B to R7-U1-P6A	10 meters
R4 U24 to Rack 8	R4-U24-P10A to R8-U1-P7A	10 meters
R4 U18 within Rack 4	R4-U18-P8A to R4-U1-P3B	3 meters
R4 U18 to Rack 1	R4-U18-P10B to R1-U1-P8B	10 meters
R4 U18 to Rack 2	R4-U18-P11A to R2-U1-P9B	5 meters
R4 U18 to Rack 3	R4-U18-P11B to R3-U1-P10B	5 meters
R4 U18 to Rack 5	R4-U18-P8B to R5-U1-P4B	5 meters
R4 U18 to Rack 6	R4-U18-P9A to R6-U1-P5B	5 meters
R4 U18 to Rack 7	R4-U18-P9B to R7-U1-P6B	10 meters
R4 U18 to Rack 8	R4-U18-P10A to R8-U1-P7B	10 meters

The following table shows the cable connections for the fifth spine switch (R5-U1) when cabling eight full racks together.

TABLE 55 Leaf Switch Connections for the Fifth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R5 U24 within Rack 5	R5-U24-P8A to R5-U1-P3A	3 meters
R5 U24 to Rack 1	R5-U24-P10A to R1-U1-P7A	10 meters

Leaf Switch	Connection	Cable Length
R5 U24 to Rack 2	R5-U24-P10B to R2-U1-P8A	10 meters
R5 U24 to Rack 3	R5-U24-P11A to R3-U1-P9A	5 meters
R5 U24 to Rack 4	R5-U24-P11B to R4-U1-P10A	5 meters
R5 U24 to Rack 6	R5-U24-P8B to R6-U1-P4A	5 meters
R5 U24 to Rack 7	R5-U24-P9A to R7-U1-P5A	5 meters
R5 U24 to Rack 8	R5-U24-P9B to R8-U1-P6A	10 meters
R5 U18 within Rack 5	R5-U18-P8A to R5-U1-P3B	3 meters
R5 U18 to Rack 1	R5-U18-P10A to R1-U1-P7B	10 meters
R5 U18 to Rack 2	R5-U24-P10B to R2-U1-P8B	10 meters
R5 U18 to Rack 3	R5-U18-P11A to R3-U1-P9B	5 meters
R5 U18 to Rack 4	R5-U18-P11B to R4-U1-P10B	5 meters
R5 U18 to Rack 6	R5-U18-P8B to R6-U1-P4B	5 meters
R5 U18 to Rack 7	R5-U18-P9A to R7-U1-P5B	5 meters
R5 U18 to Rack 8	R5-U18-P9B to R8-U1-P6B	10 meters

The following table shows the cable connections for the sixth spine switch (R6-U1) when cabling eight full racks together.

TABLE 56 Leaf Switch Connections for the Sixth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 U24 within Rack 6	R6-U24-P8A to R6-U1-P3A	3 meters
R6 U24 to Rack 1	R6-U24-P9B to R1-U1-P6A	10 meters
R6 U24 to Rack 2	R6-U24-P10A to R2-U1-P7A	10 meters
R6 U24 to Rack 3	R6-U24-P10B to R3-U1-P8A	5 meters
R6 U24 to Rack 4	R6-U24-P11A to R4-U1-P9A	5 meters
R6 U24 to Rack 5	R6-U24-P11B to R5-U1-P10A	5 meters
R6 U24 to Rack 7	R6-U24-P8B to R7-U1-P4A	5 meters
R6 U24 to Rack 8	R6-U24-P9A to R8-U1-P5A	5 meters
R6 U18 within Rack 6	R6-U18-P8A to R6-U1-P3B	3 meters
R6 U18 to Rack 1	R6-U18-P9B to R1-U1-P6B	10 meters
R6 U18 to Rack 2	R6-U24-P10A to R2-U1-P7B	10 meters
R6 U18 to Rack 3	R6-U18-P10B to R3-U1-P8B	5 meters
R6 U18 to Rack 4	R6-U18-P11A to R4-U1-P9B	5 meters
R6 U18 to Rack 5	R6-U18-P11B to R5-U1-P10B	5 meters
R6 U18 to Rack 7	R6-U24-P8B to R7-U1-P4B	5 meters
R6 U18 to Rack 8	R6-U18-P9A to R8-U1-P5B	5 meters

The following table shows the cable connections for the seventh spine switch (R7-U1) when cabling eight full racks together.

TABLE 57 Leaf Switch Connections for the Seventh Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 U24 within Rack 7	R7-U24-P8A to R7-U1-P3A	3 meters
R7 U24 to Rack 1	R7-U24-P9A to R1-U1-P5A	10 meters
R7 U24 to Rack 2	R7-U24-P9B to R2-U1-P6A	10 meters
R7 U24 to Rack 3	R7-U24-P10A to R3-U1-P7A	10 meters
R7 U24 to Rack 4	R7-U24-P10B to R4-U1-P8A	10 meters
R7 U24 to Rack 5	R7-U24-P11A to R5-U1-P9A	5 meters
R7 U24 to Rack 6	R7-U24-P11B to R6-U1-P10A	5 meters
R7 U24 to Rack 8	R7-U24-P8B to R8-U1-P4A	5 meters
R7 U18 within Rack 7	R7-U18-P8A to R7-U1-P3B	3 meters
R7 U18 to Rack 1	R7-U18-P9A to R1-U1-P5B	10 meters
R7 U18 to Rack 2	R7-U24-P9B to R2-U1-P6B	10 meters
R7 U18 to Rack 3	R7-U18-P10A to R3-U1-P7B	10 meters
R7 U18 to Rack 4	R7-U18-P10B to R4-U1-P8B	10 meters
R7 U18 to Rack 5	R7-U18-P11A to R5-U1-P9B	5 meters
R7 U18 to Rack 6	R7-U24-P11B to R6-U1-P10B	5 meters
R7 U18 to Rack 8	R7-U24-P8B to R8-U1-P4B	5 meters

The following table shows the cable connections for the eighth spine switch (R8-U1) when cabling eight full racks together.

TABLE 58 Leaf Switch Connections for the Eighth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 U24 within Rack 8	R8-U24-P8A to R8-U1-P3A	3 meters
R8 U24 to Rack 1	R8-U24-P8B to R1-U1-P4A	10 meters
R8 U24 to Rack 2	R8-U24-P9A to R2-U1-P5A	10 meters
R8 U24 to Rack 3	R8-U24-P9B to R3-U1-P6A	10 meters
R8 U24 to Rack 4	R8-U24-P10A to R4-U1-P7A	10 meters
R8 U24 to Rack 5	R8-U24-P10B to R5-U1-P8A	5 meters
R8 U24 to Rack 6	R8-U24-P11A to R6-U1-P9A	5 meters
R8 U24 to Rack 7	R8-U24-P11B to R7-U1-P10A	5 meters
R8 U18 within Rack 8	R8-U18-P8A to R8-U1-P3B	3 meters
R8 U18 to Rack 1	R8-U18-P8B to R1-U1-P4B	10 meters
R8 U18 to Rack 2	R8-U24-P9A to R2-U1-P5B	10 meters
R8 U18 to Rack 3	R8-U18-P9B to R3-U1-P6B	10 meters
R8 U18 to Rack 4	R8-U18-P10A to R4-U1-P7B	10 meters
R8 U18 to Rack 5	R8-U18-P10B to R5-U1-P8B	5 meters
R8 U18 to Rack 6	R8-U24-P11A to R6-U1-P9B	5 meters
R8 U18 to Rack 7	R8-U24-P1B to R7-U1-P10B	5 meters

Connecting Expansion Racks

These topics provide instructions for connecting a SPARC SuperCluster T4-4 to an Oracle Exadata Storage Expansion Rack.

Oracle Exadata Storage Expansion Rack Components

Oracle Exadata Storage Expansion Rack provides additional storage for SPARC SuperCluster T4-4. The additional storage can be used for backups, historical data, and unstructured data. Oracle Exadata Storage Expansion Racks can be used to add space to SPARC SuperCluster T4-4 as follows:

- Add new Exadata Storage Servers and grid disks to a new Oracle Automatic Storage Management (Oracle ASM) disk group.
- Extend existing disk groups by adding grid disks in Oracle Exadata Storage Expansion Rack.
- Split Oracle Exadata Storage Expansion Rack among multiple SPARC SuperCluster T4-4 systems.

Oracle Exadata Storage Expansion Rack is available as a full rack, half rack, or quarter rack. The following table lists the components included in each type of Oracle Exadata Storage Expansion Rack.

TABLE 59 Components of Oracle Exadata Storage Expansion Racks

Oracle Exadata Storage Expansion Full Rack	Oracle Exadata Storage Expansion Half Rack	Oracle Exadata Storage Expansion Quarter Rack
■ 18 Exadata Storage Servers with 600 GB 15 K RPM High Performance	■ 9 Exadata Storage Servers with 600 GB 15 K RPM High Performance SAS disks	■ 4 Exadata Storage Servers with 600 GB 15 K RPM High Performance SAS disks

Oracle Exadata Storage Expansion Full Rack	Oracle Exadata Storage Expansion Half Rack	Oracle Exadata Storage Expansion Quarter Rack
SAS disks or 3 TB 7.2 K RPM High Capacity SAS disks[†] ■ 3 Sun Datacenter InfiniBand Switch 36 ■ 6.75 TB high speed flash ■ Keyboard, video, and mouse (KVM) hardware ■ 2 redundant 15 kVA PDU (single-phase or three-phase, high voltage or low voltage) ■ 1 48-port Cisco Catalyst 4948, model number WS-C4948-S Ethernet switch	or 3 TB 7.2 K RPM High Capacity SAS disk * ■ 3 Sun Datacenter InfiniBand Switch 36 ■ 3.4 TB high speed flash ■ Keyboard, video, and mouse (KVM) hardware ■ 2 redundant 15 kVA PDUs (single-phase or three-phase, high voltage or low voltage) ■ 1 48-port Cisco Catalyst 4948, model number WS-C4948-S Ethernet switch	or 3 TB 7.2 K RPM High Capacity SAS disks * ■ 2 Sun Datacenter InfiniBand Switch 36 ■ 1.5 TB high speed flash ■ Keyboard, video, and mouse (KVM) hardware ■ 2 redundant 15 kVA PDUs (single-phase or three-phase, high voltage or low voltage) ■ 1 48-port Cisco Catalyst 4948, model number WS-C4948-S Ethernet switch

[†]In earlier releases, the high capacity disks were 2 TB.

Preparing for Installation

These topics provide information to prepare your site for the installation of the Oracle Exadata Storage Expansion Rack. Planning for the expansion rack is similar to planning for the SPARC SuperCluster T4-4 rack. This section contains information specific to the expansion rack, and also refers to "Preparing the Site" for general planning information.

- "Reviewing System Specifications" on page 210
- "Reviewing Power Requirements" on page 211
- "Preparing for Cooling" on page 214
- "Preparing the Unloading Route and Unpacking Area" on page 217
- "Preparing the Network" on page 218

Reviewing System Specifications

- "Physical Specifications" on page 210
- "Installation and Service Area" on page 60
- "Rack and Floor Cutout Dimensions" on page 61

Physical Specifications

TABLE 60 Exadata Expansion Rack Specifications

Parameter	Metric	English
Height	1998 mm	78.66 in.

Parameter	Metric	English
Width with side panels	600 mm	23.62 in.
Depth with front and rear doors	1200 mm	47.24 in.
Depth without doors	1112 mm	43.78 in.
Minimum ceiling height	2300 mm	90 in.
Minimum space between top of cabinet and ceiling	914 mm	36 in.
Weight (full rack)	917.6 kg	2023 lbs
Weight (half rack)	578.3 kg	1275 lbs
Weight (quarter rack)	396.8 kg	875 lbs

Related Information

- “Installation and Service Area” on page 60
- “Rack and Floor Cutout Dimensions” on page 61

Reviewing Power Requirements

- “System Power Consumption” on page 211
- “Facility Power Requirements” on page 62
- “Grounding Requirements” on page 63
- “PDU Power Requirements” on page 212

System Power Consumption

TABLE 61 Oracle Exadata Storage Expansion Rack Power Consumption

Comments	Full Rack	Half Rack	Quarter Rack
Maximum	12.6 kW	6.9 kW	3.4 kW
Typical	12.9 kVA	7.1 kVA	3.5 kVA
	8.8 kW	4.8 kW	2.4 kW
	9.0 kVA	5.0 kVA	2.5kVA

Related Information

- “PDU Power Requirements” on page 212
- “Grounding Requirements” on page 63
- “Facility Power Requirements” on page 62

PDU Power Requirements

When ordering the Oracle Exadata Storage Expansion Rack, you must provide two specifications:

- Low or high voltage
- Single or three phase power

Refer to the following tables for Oracle marketing and manufacturing part numbers.

