Oracle® SuperCluster M8 and SuperCluster M7 Overview Guide

Part No: E58633-12
March 2020

ORACLE

Oracle SuperCluster M8 and SuperCluster M7 Overview Guide

Part No: E58633-12

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: E58633-12

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

- **Overview** – Provides information about Oracle SuperCluster M8 and SuperCluster M7 configurations and components, including domain configurations, CPU and memory resources, and network requirements.
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Experience with SuperCluster systems

---

Note - All hardware-related specifications in this guide are based on information for a typical deployment provided by Oracle at the time this guide was written. Oracle is not responsible for hardware problems that might result from following the typical deployment specifications in this document. For detailed information about preparing your site for SuperCluster M7 deployment, consult your hardware specification.

Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E58626_01/index.html.

Feedback

Provide feedback about this documentation at http://www.oracle.com/goto/docfeedback.

Understanding SuperCluster Hardware Component Configurations

Oracle SuperCluster M8 and SuperCluster M7 provide a secure enterprise and cloud infrastructure for both databases and applications. Both systems are engineered to combine compute, networking, and storage hardware with virtualization, operating system, and management software into a single system that is easy to deploy, secure, manage, and maintain.

This document provides information about Oracle SuperCluster M8 and SuperCluster M7 configurations and components, including domain configurations, CPU and memory resources, and network requirements. This document provides information that can assist you when completing the preinstallation SuperCluster Configuration Worksheets document.

These topics describe the SuperCluster M8 and SuperCluster M7 hardware components.

Single Compute Server Components

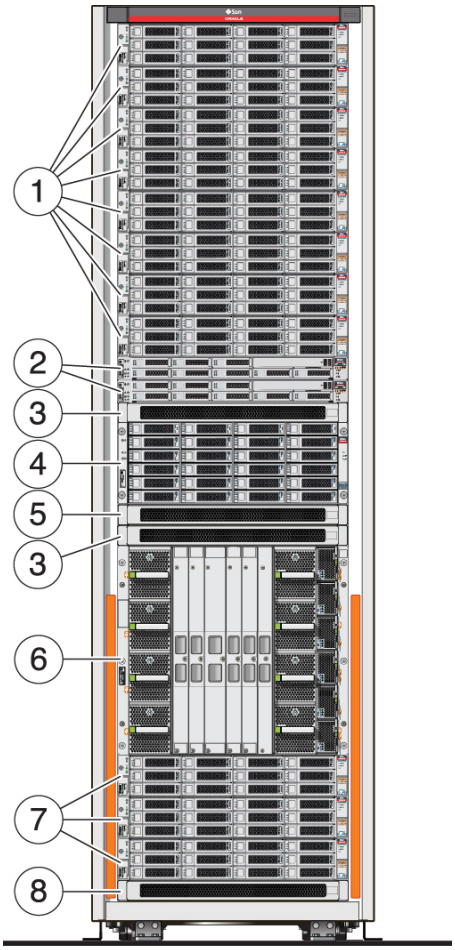

1	Space for up to eight additional storage servers
2	Storage controllers (2)
3	Sun Datacenter IB Switch 36 leaf switches (2)
4	Sun Disk Shelf
5	Ethernet management switch

6	Compute server
7	Storage servers (3)
8	IB spine switch

SuperCluster M8 and SuperCluster M7 with a single compute server comes with a minimum of three storage servers, located at the bottom of the rack. Up to eight additional storage servers can be added at the top of this rack. Two IB leaf switches, and optionally one IB spine switch are included in the system.

You can also expand a single compute server to add an additional compute server, so that you have a dual compute server system. However, these restrictions apply:

- Adding a second compute server to SuperCluster after the initial installation of the system requires a software reset and reinstallation process by an Oracle installer.
- You can only install one additional compute server to a single compute server system.
- Due to rack space limitations, you can add an additional compute server to a single server system only if you have six or fewer storage servers.
- The orderable option of an additional compute server contains two PDomains, with one CMIOU installed in PDomain 0, and with PDomain 1 empty. You can order additional CMIOUs that can be installed into the empty CMIOU slots. However, these CMIOUs follow the restrictions noted in "SuperCluster M8 and SuperCluster M7 Rules and Restrictions" on page 20, where additional CMIOUs installed after the initial installation of the system require a software reset and reinstallation process by an Oracle installer.