Note - Each Oracle Exadata Storage Expansion Rack has two power distribution units (PDUs). Both PDUs in a rack must be the same type.

TABLE 62 PDU Choices

Voltage	Phases	Reference
Low	1	Table 63
Low	3	Table 64
High	1	Table 65
High	3	Table 66

TABLE 63 Low Voltage 1 Phase PDUs for Oracle Exadata Storage Expansion Rack

High Voltage	One Phase	Comments
kVA Size	15 kVA	
Marketing part number	N/A	
Manufacturing part number	N/A	
Phase	1 ph	
Voltage	200 to 240 VAC	
Amps per PDU	72 A (3 inputs x 24 A)	
Outlets	42 C13 6 C19	
Number of inputs	3 inputs	
Input current	24 A max. per input	
Data center receptacle	NEMA 30A/250 VAC 2-pole/3-wire L6-30P	
Outlet groups per PDU	6	
Usable PDU power cord length	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 64 Low Voltage 3 Phase PDUs for Oracle Exadata Storage Expansion Rack

High Voltage	Three Phase	Comments
kVA Size	14.4 kVA	

High Voltage	Three Phase	Comments
Marketing part number	N/A	
Manufacturing part number	N/A	
Phase	3 ph	
Voltage	190 to 220 VAC	
Amps per PDU	69 A (3 inputs x 23 A)	
Outlets	42 C13 6 C19	
Number of inputs	2 inputs x 60 A, 3 ph	
Input current	40 A max. per phase	
Data center receptacle	IEC 60309 60A 4-pin 250 VAC 3 ph IP 67	
Outlet groups per PDU	6	
Usable PDU power cord length	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 65 High Voltage 1 Phase PDUs for Oracle Exadata Storage Expansion Rack

High Voltage	Single Phase	Comments
kVA Size	15 kVA	
Marketing part number	N/A	
Manufacturing part number	N/A	
Phase	1 ph	
Voltage	220 to 240 VAC	
Amps per PDU	72 A (3 inputs x24 A)	
Outlets	42 C13 6 C19	
Number of inputs	6	
Input current	25 A max. per input	
Data center receptacle	IEC 60309 32A 3-pin 250VAC IP44	
Outlet groups per PDU	6	
Usable PDU power cord length	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

TABLE 66 High Voltage 3 Phase PDUs for Oracle Exadata Storage Expansion Rack

High Voltage	Single Phase	Comments
kVA size	14.4 kVA	
Marketing part number	N/A	
Manufacturing part number	N/A	
Phase	3 ph	

High Voltage	Single Phase	Comments
Voltage	220/380 to 240/415 VAC	
Amps per PDU	62.7 A (3 inputs x 20.9. A)	
Outlets	42 C13 6 C19	
Number of inputs	2 inputs x 25 A	
Input current	25 A max. per phase	
Data center receptacle	IEC 60309 32A, 5-pin 230/400 VAC	
Outlet groups per PDU	6	
Usable PDU power cord length	2 m (6.6 feet)	PDU power cords are 4 m long (13 feet), but sections are used for internal routing in the rack.

Related Information

- “Facility Power Requirements” on page 62
- “Grounding Requirements” on page 63
- “System Power Consumption” on page 211

Preparing for Cooling

Environmental Requirements

The following table lists temperature, humidity and altitude requirements.

TABLE 67 Temperature, Humidity, and Altitude

Condition	Operating Requirement	Non-operating Requirement	Comments
Temperature	5 to 32°C (41 to 89.6°F)	-40 to 70°C (-40 to 158°F).	For optimal rack cooling, data center temperatures from 21 to 23°C (70 to 47°F)
Relative humidity	10 to 90% relative humidity, noncondensing	Up to 93% relative humidity.	For optimal data center rack cooling, 45 to 50%, noncondensing
Altitude	3048 m (10000 ft.) maximum	12000 m (40000 ft.).	Ambient temperature is reduced by 1 degree Celsius per 300 m

Condition	Operating Requirement	Non-operating Requirement	Comments
			above 900 m altitude above sea level

Related Information

- “Heat Dissipation and Airflow Requirements” on page 215
- “Perforated Floor Tiles” on page 217

Heat Dissipation and Airflow Requirements

The maximum and typical rates of heat released from the expansion racks are listed below. In order to cool the system properly, ensure that adequate airflow travels through the system.

TABLE 68 Heat Dissipation

Rack	Units	Maximum	Typical
Full rack	BTU/hour kJ/hour	43,000 45,400	30,100 31,800
Half rack	BTU/hour kJ/hour	23,600 24,900	16,500 17,400
Quarter rack	BTU/hour kJ/hour	11,600 12,250	8,100 8,600

Caution - Do not restrict the movement of cool air from the air conditioner to the cabinet, or the movement of hot air out of the rear of the cabinet.

Observe these additional requirements:

- Allow a minimum clearance of 914 mm (36 inches) at the front of the rack, and 914 mm (36 inches) at the rear of the rack for ventilation. There is no airflow requirement for the left and right sides, or the top of the rack.
- If the rack is not completely filled with components, cover the empty sections with filler panels.

FIGURE 18 Direction of Airflow Is Front to Back

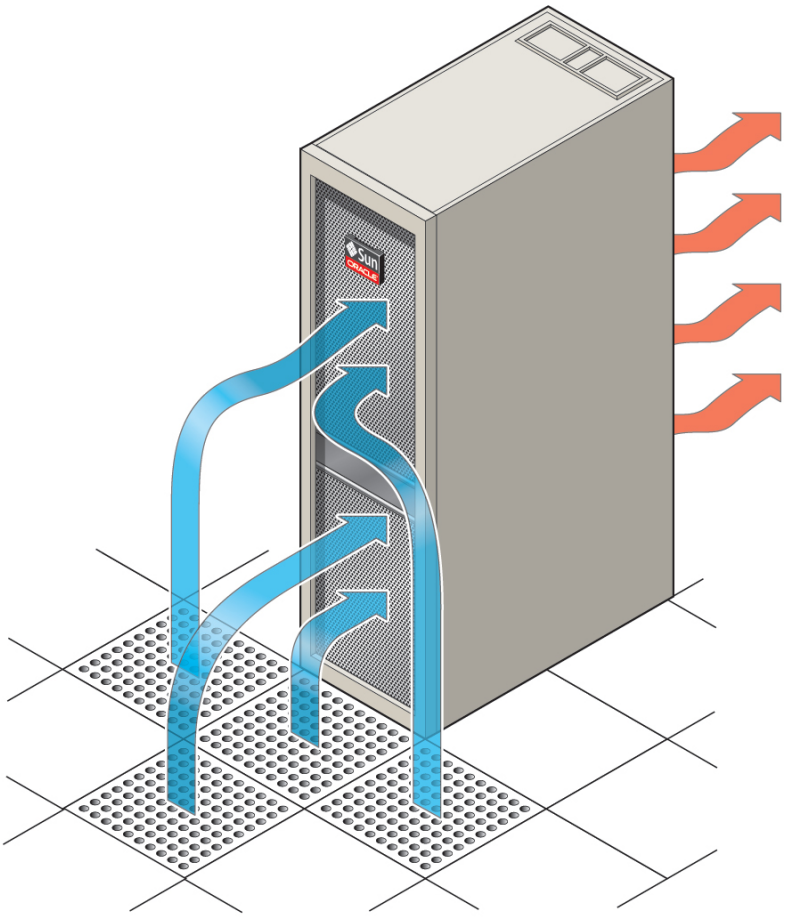

TABLE 69 Airflow (listed quantities are approximate)

Rack		Maximum	Typical
Full rack	CFM	2,000	1,390
Half rack	CFM	1,090	760
Quarter rack	CFM	530	375

Related Information

- “Environmental Requirements” on page 214

- "Perforated Floor Tiles" on page 217

Perforated Floor Tiles

If you install the server on a raised floor, use perforated tiles in front of the rack to supply cold air to the system. Each tile should support an airflow of approximately 400 CFM.

Perforated 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.

The following is the recommended number of floor tiles:

TABLE 70 Perforated Floor Tile Requirements

Rack	Number of Tiles
Full rack	4
Half rack	3
Quarter rack	2

Related Information

- "Heat Dissipation and Airflow Requirements" on page 215
- "Environmental Requirements" on page 214

Preparing the Unloading Route and Unpacking Area

- "Shipping Package Dimensions" on page 217
- "Loading Dock and Receiving Area Requirements" on page 70
- "Access Route Guidelines" on page 70
- "Unpacking Area" on page 71

Shipping Package Dimensions

TABLE 71 Exadata Expansion Rack Shipping Package Dimensions and Weights

Parameter	Metric	English
Height	2159 mm	85 in.
Width	1219 mm	48 in.
Depth	1575 mm	62 in.
Shipping weight (full rack)	1001.1 kg	2207 lbs

Parameter	Metric	English
Shipping weight (half rack)	659.6 kg	1454 lbs
Shipping weight (quarter rack)	475.3 kg	1048 lbs

Related Information

- “Loading Dock and Receiving Area Requirements” on page 70
- “Access Route Guidelines” on page 70
- “Unpacking Area” on page 71

Preparing the Network

- “Prepare DNS for the System” on page 73
- “Network Requirements Overview” on page 218
- “Network Connection and IP Address Requirements” on page 219

Network Requirements Overview

The Oracle Exadata Storage Expansion Rack includes Exadata Storage Servers, as well as equipment to connect the servers to your network. The network connections enable the servers to be administered remotely.

Each Exadata Storage Server consists of the following network components and interfaces:

- 1 embedded Gigabit Ethernet port (NET0) for connection to the host management network
- 1 dual-ported Sun QDR InfiniBand PCIe Low Profile host channel adapter (HCA) for connection to the InfiniBand private network
- 1 Ethernet port (NET MGT) for Oracle ILOM remote management

The Cisco Catalyst 4948 Ethernet switch supplied with the Oracle Exadata Storage Expansion Rack 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

To deploy the Oracle Exadata Storage Expansion Rack, ensure that you meet the minimum network requirements. There are two networks used for the Oracle Exadata Storage Expansion Rack. 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 the Exadata Storage Servers. It connects the servers, Oracle ILOM, and switches connected to the Ethernet switch in the rack. There is one uplink from the Ethernet switch in the rack to your existing management network.

  Note - Network connectivity to the PDUs is only required if the electric current will be monitored remotely.

  Each Exadata Storage Server use two network interfaces for management. One provides management access to the operating system through the 1 GbE host management interface (s), and the other provides access to the Oracle Integrated Lights Out Manager through the Oracle ILOM Ethernet interface.
- **InfiniBand private network** – This network connects the Exadata Storage Servers using the InfiniBand switches on the rack. This non-routable network is fully contained in the Oracle Exadata Storage Expansion Rack, and does not connect to your existing network. This network is automatically configured during installation.

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

Related Information

- “Network Connection and IP Address Requirements” on page 219
- “Prepare DNS for the System” on page 73

Network Connection and IP Address Requirements

Prior to installation, network cables must be run from your existing network infrastructure to the installation site. The network requirements for an Oracle Exadata Storage Expansion Rack are as follows:

TABLE 72 Oracle Exadata Storage Expansion Rack Network IP Address Requirements

Rack Type	Minimum Quantity	Comments
Full expansion rack	42	■ 18 for administration (one per Exadata Storage Server) ■ 18 for Oracle ILOM (one per Exadata Storage Server) ■ 4 for switches (3 for InfiniBand and 1 for Ethernet)

Rack Type	Minimum Quantity	Comments
		■ 2 for monitoring electric current of the PDUs (PDU connectivity is only required if the electric current will be monitored remotely.)
Half expansion rack	24	■ 9 for administration (one per Exadata Storage Server) ■ 9 for Oracle ILOM (one per Exadata Storage Server) ■ 4 for switches (3 for InfiniBand and 1 for Ethernet) ■ 2 for monitoring electric current of the PDUs (PDU connectivity is only required if the electric current will be monitored remotely.)
Quarter expansion rack	13	■ 4 for administration (one per Exadata Storage Server) ■ 4 for Oracle ILOM (one per Exadata Storage Server) ■ 3 for switches (2 for InfiniBand and 1 for Ethernet) ■ 2 for monitoring electric current of the PDUs (PDU connectivity is required only if the electric current will be monitored remotely.)

Related Information

- “Prepare DNS for the System” on page 73
- “Network Requirements Overview” on page 218
- “Default IP Addresses” on page 220
- “Cabling Tables (Internal)” on page 221

Installing the Oracle Exadata Storage Expansion Rack

The procedures for installing the Oracle Exadata Storage Expansion Rack are the same as those for installing the SPARC SuperCluster T4-4. See “Installing the System” for those procedure, then return here. Note that the sections on connecting to a 10 GbE client access network and using the optional Fibre Channel express modules do not apply for the Oracle Exadata Storage Expansion Rack.