Refer to the *Oracle SuperCluster M7 Series Upgrade Configuration Worksheets* for information on upgrading your SuperCluster.

You can expand the amount of disk storage for your system using the expansion rack. See "Expansion Rack Components" on page 18 for more information.

You can connect up to eighteen SuperCluster M8 and SuperCluster M7 systems together, or a combination of SuperCluster M8, SuperCluster M7, Oracle Exadata, Oracle Big Data Appliance, or Oracle Exalogic systems on the same IB fabric, without the need for any external switches. However, you need the IB spine switch to connect additional systems to your SuperCluster M8 or SuperCluster M7. Refer to the *Oracle SuperCluster M8 and SuperCluster M7 Installation Guide* for more information.

Related Information

- "Dual Compute Server Components" on page 14
- "Compute Server Overview" on page 15

- “Understanding Storage Servers” on page 16
- “Expansion Rack Components” on page 18

Dual Compute Server Components

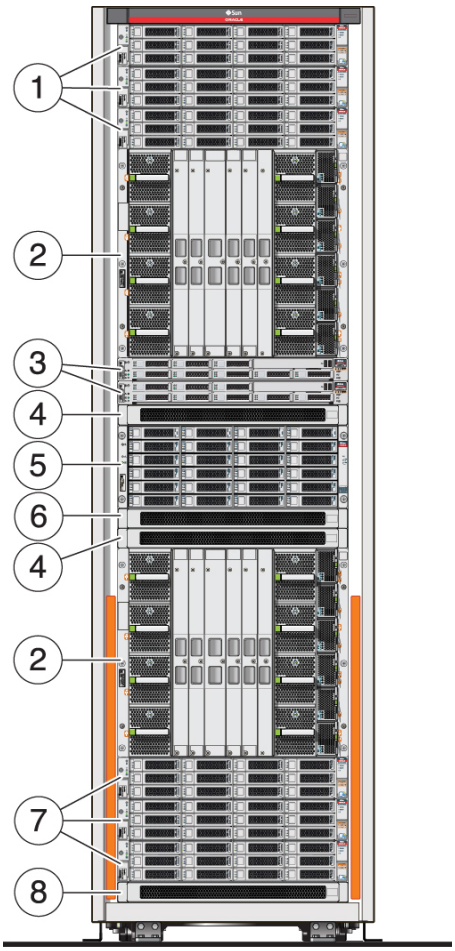

1	Space for up to three additional storage servers

2	Compute servers (2)
3	Storage controllers (2)
4	IB leaf switches (2)
5	Sun Disk Shelf
6	Ethernet management switch
7	Storage servers (3)
8	IB spine switch

SuperCluster M8 and SuperCluster M7 with two compute servers comes with a minimum of three storage servers, located at the bottom of the rack. Up to three additional storage servers can be added at the top of this rack. Two IB leaf switches, and one optional IB spine switch are included in the system.

You can expand the amount of disk storage for your system using the expansion rack. See "Expansion Rack Components" on page 18 for more information.

You can connect up to eighteen SuperCluster M8 and SuperCluster M7 systems together, or a combination of SuperCluster M8, SuperCluster M7, Oracle Exadata, Oracle Big Data Appliance,or Oracle Exalogic systems on the same IB fabric, without the need for any external switches. However, you need the IB spine switch to connect additional systems to your SuperCluster M8 or SuperCluster M7. Refer to the *Oracle SuperCluster M8 and SuperCluster M7 Installation Guide* for more information.

Related Information

- "Single Compute Server Components" on page 12
- "Compute Server Overview" on page 15
- "Understanding Storage Servers" on page 16
- "Expansion Rack Components" on page 18

Compute Server Overview

One or two SPARC M8 or SPARC M7 compute servers are installed in SuperCluster M8 or SuperCluster M7 respectively.