Default IP Addresses

TABLE 73 Default IP Addresses for Oracle Exadata Storage Expansion Rack

Component	NET0 IP Addresses	Oracle ILOM IP Addresses	InfiniBand Bonded IP Addresses
Exadata Storage Server 18	192.168.1.68	192.168.1.168	192.168.10.68
Exadata Storage Server 17	192.168.1.67	192.168.1.167	192.168.10.67
Exadata Storage Server 16	192.168.1.66	192.168.1.166	192.168.10.66
Exadata Storage Server 15	192.168.1.65	192.168.1.165	192.168.10.65
Exadata Storage Server 14	192.168.1.64	192.168.1.164	192.168.10.64
Exadata Storage Server 13	192.168.1.63	192.168.1.163	192.168.10.63

Component	NET0 IP Addresses	Oracle ILOM IP Addresses	InfiniBand Bonded IP Addresses
Exadata Storage Server 12	192.168.1.62	192.168.1.162	192.168.10.62
Exadata Storage Server 11	192.168.1.61	192.168.1.161	192.168.10.61
Exadata Storage Server 10	192.168.1.60	192.168.1.160	192.168.10.60
Exadata Storage Server 9	192.168.1.59	192.168.1.159	192.168.10.59
Exadata Storage Server 8	192.168.1.58	192.168.1.158	192.168.10.58
Exadata Storage Server 7	192.168.1.57	192.168.1.157	192.168.10.57
Exadata Storage Server 6	192.168.1.56	192.168.1.156	192.168.10.56
Exadata Storage Server 5	192.168.1.55	192.168.1.155	192.168.10.55
Exadata Storage Server 4	192.168.1.54	192.168.1.154	192.168.10.54
Exadata Storage Server 3	192.168.1.53	192.168.1.153	192.168.10.53
Exadata Storage Server 2	192.168.1.52	192.168.1.152	192.168.10.52
Exadata Storage Server 1	192.168.1.51	192.168.1.151	192.168.10.51
Sun Datacenter InfiniBand Switch 36 switch 3	192.168.1.223	NA	NA
Sun Datacenter InfiniBand Switch 36 switch 2	192.168.1.222	NA	NA
Sun Datacenter InfiniBand Switch 36 switch 1	192.168.1.221	NA	NA
Ethernet switch	192.168.1.220	NA	NA
PDU-A	192.168.1.212	NA	NA
PDU-B	192.168.1.213	NA	NA

Cabling Tables (Internal)

These topics show the cable layouts for Oracle Exadata Storage Expansion Rack.

Front and Rear Views of the Oracle Exadata Storage Expansion Rack

The following figure shows the front and rear views of Oracle Exadata Storage Expansion Full Rack. The front is shown on the left, and the rear is shown on the right.

FIGURE 19 Rack Layout of Oracle Exadata Storage Expansion Rull Rack

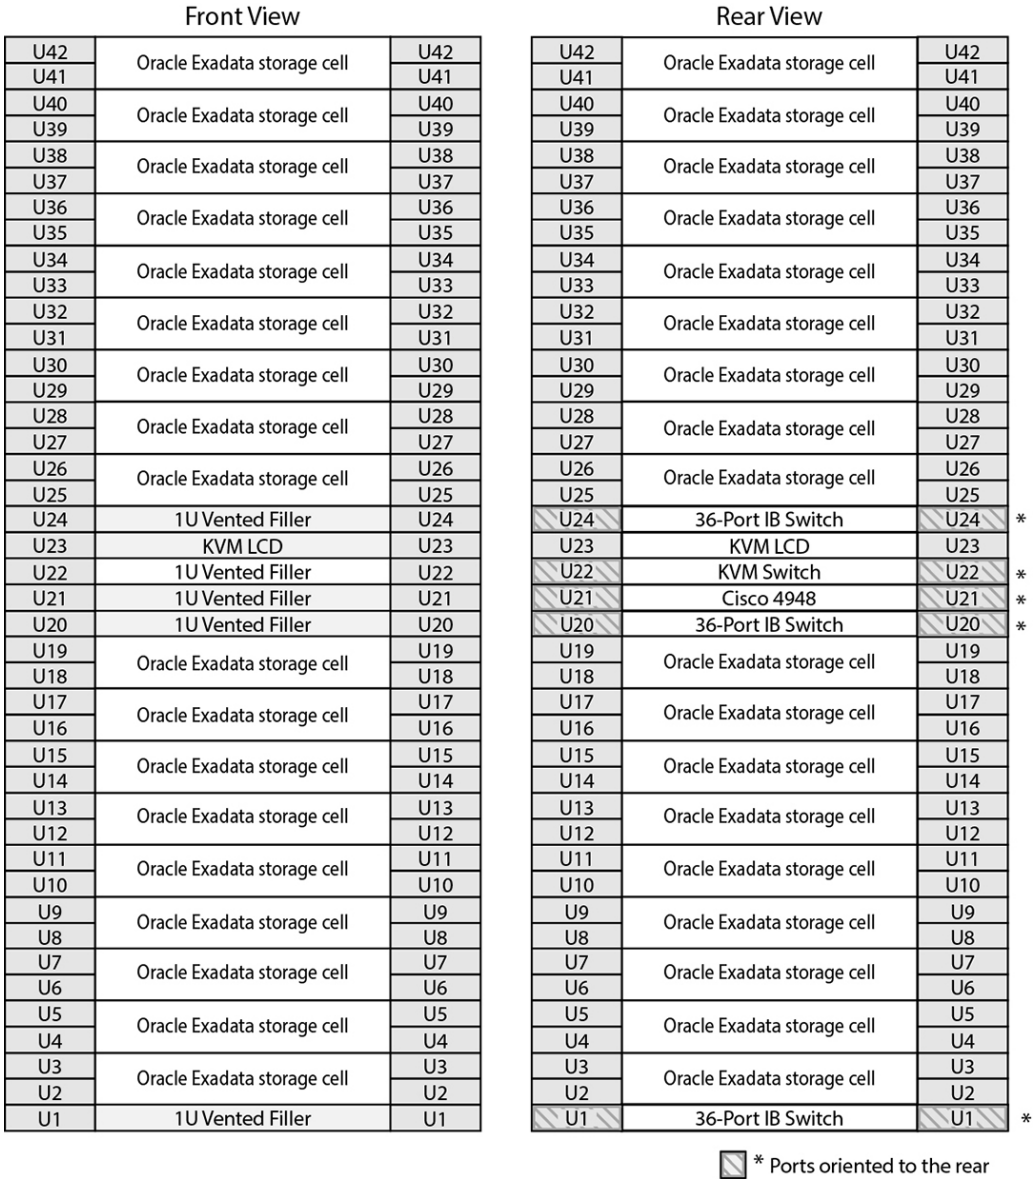

The following figure shows the front and rear views of Oracle Exadata Storage Expansion Half Rack. The front is shown on the left, and the rear is shown on the right.

FIGURE 20 Rack Layout of Oracle Storage Expansion Half Rack

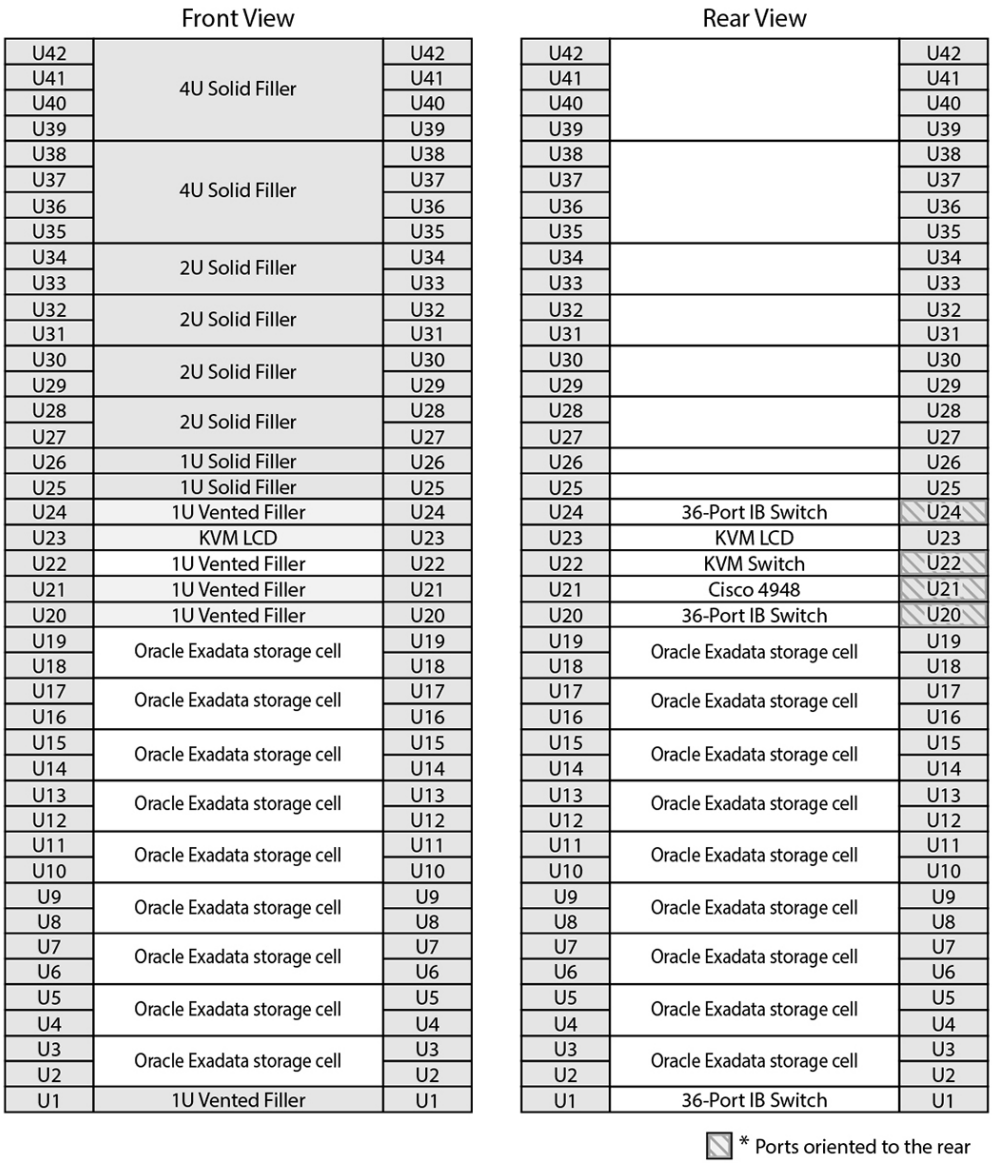

The following figure shows the front and rear views of Oracle Exadata Storage Expansion Quarter Rack. The front is shown on the left, and the rear is shown on the right.

FIGURE 21 Rack Layout of Oracle Exadata Storage Expansion Quarter Rack

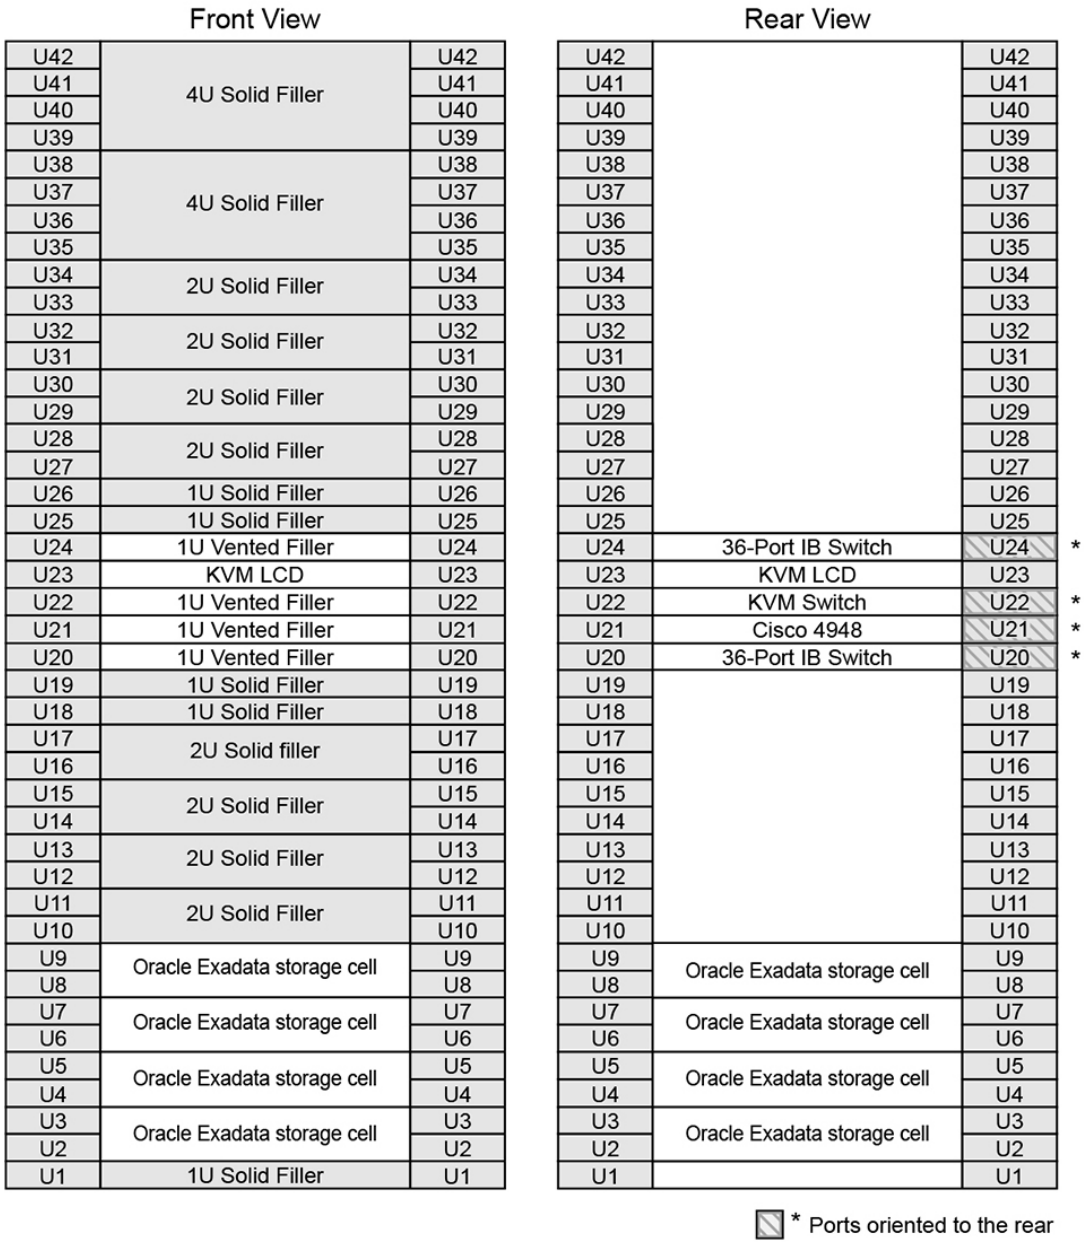

zz

Oracle ILOM Cabling Tables

This section contains the tables for the Oracle ILOM network cabling. The Oracle ILOM port on the servers is labeled NET MGT, and connects to the Gigabit Ethernet port located in rack unit 21 on Oracle Exadata Storage Expansion Racks.

The following table shows the cable connections from the servers to the Oracle ILOM switch in Oracle Exadata Storage Expansion Full Rack:

TABLE 74 Oracle ILOM Cabling for the Oracle Exadata Storage Expansion Full Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U41	Exadata Storage Server	2
U39	Exadata Storage Server	4
U37	Exadata Storage Server	6
U35	Exadata Storage Server	8
U33	Exadata Storage Server	10
U31	Exadata Storage Server	12
U29	Exadata Storage Server	14
U27	Exadata Storage Server	18
U25	Exadata Storage Server	22
U18	Exadata Storage Server	26
U16	Exadata Storage Server	30
U14	Exadata Storage Server	32
U12	Exadata Storage Server	34
U10	Exadata Storage Server	36
U8	Exadata Storage Server	38
U6	Exadata Storage Server	40
U4	Exadata Storage Server	42
U2	Exadata Storage Server	44

The following table shows the cable connections from the servers to the Oracle ILOM switch in Oracle Exadata Storage Expansion Half Rack:

TABLE 75 Oracle ILOM Cabling for the Oracle Exadata Storage Expansion Half Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U18	Exadata Storage Server	26
U16	Exadata Storage Server	30
U14	Exadata Storage Server	32
U12	Exadata Storage Server	34
U10	Exadata Storage Server	36

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U8	Exadata Storage Server	38
U6	Exadata Storage Server	40
U4	Exadata Storage Server	42
U2	Exadata Storage Server	44

The following table shows the cable connections from the servers to the Oracle ILOM switch in Oracle Exadata Storage Expansion Quarter Rack:

TABLE 76 Oracle ILOM Cabling for the Oracle Exadata Storage Expansion Quarter Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U8	Exadata Storage Server	38
U6	Exadata Storage Server	40
U4	Exadata Storage Server	42
U2	Exadata Storage Server	44

Administrative Gigabit Ethernet Port Cabling Tables

This section contains the tables for the administrative Gigabit Ethernet network cabling. The port on the servers is labeled Net-0, and connects to the Gigabit Ethernet port located in rack unit 21 on Oracle Exadata Storage Expansion Racks.

The following table shows the cable connections from the servers to the Gigabit Ethernet switch in Oracle Exadata Storage Expansion Full Rack:

TABLE 77 Gigabit Ethernet Cabling for the Oracle Exadata Storage Expansion Full Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U41	Exadata Storage Server	1
U39	Exadata Storage Server	3
U37	Exadata Storage Server	5
U35	Exadata Storage Server	7
U33	Exadata Storage Server	9
U31	Exadata Storage Server	11
U29	Exadata Storage Server	13
U27	Exadata Storage Server	17
U25	Exadata Storage Server	21
U24	Sun Datacenter InfiniBand Switch 36 switch	45
U20	Sun Datacenter InfiniBand Switch 36 switch	46

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U18	Exadata Storage Server	25
U16	Exadata Storage Server	29
U14	Exadata Storage Server	31
U12	Exadata Storage Server	33
U10	Exadata Storage Server	35
U8	Exadata Storage Server	37
U6	Exadata Storage Server	39
U4	Exadata Storage Server	41
U2	Exadata Storage Server	43
U1	Sun Datacenter InfiniBand Switch 36 switch	47
PDU-A	PDU	15
PDU-B	PDU	19

The following table shows the cable connections from the servers to the Gigabit Ethernet switch in Oracle Exadata Storage Expansion Half Rack:

TABLE 78 Gigabit Ethernet Cabling for the Oracle Exadata Storage Expansion Half Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U24	Sun Datacenter InfiniBand Switch 36 switch	45
U20	Sun Datacenter InfiniBand Switch 36 switch	46
U18	Exadata Storage Server	25
U16	Exadata Storage Server	29
U14	Exadata Storage Server	31
U12	Exadata Storage Server	33
U10	Exadata Storage Server	35
U8	Exadata Storage Server	37
U6	Exadata Storage Server	39
U4	Exadata Storage Server	41
U2	Exadata Storage Server	43
U1	Sun Datacenter InfiniBand Switch 36 switch	47
PDU-A	PDU	15
PDU-B	PDU	19

The following table shows the cable connections from the servers to the Gigabit Ethernet switch in Oracle Exadata Storage Expansion Quarter Rack:

TABLE 79 Gigabit Ethernet Cabling for the Oracle Exadata Storage Expansion Quarter Rack

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U24	Sun Datacenter InfiniBand Switch 36 switch	45

From Rack Unit	Type of Equipment	Gigabit Ethernet Port
U20	Sun Datacenter InfiniBand Switch 36 switch	46
U8	Exadata Storage Server	37
U6	Exadata Storage Server	39
U4	Exadata Storage Server	41
U2	Exadata Storage Server	43
PDU-A	PDU	15
PDU-B	PDU	19

Single Phase Power Distribution Unit Cabling Tables

This section contains the tables for single phase cabling from each power distribution unit (PDU) to the power supplies configured in each rack. The cables are terminated to PDU-A on the left, and routed to the right to enter CMA, and are bundled in groups of four.

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Full Rack.

TABLE 80 Single Phase PDU Cabling for the Oracle Exadata Storage Expansion Full Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U41	G5-6	G0-0	2 meters
U39	G5-3	G0-3	2 meters
U37	G5-0	G0-6	2 meters
U35	G4-6	G1-0	2 meters
U33	G4-4	G1-2	2 meters
U31	G4-2	G1-4	2 meters
U29	G3-6	G2-0	2 meters
U27	G3-5	G2-1	2 meters
U25	G3-3	G2-3	2 meters
U24	G3-1	G2-5	2 meters
U23	N/A	G3-0	included
U22	G2-5	G3-1	1 meter
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U18	G2-2	G3-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Half Rack.

TABLE 81 Single Phase PDU Cabling for the Oracle Exadata Storage Expansion Half Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U24	G3-1	G2-5	2 meters
U23	N/A	G3-0	included
U22	G2-5	G3-1	1 meter
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U18	G2-2	G3-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Quarter Rack.

TABLE 82 Single Phase PDU Cabling for the Oracle Exadata Storage Expansion Quarter Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U24	G3-1	G2-5	2 meters
U23	N/A	G3-0	included
U22	G2-5	G3-1	1 meter
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U8	G1-0	G4-6	2 meters

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters

Three Phase Power Distribution Unit Cabling Tables

This section contains the tables for three phase cabling from each power distribution unit (PDU) to the power supplies configured in each rack. The cables are terminated to PDU-A on the left, and routed to the right to enter CMA, and are bundled in groups of four.

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Full Rack.

TABLE 83 Three Phase PDU Cabling for the Oracle Exadata Storage Expansion Full Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U41	G5-6	G2-0	2 meters
U39	G5-3	G2-3	2 meters
U37	G5-0	G2-6	2 meters
U35	G4-6	G1-0	2 meters
U33	G4-4	G1-2	2 meters
U31	G4-2	G1-4	2 meters
U29	G3-6	G0-0	2 meters
U27	G3-5	G0-1	2 meters
U25	G3-3	G0-3	2 meters
U24	G3-1	G0-5	2 meters
U23	N/A	G5-0	included
U22	G2-5	G5-1	1 meter
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G5-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Half Rack.

TABLE 84 Three Phase PDU Cabling for the Oracle Exadata Storage Expansion Half Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U24	G3-1	G0-5	2 meters
U23	N/A	G5-0	included
U22	G2-5	G5-1	1 meter
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G5-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

The following table shows the cable connections from the PDUs to the rack units in Oracle Exadata Storage Expansion Quarter Rack.

TABLE 85 Three Phase PDU Cabling for the Oracle Exadata Storage Expansion Quarter Rack

Rack Unit	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U24	G3-1	G0-5	2 meters
U23	N/A	G5-0	included
U22	G2-5	G5-1	1 meter
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters

InfiniBand Network Cabling Tables

This section contains the tables for the InfiniBand network cabling. The Sun Datacenter InfiniBand Switch 36 switches are located in located in rack unit 1, 20 and 24.

The following table lists the location, ports and cables for the InfiniBand connections for Oracle Exadata Storage Expansion Full Rack.

TABLE 86 InfiniBand Network Cabling for the Oracle Exadata Storage Expansion Full Rack

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U24	0A	U41	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	0B	U39	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	1A	U37	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	1B	U35	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	2A	U33	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	2B	U31	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	3A	U29	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U24	4A	U27	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	5A	U25	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U24	13A	U18	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	14A	U16	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	14B	U14	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	15A	U12	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	15B	U10	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	16A	U8	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	16B	U6	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17A	U4	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17B	U2	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	0A	U41	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	0B	U39	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	1A	U37	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	1B	U35	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	2A	U33	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	2B	U31	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	3A	U29	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	4A	U27	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U20	5A	U25	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U20	13A	U18	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U20	14A	U16	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U20	14B	U14	Exadata Storage Server	PCIe 3, P1	2 meter QDR InfiniBand cable

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U20	15A	U12	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	15B	U10	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	16A	U8	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	16B	U6	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17A	U4	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17B	U2	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	9B	U24	Sun Datacenter InfiniBand Switch 36 switch	9A	2 meter QDR InfiniBand cable
U20	10B	U24	Sun Datacenter InfiniBand Switch 36 switch	10A	2 meter QDR InfiniBand cable
U20	11B	U24	Sun Datacenter InfiniBand Switch 36 switch	11A	2 meter QDR InfiniBand cable
U20	8A	U24	Sun Datacenter InfiniBand Switch 36 switch	8A	2 meter QDR InfiniBand cable
U20	9A	U24	Sun Datacenter InfiniBand Switch 36 switch	9B	2 meter QDR InfiniBand cable
U20	10A	U24	Sun Datacenter InfiniBand Switch 36 switch	10B	2 meter QDR InfiniBand cable
U20	11A	U24	Sun Datacenter InfiniBand Switch 36 switch	11B	2 meter QDR InfiniBand cable
U1	1B	U20	Sun Datacenter InfiniBand Switch 36 switch	8B	3 meter QDR InfiniBand cable
U1	0B	U24	Sun Datacenter InfiniBand Switch 36 switch	8B	3 meter QDR InfiniBand cable

The following table lists the location, ports and cables for the InfiniBand connections for Oracle Exadata Storage Expansion Half Rack.

TABLE 87 InfiniBand Network Cabling for the Oracle Exadata Storage Expansion Half Rack

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U24	13A	U18	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	14A	U16	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	14B	U14	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	15A	U12	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	15B	U10	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	16A	U8	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	16B	U6	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17A	U4	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17B	U2	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	13A	U18	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U20	14A	U16	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U20	14B	U14	Exadata Storage Server	PCIe 3, P1	2 meter QDR InfiniBand cable
U20	15A	U12	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	15B	U10	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	16A	U8	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	16B	U6	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17A	U4	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17B	U2	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	9B	U24	Sun Datacenter InfiniBand Switch 36 switch	9A	2 meter QDR InfiniBand cable
U20	10B	U24	Sun Datacenter InfiniBand Switch 36 switch	10A	2 meter QDR InfiniBand cable
U20	11B	U24	Sun Datacenter InfiniBand Switch 36 switch	11A	2 meter QDR InfiniBand cable
U20	8A	U24	Sun Datacenter InfiniBand Switch 36 switch	8A	2 meter QDR InfiniBand cable
U20	9A	U24	Sun Datacenter InfiniBand Switch 36 switch	9B	2 meter QDR InfiniBand cable
U20	10A	U24	Sun Datacenter InfiniBand Switch 36 switch	10B	2 meter QDR InfiniBand cable
U20	11A	U24	Sun Datacenter InfiniBand Switch 36 switch	11B	2 meter QDR InfiniBand cable
U1	1B	U20	Sun Datacenter InfiniBand Switch 36 switch	8B	3 meter QDR InfiniBand cable
U1	0B	U24	Sun Datacenter InfiniBand Switch 36 switch	8B	3 meter QDR InfiniBand cable

The following table lists the location, ports and cables for the InfiniBand connections for Oracle Exadata Storage Expansion Quarter Rack.

TABLE 88 InfiniBand Network Cabling for the Oracle Exadata Storage Expansion Quarter Rack

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U24	16A	U8	Exadata Storage Server	PCIe 2, P2	2 meter QDR InfiniBand cable
U24	16B	U6	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17A	U4	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U24	17B	U2	Exadata Storage Server	PCIe 3, P2	3 meter QDR InfiniBand cable
U20	16A	U8	Exadata Storage Server	PCIe 2, P1	2 meter QDR InfiniBand cable
U20	16B	U6	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17A	U4	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	17B	U2	Exadata Storage Server	PCIe 3, P1	3 meter QDR InfiniBand cable
U20	9B	U24	Sun Datacenter InfiniBand Switch 36 switch	9A	2 meter QDR InfiniBand cable

From InfiniBand Switch Rack Unit	Port	To Rack Unit	Type of Equipment	Port	Cable Description
U20	10B	U24	Sun Datacenter InfiniBand Switch 36 switch	10A	2 meter QDR InfiniBand cable
U20	11B	U24	Sun Datacenter InfiniBand Switch 36 switch	11A	2 meter QDR InfiniBand cable
U20	8A	U24	Sun Datacenter InfiniBand Switch 36 switch	8A	2 meter QDR InfiniBand cable
U20	9A	U24	Sun Datacenter InfiniBand Switch 36 switch	9B	2 meter QDR InfiniBand cable
U20	10A	U24	Sun Datacenter InfiniBand Switch 36 switch	10B	2 meter QDR InfiniBand cable
U20	11A	U24	Sun Datacenter InfiniBand Switch 36 switch	11B	2 meter QDR InfiniBand cable

Connecting an Oracle Exadata Storage Expansion Rack to the SPARC SuperCluster T4-4

Three Sun Datacenter InfiniBand Switch 36 switches are needed when cabling racks together. The SPARC SuperCluster T4-4 and the Oracle Exadata Storage Expansion Half Rack or Oracle Exadata Storage Expansion Full Rack include three Sun Datacenter InfiniBand Switch 36 switches. Two switches are used as leaf switches, and the third switch is used as a spine switch. These switches attach to standard Quad Small Form-factor Pluggable (QSFP) connectors at the end of the InfiniBand cables. The procedures in this section assume the racks are adjacent to each other. If they are not, then longer cables may be required for the connections.

Note - The Oracle Exadata Storage Expansion Quarter Rack contains only two Sun Datacenter InfiniBand Switch 36 switches (the leaf switches), so no spine switch is available in the Oracle Exadata Storage Expansion Quarter Rack. See “Connecting an Oracle Exadata Storage Expansion Quarter Rack to the SPARC SuperCluster T4-4” on page 236 for more information.

InfiniBand Switch Information for the SPARC SuperCluster T4-4

In the SPARC SuperCluster T4-4, the switch at rack unit 1 (U1) is referred to as the spine switch. The switches at rack unit 18 (U18) and rack unit 24 (U24) are referred to as leaf switches. In a single rack, the two leaf switches are interconnected using seven connections. In addition, each leaf switch has one connection to the spine switch. The leaf switches connect to the spine switch as shown in the following graphic.

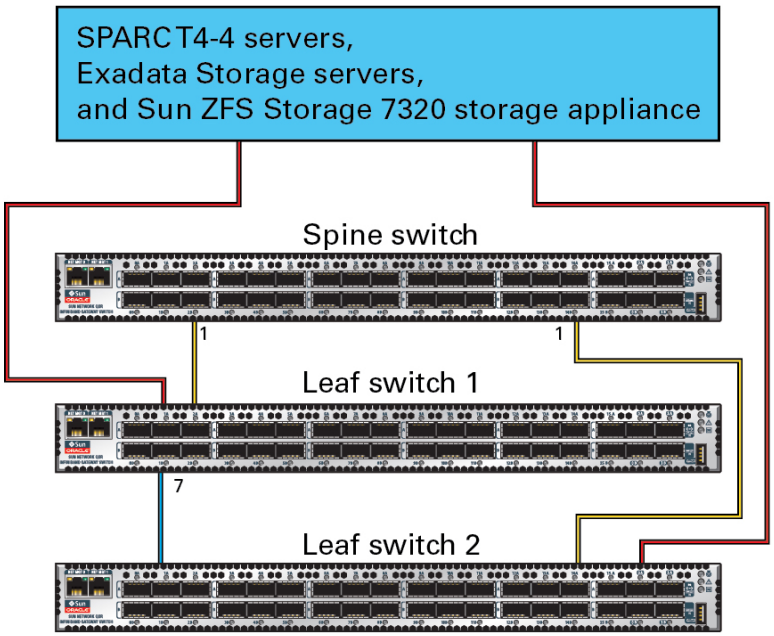

InfiniBand Switch Information for the Oracle Exadata Storage Expansion Rack

In the Oracle Exadata Storage Expansion Rack, the switch at rack unit 1 (U1) is referred to as the spine switch, except for the Oracle Exadata Storage Expansion Quarter Rack which has no spine switch. The switches at rack unit 20 (U20) and rack unit 24 (U24) are referred to as leaf switches.

Connecting an Oracle Exadata Storage Expansion Quarter Rack to the SPARC SuperCluster T4-4

These topics describe how to connect an Oracle Exadata Storage Expansion Quarter Rack to your SPARC SuperCluster T4-4 from Oracle.

Note - Only one Oracle Exadata Storage Expansion Quarter Rack can be connected to a SPARC SuperCluster T4-4 Half Rack or Full Rack.

Note - For instructions on connecting an Oracle Exadata Storage Expansion Half Rack or Oracle Exadata Storage Expansion Full Rack to the SPARC SuperCluster T4-4, see "Connecting an Oracle Exadata Storage Expansion Half Rack or Oracle Exadata Storage Expansion Full Rack to the SPARC SuperCluster T4-4" on page 239.

The following graphic shows the cable connections from an Oracle Exadata Storage Expansion Quarter Rack to a SPARC SuperCluster T4-4. The leaf switches within each rack maintain their existing seven connections. The leaf switches interconnect between the racks with two links each using the ports reserved for external connectivity.

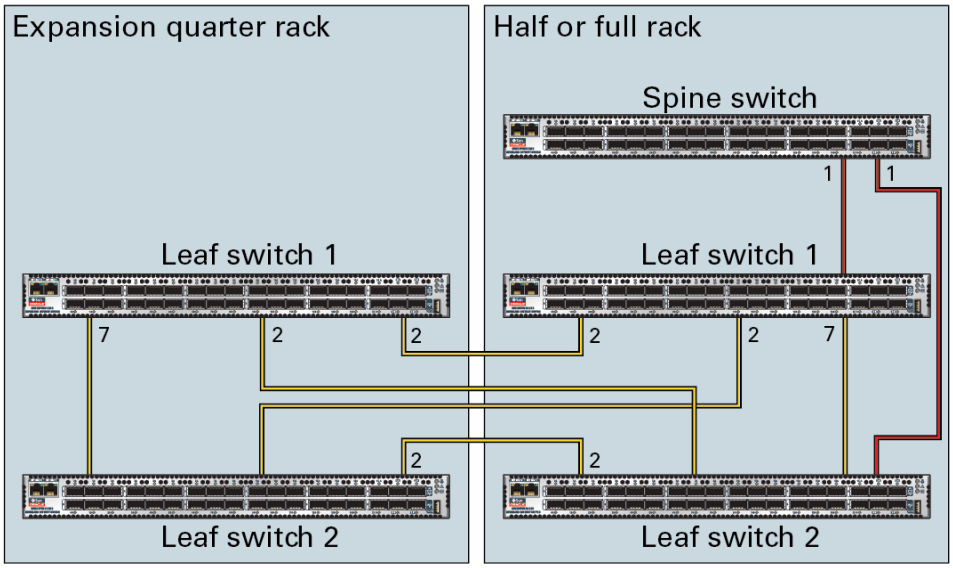

The following graphic shows the cable connections from an Oracle Exadata Storage Expansion Quarter Rack to two or more racks. The following racks can connect to an Oracle Exadata Storage Expansion Quarter Rack:

- SPARC SuperCluster T4-4 Half Rack
- SPARC SuperCluster T4-4 Full Rack
- Oracle Exadata Storage Expansion Half Rack
- Oracle Exadata Storage Expansion Full Rack

The racks are interconnected using a fat-tree topology. Each leaf switch in the Oracle Exadata Storage Expansion Quarter Rack connects to the spine switches in the other Half Racks or Full Racks with two links each. If there are more than four racks, then use one link instead of two.

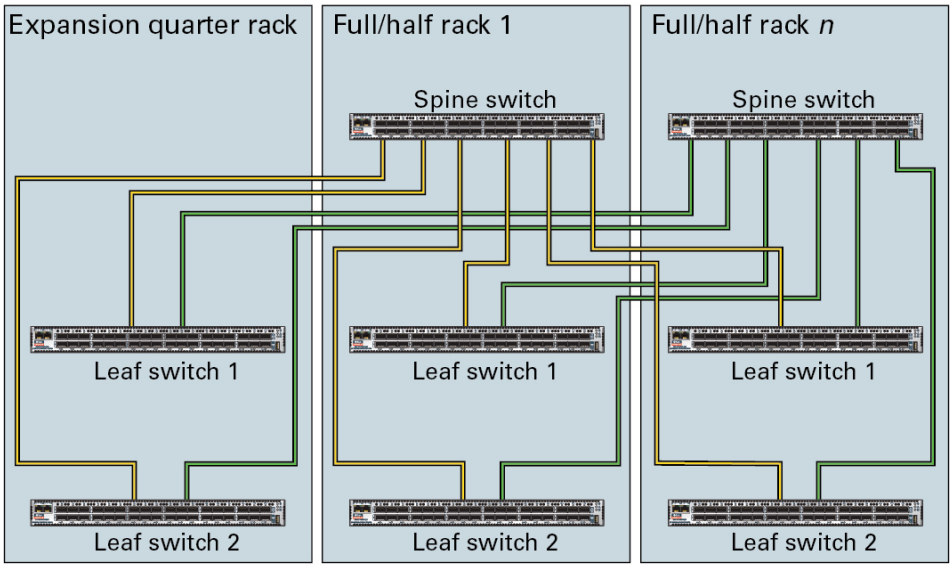

The following terms are used when referring to the two racks:

- Rack 1 (R1) refers to the SPARC SuperCluster T4-4
- Rack 2 (R2) refers to the Oracle Exadata Storage Expansion Quarter Rack

In addition, because the InfiniBand switches are physically located in different rack units in the two racks, the following terms are used when referring to the InfiniBand switches:

- InfiniBand 1 (IB1) refers to the spine switch, located in U1 in the SPARC SuperCluster T4-4
- InfiniBand 2 (IB2) refers to the first leaf switch, located in:
  - U18 in the SPARC SuperCluster T4-4
  - U20 in the Oracle Exadata Storage Expansion Rack
- InfiniBand 3 (IB3) refers to the second leaf switch, located in U24 in both racks

The following table shows the cable connections for a single Oracle Exadata Storage Expansion Quarter Rack to a SPARC SuperCluster T4-4.

TABLE 89 Leaf Switch Connections for the First Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R1-IB3 to Rack 2	R1-IB3-P2A to R2-IB3-P2A	5 meters
	R1-IB3-P2B to R2-IB3-P2B	
	R1-IB3-P7B to R2-IB2-P7B	
	R1-IB3-P12A to R2-IB2-P12A	
R1-IB2 to Rack 2	R1-IB2-P7B to R2-IB3-P7B	5 meters

Leaf Switch	Connection	Cable Length
	R1-IB2-P12A to R2-IB3-P12A R1-IB2-P2A to R2-IB2-P2A R1-IB2-P2B to R2-IB2-P2B	

Connecting an Oracle Exadata Storage Expansion Half Rack or Oracle Exadata Storage Expansion Full Rack to the SPARC SuperCluster T4-4

These topics describe how to connect an Oracle Exadata Storage Expansion Half Rack or Oracle Exadata Storage Expansion Full Rack to your SPARC SuperCluster T4-4 from Oracle.

Note - For instructions on connecting an Oracle Exadata Storage Expansion Quarter Rack to the SPARC SuperCluster T4-4, see "Connecting an Oracle Exadata Storage Expansion Quarter Rack to the SPARC SuperCluster T4-4" on page 236.

When connecting multiple racks together, remove the seven existing inter-switch connections between each leaf switches, as well as the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multi-rack environments, the leaf switches inside a rack are no longer directly interconnected, as shown in the following graphic.

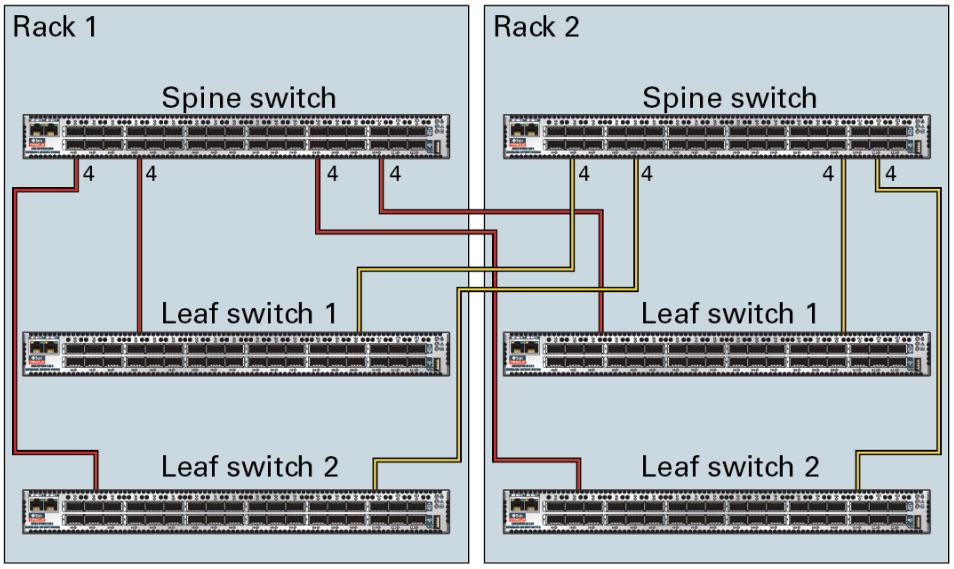

As shown in the preceding graphic, each leaf switch in rack 1 connects to the following switches:

- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2

The spine switch in rack 1 connects to the following switches:

- Eight connections to both internal leaf switches
- Eight connections to both leaf switches in rack 2

The following terms are used when referring to the two racks:

- Rack 1 (R1) refers to the SPARC SuperCluster T4-4
- Rack 2 (R2) refers to the Oracle Exadata Storage Expansion Rack

In addition, because the InfiniBand switches are physically located in different rack units in the two racks, the following terms are used when referring to the InfiniBand switches:

- InfiniBand 1 (IB1) refers to the spine switch, located in U1 in both racks
- InfiniBand 2 (IB2) refers to the first leaf switch, located in:
  - U18 in the SPARC SuperCluster T4-4
  - U20 in the Oracle Exadata Storage Expansion Rack
- InfiniBand 3 (IB3) refers to the second leaf switch, located in U24 in both racks

Two-Rack C20abling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling two racks together.

TABLE 90 Leaf Switch Connections for the First Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R1-IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	5 meters
	R1-IB3-P8B to R1-IB1-P4A	
	R1-IB3-P9A to R1-IB1-P5A	
	R1-IB3-P9B to R1-IB1-P6A	
R1-IB3 to Rack 2	R1-IB3-P10A to R2-IB1-P7A	5 meters
	R1-IB3-P10B to R2-IB1-P8A	
	R1-IB3-P11A to R2-IB1-P9A	
	R1-IB3-P11B to R2-IB1-P10A	
R1-IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	5 meters

Leaf Switch	Connection	Cable Length
	R1-IB2-P8B to R1-IB1-P4B	
	R1-IB2-P9A to R1-IB1-P5B	
	R1-IB2-P9B to R1-IB1-P6B	
R1-IB2 to Rack 2	R1-IB2-P10A to R2-IB1-P7B	5 meters
	R1-IB2-P10B to R2-IB1-P8B	
	R1-IB2-P11A to R2-IB1-P9B	
	R1-IB2-P11B to R2-IB1-P10B	

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling two racks together.

TABLE 91 Leaf Switch Connections for the Second Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R2-IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	5 meters
	R2-IB3-P8B to R2-IB1-P4A	
	R2-IB3-P9A to R2-IB1-P5A	
	R2-IB3-P9B to R2-IB1-P6A	
R2-IB3 to Rack 1	R2-IB3-P10A to R1-IB1-P7A	5 meters
	R2-IB3-P10B to R1-IB1-P8A	
	R2-IB3-P11A to R1-IB1-P9A	
	R2-IB3-P11B to R1-IB1-P10A	
R2-IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	5 meters
	R2-IB2-P8B to R2-IB1-P4B	
	R2-IB2-P9A to R2-IB1-P5B	
	R2-IB2-P9B to R2-IB1-P6B	
R2-IB2 to Rack 1	R2-IB2-P10A to R1-IB1-P7B	5 meters
	R2-IB2-P10B to R1-IB1-P8B	
	R2-IB2-P11A to R1-IB1-P9B	
	R2-IB2-P11B to R1-IB1-P10B	

Three-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling three full racks together.

TABLE 92 Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1-IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	5 meters
	R1-IB3-P8B to R1-IB1-P4A	
	R1-IB3-P9A to R1-IB1-P5A	
R1-IB3 to Rack 2	R1-IB3-P9B to R2-IB1-P6A	5 meters
	R1-IB3-P10A to R2-IB1-P7A	
	R1-IB3-P10B to R2-IB1-P8A	
R1-IB3 to Rack 3	R1-IB3-P11A to R3-IB1-P9A	5 meters
	R1-IB3-P11B to R3-IB1-P10A	
R1-IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	5 meters
	R1-IB2-P8B to R1-IB1-P4B	
	R1-IB2-P9A to R1-IB1-P5B	
R1-IB2 to Rack 2	R1-IB2-P9B to R2-IB1-P6B	5 meters
	R1-IB2-P10A to R2-IB1-P7B	
	R1-IB2-P10B to R2-IB1-P8B	
R1-IB2 to Rack 3	R1-IB2-P11A to R3-IB1-P9B	5 meters
	R1-IB2-P11B to R3-IB1-P10B	

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling three racks together.

TABLE 93 Leaf Switch Connections for the Second Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2-IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	5 meters
	R2-IB3-P8B to R2-IB1-P4A	
	R2-IB3-P9A to R2-IB1-P5A	
R2-IB3 to Rack 1	R2-IB3-P11A to R1-IB1-P9A	5 meters
	R2-IB3-P11B to R1-IB1-P10A	
R2-IB3 to Rack 3	R2-IB3-P9B to R3-IB1-P6A	5 meters
	R2-IB3-P10A to R3-IB1-P7A	
	R2-IB3-P10B to R3-IB1-P8A	
R2-IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	5 meters
	R2-IB2-P8B to R2-IB1-P4B	
	R2-IB2-P9A to R2-IB1-P5B	

Leaf Switch	Connection	Cable Length
R2-IB2 to Rack 1	R2-IB2-P11A to R1-IB1-P9B	5 meters
	R2-IB2-P11B to R1-IB1-P10B	
R2-IB2 to Rack 3	R2-IB2-P9B to R3-IB1-P6B	5 meters
	R2-IB2-P10A to R3-IB1-P7B	
	R2-IB2-P10B to R3-IB1-P8B	

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling three full racks together.

TABLE 94 Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3-IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	5 meters
	R3-IB3-P8B to R3-IB1-P4A	
	R3-IB3-P9A to R3-IB1-P5A	
R3-IB3 to Rack 1	R3-IB3-P9B to R1-IB1-P6A	5 meters
	R3-IB3-P10A to R1-IB1-P7A	
	R3-IB3-P10B to R1-IB1-P8A	
R3-IB3 to Rack 2	R3-IB3-P11A to R2-IB1-P9A	5 meters
	R3-IB3-P11B to R2-IB1-P10A	
R3-IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	5 meters
	R3-IB2-P8B to R3-IB1-P4B	
	R3-IB2-P9A to R3-IB1-P5B	
R3-IB2 to Rack 1	R3-IB2-P9B to R1-IB1-P6B	5 meters
	R3-IB2-P10A to R1-IB1-P7B	
	R3-IB2-P10B to R1-IB1-P8B	
R3-IB2 to Rack 2	R3-IB2-P11A to R2-IB1-P9B	5 meters
	R3-IB2-P11B to R2-IB1-P10B	

Four-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling four full racks together.

TABLE 95 Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1-IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	5 meters
	R1-IB3-P8B to R1-IB1-P4A	
R1-IB3 to Rack 2	R1-IB3-P9A to R2-IB1-P5A	5 meters
	R1-IB3-P9B to R2-IB1-P6A	
R1-IB3 to Rack 3	R1-IB3-P10A to R3-IB1-P7A	5 meters
	R1-IB3-P10B to R3-IB1-P8A	
R1-IB3 to Rack 4	R1-IB3-P11A to R4-IB1-P9A	10 meters
	R1-IB3-P11B to R4-IB1-P10A	
R1-IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	5 meters
	R1-IB2-P8B to R1-IB1-P4B	
R1-IB2 to Rack 2	R1-IB2-P9A to R2-IB1-P5B	5 meters
	R1-IB2-P9B to R2-IB1-P6B	
R1-IB2 to Rack 3	R1-IB2-P10A to R3-IB1-P7B	5 meters
	R1-IB2-P10B to R3-IB1-P8B	
R1-IB2 to Rack 4	R1-IB2-P11A to R4-IB1-P9B	10 meters
	R1-IB2-P11B to R4-IB1-P10B	

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling four full racks together.

TABLE 96 Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2-IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	5 meters
	R2-IB3-P8B to R2-IB1-P4A	
R2-IB3 to Rack 1	R2-IB3-P11A to R1-IB1-P9A	5 meters
	R2-IB3-P11B to R1-IB1-P10A	
R2-IB3 to Rack 3	R2-IB3-P9A to R3-IB1-P5A	5 meters
	R2-IB3-P9B to R3-IB1-P6A	
R2-IB3 to Rack 4	R2-IB3-P10A to R4-IB1-P7A	5 meters
	R2-IB3-P10B to R4-IB1-P8A	
R2-IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	5 meters
	R2-IB2-P8B to R2-IB1-P4B	
R2-IB2 to Rack 1	R2-IB2-P11A to R1-IB1-P9B	5 meters
	R2-IB2-P11B to R1-IB1-P10B	

Leaf Switch	Connection	Cable Length
R2-IB2 to Rack 3	R2-IB2-P9A to R3-IB1-P5B	5 meters
	R2-IB2-P9B to R3-IB1-P6B	
R2-IB2 to Rack 4	R2-IB2-P10A to R4-IB1-P7B	5 meters
	R2-IB2-P10B to R4-IB1-P8B	

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling four full racks together.

TABLE 97 Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R3-IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	5 meters
	R3-IB3-P8B to R3-IB1-P4A	
R3-IB3 to Rack 1	R3-IB3-P10A to R1-IB1-P7A	5 meters
	R3-IB3-P10B to R1-IB1-P8A	
R3-IB3 to Rack 2	R3-IB3-P11A to R2-IB1-P9A	5 meters
	R3-IB3-P11B to R2-IB1-P10A	
R3-IB3 to Rack 4	R3-IB3-P9A to R4-IB1-P5A	5 meters
	R3-IB3-P9B to R4-IB1-P6A	
R3-IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	5 meters
	R3-IB2-P8B to R3-IB1-P4B	
R3-IB2 to Rack 1	R3-IB2-P10A to R1-IB1-P7B	5 meters
	R3-IB2-P10B to R1-IB1-P8B	
R3-IB2 to Rack 2	R3-IB2-P11A to R2-IB1-P9B	5 meters
	R3-IB2-P11B to R2-IB1-P10B	
R3-IB2 to Rack 4	R3-IB2-P9A to R4-IB1-P5B	5 meters
	R3-IB2-P9B to R4-IB1-P6B	

The following table shows the cable connections for the fourth spine switch (R4-IB1) when cabling four full racks together.

TABLE 98 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4-IB3 within Rack 4	R4-IB3-P8A to R4-IB1-P3A	5 meters
	R4-IB3-P8B to R4-IB1-P4A	
R4-IB3 to Rack 1	R4-IB3-P9A to R1-IB1-P5A	10 meters

Leaf Switch	Connection	Cable Length
	R4-IB3-P9B to R1-IB1-P6A	
R4-IB3 to Rack 2	R4-IB3-P10A to R2-IB1-P7A	5 meters
	R4-IB3-P10B to R2-IB1-P8A	
R4-IB3 to Rack 3	R4-IB3-P11A to R3-IB1-P9A	5 meters
	R4-IB3-P11B to R3-IB1-P10A	
R4-IB2 within Rack 4	R4-IB2-P8A to R4-IB1-P3B	5 meters
	R4-IB2-P8B to R4-IB1-P4B	
R4-IB2 to Rack 1	R4-IB2-P9A to R1-IB1-P5B	10 meters
	R4-IB2-P9B to R1-IB1-P6B	
R4-IB2 to Rack 2	R4-IB2-P10A to R2-IB1-P7B	5 meters
	R4-IB2-P10B to R2-IB1-P8B	
R4-IB2 to Rack 3	R4-IB2-P11A to R3-IB1-P9B	5 meters
	R4-IB2-P11B to R3-IB1-P10B	

Five-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling five full racks together.

TABLE 99 Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	3 meters
	R1-IB3-P8B to R1-IB1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-IB1-P5A	5 meters
	R1-IB3-P9B to R2-IB1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-IB1-P7A	5 meters
	R1-IB3-P10B to R3-IB1-P8A	
R1 IB3 to Rack 4	R1-IB3-P11A to R4-IB1-P9A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11B to R5-IB1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	3 meters
	R1-IB2-P8B to R1-IB1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-IB1-P5B	3 meters
	R1-IB2-P9B to R2-IB1-P6B	
R1 IB2 to Rack 3	R1-IB2-P10A to R3-IB1-P7B	5 meters
	R1-IB2-P10B to R3-IB1-P8B	

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 4	R1-IB2-P11A to R4-IB1-P9B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11B to R5-IB1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling five full racks together.

TABLE 100 Leaf Switch Connections for the Second Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	3 meters
	R2-IB3-P8B to R2-IB1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-IB1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-IB1-P5A	5 meters
	R2-IB3-P9B to R3-IB1-P6A	
R2 IB3 to Rack 4	R2-IB3-P10A to R4-IB1-P7A	5 meters
	R2-IB3-P10B to R4-IB1-P8A	
R2 IB3 to Rack 5	R2-IB3-P11A to R5-IB1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	3 meters
	R2-IB2-P8B to R2-IB1-P4B	
R2 IB2 to Rack 1	R2-IB2-P11B to R1-IB1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-IB1-P5B	5 meters
	R2-IB2-P9B to R3-IB1-P6B	
R2 IB2 to Rack 4	R2-IB2-P10A to R4-IB1-P7B	5 meters
	R2-IB2-P10B to R4-IB1-P8B	
R2 IB2 to Rack 5	R2-IB2-P11A to R5-IB1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling five full racks together.

TABLE 101 Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	3 meters
	R3-IB3-P8B to R3-IB1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-IB1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-IB1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-IB1-P5A	5 meters
	R3-IB3-P9B to R4-IB1-P6A	

Leaf Switch	Connection	Cable Length
R3 IB3 to Rack 5	R3-IB3-P10A to R5-IB1-P7A	5 meters
	R3-IB3-P10B to R5-IB1-P8A	
R3 IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	3 meters
	R3-IB2-P8B to R3-IB1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-IB1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-IB1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-IB1-P5B	5 meters
	R3-IB2-P9B to R4-IB1-P6B	
R3 IB2 to Rack 5	R3-IB2-P10A to R5-IB1-P7B	5 meters
	R3-IB2-P10B to R5-IB1-P8B	

The following table shows the cable connections for the fourth spine switch (R4-IB1) when cabling five full racks together.

TABLE 102 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-IB1-P3A	3 meters
	R4-IB3-P8B to R4-IB1-P4A	
R4 IB3 to Rack 1	R4-IB3-P10A to R1-IB1-P7A	10 meters
	R4-IB3-P10B to R1-IB1-P8A	
R4 IB3 to Rack 2	R4-IB3-P11A to R2-IB1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-IB1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-IB1-P5A	5 meters
	R4-IB3-P9B to R5-IB1-P6A	
R4 IB2 within Rack 4	R4-IB2-P8A to R4-IB1-P3B	3 meters
	R4-IB2-P8B to R4-IB1-P4B	
R4 IB2 to Rack 1	R4-IB2-P10A to R1-IB1-P7B	10 meters
	R4-IB2-P10B to R1-IB1-P8B	
R4 IB2 to Rack 2	R4-IB2-P11A to R2-IB1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-IB1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-IB1-P5B	5 meters
	R4-IB2-P9B to R5-IB1-P6B	

The following table shows the cable connections for the fifth spine switch (R5-IB1) when cabling five full racks together.

TABLE 103 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-IB1-P3A	3 meters
	R5-IB3-P8B to R5-IB1-P4A	
R5 IB3 to Rack 1	R5-IB3-P9A to R1-IB1-P5A	10 meters
	R5-IB3-P9B to R1-IB1-P6A	
R5 IB3 to Rack 2	R5-IB3-P10A to R2-IB1-P7A	10 meters
	R5-IB3-P10B to R2-IB1-P8A	
R5 IB3 to Rack 3	R5-IB3-P11A to R3-IB1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-IB1-P10A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-IB1-P3B	3 meters
	R5-IB2-P8B to R5-IB1-P4B	
R5 IB2 to Rack 1	R5-IB2-P9A to R1-IB1-P5B	10 meters
	R5-IB2-P9B to R1-IB1-P6B	
R5 IB2 to Rack 2	R5-IB2-P10A to R2-IB1-P7B	10 meters
	R5-IB3-P10B to R2-IB1-P8B	
R5 IB2 to Rack 3	R5-IB2-P11A to R3-IB1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-IB1-P10B	5 meters

Six-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling six full racks together.

TABLE 104 Leaf Switch Connections for the First Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	3 meters
	R1-IB3-P8B to R1-IB1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-IB1-P5A	5 meters
	R1-IB3-P9B to R2-IB1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-IB1-P7A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10B to R4-IB1-P8A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11A to R5-IB1-P9A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11B to R6-IB1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	3 meters
	R1-IB2-P8B to R1-IB1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-IB1-P5B	5 meters

Leaf Switch	Connection	Cable Length
	R1-IB2-P9B to R2-IB1-P6B	
R1 IB2 to Rack 3	R1-IB2-P10A to R3-IB1-P7B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10B to R4-IB1-P8B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11A to R5-IB1-P9B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11B to R6-IB1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling six full racks together.

TABLE 105 Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	3 meters
	R2-IB3-P8B to R2-IB1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-IB1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-IB1-P5A	5 meters
	R2-IB3-P9B to R3-IB1-P6A	
R2 IB3 to Rack 4	R2-IB3-P10A to R4-IB1-P7A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10B to R5-IB1-P8A	10 meters
R2 IB3 to Rack 6	R2-IB3-P11Ato R6-IB1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	3 meters
	R2-IB2-P8B to R2-IB1-P4B	
R2 IB2 to Rack 1	R2-IB2-P11B to R1-IB1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-IB1-P5B	5 meters
	R2-IB2-P9B to R3-IB1-P6B	
R2 IB2 to Rack 4	R2-IB2-P10A to R4-IB1-P7B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10B to R5-IB1-P8B	10 meters
R2 IB2 to Rack 6	R2-IB2-P11Ato R6-IB1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling six full racks together.

TABLE 106 Leaf Switch Connections for the Third Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	3 meters
	R3-IB3-P8B to R3-IB1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-IB1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-IB1-P10A	5 meters

Leaf Switch	Connection	Cable Length
R3 IB3 to Rack 4	R3-IB3-P9A to R4-IB1-P5A	5 meters
	R3-IB3-P9B to R4-IB1-P6A	
R3 IB3 to Rack 5	R3-IB3-P10A to R5-IB1-P7A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10B to R6-IB1-P8A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	3 meters
	R3-IB2-P8B to R3-IB1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-IB1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-IB1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-IB1-P5B	5 meters
	R3-IB2-P9B to R4-IB1-P6B	
R3 IB2 to Rack 5	R3-IB2-P10A to R5-IB1-P7B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10B to R6-IB1-P8B	5 meters

The following table shows the cable connections for the fourth spine switch (R4-IB1) when cabling six full racks together.

TABLE 107 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-IB1-P3A	3 meters
	R4-IB3-P8B to R4-IB1-P4A	
R4 IB3 to Rack 1	R4-IB3-P10B to R1-IB1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-IB1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-IB1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-IB1-P5A	5 meters
	R4-IB3-P9B to R5-IB1-P6A	
R4 IB3 to Rack 6	R4-IB3-P10A to R6-IB1-P7A	5 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-IB1-P3B	3 meters
	R4-IB2-P8B to R4-IB1-P4B	
R4 IB2 to Rack 1	R4-IB2-P10B to R1-IB1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-IB1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-IB1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-IB1-P5B	5 meters
	R4-IB2-P9B to R5-IB1-P6B	
R4 IB2 to Rack 6	R4-IB2-P10A to R6-IB1-P7B	5 meters

The following table shows the cable connections for the fifth spine switch (R5-IB1) when cabling six full racks together.

TABLE 108 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-IB1-P3A	3 meters
	R5-IB3-P8B to R5-IB1-P4A	
R5 IB3 to Rack 1	R5-IB3-P10A to R1-IB1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-IB1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-IB1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-IB1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-IB1-P5A	5 meters
	R5-IB3-P9B to R6-IB1-P6A	
R5 IB2 within Rack 5	R5-IB2-P8A to R5-IB1-P3B	3 meters
	R5-IB2-P8B to R5-IB1-P4B	
R5 IB2 to Rack 1	R5-IB2-P10A to R1-IB1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-IB1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-IB1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-IB1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-IB1-P5B	5 meters
	R5-IB2-P9B to R6-IB1-P6B	

The following table shows the cable connections for the sixth spine switch (R6-IB1) when cabling six full racks together.

TABLE 109 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-IB1-P3A	3 meters
	R6-IB3-P8B to R6-IB1-P4A	
R6 IB3 to Rack 1	R6-IB3-P9A to R1-IB1-P5A	10 meters
	R6-IB3-P9B to R1-IB1-P6A	
R6 IB3 to Rack 2	R6-IB3-P10A to R2-IB1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-IB1-P8A	5 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-IB1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-IB1-P10A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-IB1-P3B	3 meters
	R6-IB2-P8B to R6-IB1-P4B	
R6 IB2 to Rack 2	R6-IB3-P10A to R2-IB1-P7B	10 meters
R6 IB2 to Rack 1	R6-IB2-P9A to R1-IB1-P5B	10 meters
	R6-IB2-P9B to R1-IB1-P6B	
R6 IB2 to Rack 3	R6-IB2-P10B to R3-IB1-P8B	5 meters

Leaf Switch	Connection	Cable Length
R6 IB2 to Rack 4	R6-IB2-P11A to R4-IB1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-IB1-P10B	5 meters

Seven-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling seven full racks together.

TABLE 110 Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	3 meters
	R1-IB3-P8B to R1-IB1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-IB1-P5A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9B to R3-IB1-P6A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10A to R4-IB1-P7A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10B to R5-IB1-P8A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11A to R6-IB1-P9A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11B to R7-IB1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	3 meters
	R1-IB2-P8B to R1-IB1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-IB1-P5B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9B to R3-IB1-P6B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10A to R4-IB1-P7B	10 meters
R1 IB2 to Rack 5	R1-IB2-P10B to R5-IB1-P8B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11A to R6-IB1-P9B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11B to R7-IB1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling seven full racks together.

TABLE 111 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	3 meters
	R2-IB3-P8B to R2-IB1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-IB1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-IB1-P5A	5 meters

Leaf Switch	Connection	Cable Length
R2 IB3 to Rack 4	R2-IB3-P9B to R4-IB1-P6A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10A to R5-IB1-P7A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10B to R6-IB1-P8A	10 meters
R2 IB3 to Rack 7	R2-IB3-P11A to R7-IB1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	3 meters
	R2-IB2-P8B to R2-IB1-P4B	
R2 IB2 to Rack 1	R2-IB2-P11B to R1-IB1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-IB1-P5B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9B to R4-IB1-P6B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10A to R5-IB1-P7B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10Bto R6-IB1-P8B	10 meters
R2 IB2 to Rack 7	R2-IB2-P11A to R7-IB1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling seven full racks together.

TABLE 112 Leaf Switch Connections for the Third Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	3 meters
	R3-IB3-P8B to R3-IB1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-IB1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-IB1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-IB1-P5A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9B to R5-IB1-P6A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10A to R6-IB1-P7A	10 meters
R3 IB3 to Rack 7	R3-IB3-P10B to R7-IB1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	3 meters
	R3-IB2-P8B to R3-IB1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-IB1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-IB1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-IB1-P5B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9B to R5-IB1-P6B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10A to R6-IB1-P7B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10B to R7-IB1-P8B	10 meters

The following table shows the cable connections for the fourth spine switch (R4-IB1) when cabling seven full racks together.

TABLE 113 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-IB1-P3A	3 meters
	R4-IB3-P8B to R4-IB1-P4A	
R4 IB3 to Rack 1	R4-IB3-P10B to R1-IB1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-IB1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-IB1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-IB1-P5A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9B to R6-IB1-P6A	5 meters
R4 IB3 to Rack 7	R4-IB3-P10A to R7-IB1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-IB1-P3B	3 meters
	R4-IB2-P8B to R4-IB1-P4B	
R4 IB2 to Rack 1	R4-IB2-P10B to R1-IB1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-IB1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-IB1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-IB1-P5B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9B to R6-IB1-P6B	5 meters
R4 IB2 to Rack 7	R4-IB2-P10A to R7-IB1-P7B	10 meters

The following table shows the cable connections for the fifth spine switch (R5-IB1) when cabling seven full racks together.

TABLE 114 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-IB1-P3A	3 meters
	R5-IB3-P8B to R5-IB1-P4A	
R5 IB3 to Rack 1	R5-IB3-P10A to R1-IB1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-IB1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-IB1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-IB1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-IB1-P5A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9B to R7-IB1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-IB1-P3B	3 meters
	R5-IB2-P8B to R5-IB1-P4B	
R5 IB2 to Rack 1	R5-IB2-P10A to R1-IB1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-IB1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-IB1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-IB1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-IB1-P5B	5 meters

Leaf Switch	Connection	Cable Length
R5 IB2 to Rack 7	R5-IB2-P9B to R7-IB1-P6B	5 meters

The following table shows the cable connections for the sixth spine switch (R6-IB1) when cabling seven full racks together.

TABLE 115 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-IB1-P3A	3 meters
	R6-IB3-P8B to R6-IB1-P4A	
R6 IB3 to Rack 1	R6-IB3-P9B to R1-IB1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-IB1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-IB1-P8A	5 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-IB1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-IB1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P9A to R7-IB1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-IB1-P3B	3 meters
	R6-IB2-P8B to R6-IB1-P4B	
R6 IB2 to Rack 1	R6-IB2-P9B to R1-IB1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-IB1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-IB1-P8B	5 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-IB1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-IB1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P9A to R7-IB1-P5B	5 meters

The following table shows the cable connections for the seventh spine switch (R7-IB1) when cabling seven full racks together.

TABLE 116 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-IB1-P3A	3 meters
	R7-IB3-P8B to R7-IB1-P4A	
R7 IB3 to Rack 1	R7-IB3-P9A to R1-IB1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-IB1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-IB1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-IB1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-IB1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-IB1-P10A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-IB1-P3B	3 meters

Leaf Switch	Connection	Cable Length
	R7-IB2-P8B to R7-IB1-P4B	
R7 IB2 to Rack 1	R7-IB2-P9A to R1-IB1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB3-P9B to R2-IB1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-IB1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-IB1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-IB1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB3-P11B to R6-IB1-P10B	5 meters

Eight-Rack Cabling

The following table shows the cable connections for the first spine switch (R1-IB1) when cabling eight full racks together.

TABLE 117 Leaf Switch Connections for the First Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-IB1-P3A	3 meters
R1 IB3 to Rack 2	R1-IB3-P8B to R2-IB1-P4A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9A to R3-IB1-P5A	5 meters
R1 IB3 to Rack 4	R1-IB3-P9B to R4-IB1-P6A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10A to R5-IB1-P7A	10 meters
R1 IB3 to Rack 6	R1-IB3-P10B to R6-IB1-P8A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11A to R7-IB1-P9A	10 meters
R1 IB3 to Rack 8	R1-IB3-P11B to R8-IB1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-IB1-P3B	3 meters
R1 IB2 to Rack 2	R1-IB2-P8B to R2-IB1-P4B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9A to R3-IB1-P5B	5 meters
R1 IB2 to Rack 4	R1-IB2-P9B to R4-IB1-P6B	10 meters
R1 IB2 to Rack 5	R1-IB2-P10A to R5-IB1-P7B	10 meters
R1 IB2 to Rack 6	R1-IB2-P10B to R6-IB1-P8B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11A to R7-IB1-P8B	10 meters
R1 IB2 to Rack 8	R1-IB2-P11B to R8-IB1-P10B	10 meters

The following table shows the cable connections for the second spine switch (R2-IB1) when cabling eight full racks together.

TABLE 118 Leaf Switch Connections for the Second Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-IB1-P3A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-IB1-P10A	5 meters

Leaf Switch	Connection	Cable Length
R2 IB3 to Rack 3	R2-IB3-P8B to R3-IB1-P4A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9A to R4-IB1-P5A	5 meters
R2 IB3 to Rack 5	R2-IB3-P9B to R5-IB1-P6A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10A to R6-IB1-P7A	10 meters
R2 IB3 to Rack 7	R2-IB3-P10B to R7-IB1-P8A	10 meters
R2 IB3 to Rack 8	R2-IB3-P11A to R8-IB1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-IB1-P3B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-IB1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P8B to R3-IB1-P4B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9A to R4-IB1-P5B	5 meters
R2 IB2 to Rack 5	R2-IB2-P9B to R5-IB1-P6B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10A to R6-IB1-P7B	10 meters
R2 IB2 to Rack 7	R2-IB2-P10B to R7-IB1-P8B	10 meters
R2 IB2 to Rack 8	R2-IB2-P11A to R8-IB1-P9B	10 meters

The following table shows the cable connections for the third spine switch (R3-IB1) when cabling eight full racks together.

TABLE 119 Leaf Switch Connections for the Third Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-IB1-P3A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-IB1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-IB1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P8B to R4-IB1-P4A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9A to R5-IB1-P5A	5 meters
R3 IB3 to Rack 6	R3-IB3-P9B to R6-IB1-P6A	5 meters
R3 IB3 to Rack 7	R3-IB3-P10A to R7-IB1-P7A	10 meters
R3 IB3 to Rack 8	R3-IB3-P10B to R8-IB1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-IB1-P3B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-IB1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-IB1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P8B to R4-IB1-P4B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9A to R5-IB1-P5B	5 meters
R3 IB2 to Rack 6	R3-IB2-P9B to R6-IB1-P6B	5 meters
R3 IB2 to Rack 7	R3-IB2-P10A to R7-IB1-P7B	10 meters
R3 IB2 to Rack 8	R3-IB2-P10B to R8-IB1-P8B	10 meters

The following table shows the cable connections for the fourth spine switch (R4-IB1) when cabling eight full racks together.

TABLE 120 Leaf Switch Connections for the Fourth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-IB1-P3A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-IB1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-IB1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-IB1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P8B to R5-IB1-P4A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9A to R6-IB1-P5A	5 meters
R4 IB3 to Rack 7	R4-IB3-P9B to R7-IB1-P6A	10 meters
R4 IB3 to Rack 8	R4-IB3-P10A to R8-IB1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-IB1-P3B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-IB1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-IB1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-IB1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P8B to R5-IB1-P4B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9A to R6-IB1-P5B	5 meters
R4 IB2 to Rack 7	R4-IB2-P9B to R7-IB1-P6B	10 meters
R4 IB2 to Rack 8	R4-IB2-P10A to R8-IB1-P7B	10 meters

The following table shows the cable connections for the fifth spine switch (R5-IB1) when cabling eight full racks together.

TABLE 121 Leaf Switch Connections for the Fifth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-IB1-P3A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-IB1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-IB1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-IB1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-IB1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P8B to R6-IB1-P4A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9A to R7-IB1-P5A	5 meters
R5 IB3 to Rack 8	R5-IB3-P9B to R8-IB1-P6A	10 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-IB1-P3B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-IB1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-IB1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-IB1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-IB1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P8B to R6-IB1-P4B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9A to R7-IB1-P5B	5 meters
R5 IB2 to Rack 8	R5-IB2-P9B to R8-IB1-P6B	10 meters

The following table shows the cable connections for the sixth spine switch (R6-IB1) when cabling eight full racks together.

TABLE 122 Leaf Switch Connections for the Sixth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-IB1-P3A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-IB1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-IB1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-IB1-P8A	5 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-IB1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-IB1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P8B to R7-IB1-P4A	5 meters
R6 IB3 to Rack 8	R6-IB3-P9A to R8-IB1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-IB1-P3B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-IB1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-IB1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-IB1-P8B	5 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-IB1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-IB1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P8B to R7-IB1-P4B	5 meters
R6 IB2 to Rack 8	R6-IB2-P9A to R8-IB1-P5B	5 meters

The following table shows the cable connections for the seventh spine switch (R7-IB1) when cabling eight full racks together.

TABLE 123 Leaf Switch Connections for the Seventh Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-IB1-P3A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-IB1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-IB1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-IB1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-IB1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-IB1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-IB1-P10A	5 meters
R7 IB3 to Rack 8	R7-IB3-P8B to R8-IB1-P4A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-IB1-P3B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-IB1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB3-P9B to R2-IB1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-IB1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-IB1-P8B	10 meters

Leaf Switch	Connection	Cable Length
R7 IB2 to Rack 5	R7-IB2-P11A to R5-IB1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB3-P11B to R6-IB1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB3-P8B to R8-IB1-P4B	5 meters

The following table shows the cable connections for the eighth spine switch (R8-IB1) when cabling eight full racks together.

TABLE 124 Leaf Switch Connections for the Eighth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 IB3 within Rack 8	R8-IB3-P8A to R8-IB1-P3A	3 meters
R8 IB3 to Rack 1	R8-IB3-P8B to R1-IB1-P4A	10 meters
R8 IB3 to Rack 2	R8-IB3-P9A to R2-IB1-P5A	10 meters
R8 IB3 to Rack 3	R8-IB3-P9B to R3-IB1-P6A	10 meters
R8 IB3 to Rack 4	R8-IB3-P10A to R4-IB1-P7A	10 meters
R8 IB3 to Rack 5	R8-IB3-P10B to R5-IB1-P8A	5 meters
R8 IB3 to Rack 6	R8-IB3-P11A to R6-IB1-P9A	5 meters
R8 IB3 to Rack 7	R8-IB3-P11B to R7-IB1-P10A	5 meters
R8 IB2 within Rack 8	R8-IB2-P8A to R8-IB1-P3B	3 meters
R8 IB2 to Rack 1	R8-IB2-P8B to R1-IB1-P4B	10 meters
R8 IB2 to Rack 2	R8-IB3-P9A to R2-IB1-P5B	10 meters
R8 IB2 to Rack 3	R8-IB2-P9B to R3-IB1-P6B	10 meters
R8 IB2 to Rack 4	R8-IB2-P10A to R4-IB1-P7B	10 meters
R8 IB2 to Rack 5	R8-IB2-P10B to R5-IB1-P8B	5 meters
R8 IB2 to Rack 6	R8-IB3-P11A to R6-IB1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB3-P1B to R7-IB1-P10B	5 meters

Index

T

U