Each compute server provides these features:

- SPARC M8 or SPARC M7 32-core processors
- 64 GB or 32 GB DDR4-2400 memory DIMMs

- PCIe 3.0 capable expansion slots (some preinstalled with Ethernet and IB PCIe cards)
- Redundant SPs to monitor and control the system remotely

Each compute server is divided into two separate physical domains (PDomains). Each PDomain includes half of the possible processors, memory, and PCIe expansion slots in the chassis. Both PDomains operate as a separate server within the same chassis.

For details about the compute servers, refer to the SPARC M8 and SPARC M7 servers documentation library at: `https://docs.oracle.com/cd/E55211_01`

Related Information

- "Single Compute Server Components" on page 12
- "Dual Compute Server Components" on page 14

Understanding Storage Servers

Every SuperCluster M8 and SuperCluster M7 has a minimum of three storage servers installed in rack slots U2, U4, and U6. With elastic configurations, you can install additional storage servers in the rack, starting at rack slot U41 and moving down.

- Oracle Exadata X7-2L Storage Servers Storage Servers are supported in SuperCluster M8.
- Oracle Exadata X5-2L Storage Servers and Oracle Exadata X6-2L Storage Servers are supported in SuperCluster M7. You can install a combination of those storage server models in SuperCluster M7.

The storage servers are available with these types of storage:

- "Extreme Flash Storage" on page 16
- "High Capacity Storage" on page 17

Extreme Flash Storage

These are the components in the Extreme Flash version of the storage server:

- 8 NVMe PCIe 3.0 SSD Extreme Flash disks. Capacities of those Extreme Flash disks vary, depending on the type of storage server:

  - 1.6 TB (X5-2L)
  - 3.2 TB (X6-2L)
  - 6.4 TB (X7-2L)
- 2 IB 4 X QDR (40 Gb/s) IB ports (1 dual-port PCIe 3.0 HCA)
- 4 embedded Gigabit Ethernet ports
- 1 Ethernet port for ILOM for remote management
- Oracle Linux with Unbreakable Enterprise Kernel 2
- Oracle Exadata Storage Server Software

This table lists the storage capacities for a single storage server with Extreme Flash drives. To determine the system's total storage server capacity, multiply the single storage server capacity with the total number of storage servers in the system.

TABLE 1 Single Storage Server Capacity, Extreme Flash Version

Capacity Type	8 x 1.6 TB (X5-2L)	8 x 3.2 TB (X6-2L)	8 x 6.4 TB (X7-2L)
Raw capacity	12.8 TB	25.6 TB	51.2 TB
Usable mirrored capacity (ASM normal redundancy)	5 TB	10 TB	20 TB
Usable triple-mirrored capacity (ASM high redundancy)	4.3 TB	8.6 TB	17.2 TB

Related Information

- “High Capacity Storage” on page 17
- “Expansion Rack Components” on page 18

High Capacity Storage

Following are the components in the High Capacity version of the storage server:

- 12 7.2 K RPM High Capacity SAS drives. Capacities of those High Capacity disks vary, depending on the type of storage server:
  - 8 TB (X5-2L and X6-2L)
  - 10 TB (X7-2L)
- 4 flash accelerator PCIe cards disk controller HBA with 1 GB supercap-backed write cache
  - 4 x 1.6 TB (X5-2L and X6-2L)

  - 4 x 6.4 TB (X7-2L)
- 2 IB 4 X QDR (40 Gb/s) IB ports (1 dual-port PCIe 3.0 HCA)
- 4 embedded Gigabit Ethernet ports
- 1 Ethernet port for ILOM for remote management
- Oracle Linux with Unbreakable Enterprise Kernel 2
- Oracle Exadata Storage Server Software

This table lists the storage capacities for a single storage server with High Capacity drives. To determine the system's total storage server capacity, multiply the single storage server capacity with the total number of storage servers in the system.

TABLE 2 Storage Server Capacity, High Capacity Version

Capacity Type	12 x 8 TB (X5-2L or X6-2L)	12 x 10 TB (X7-2L)
Raw capacity	96 TB	120 TB
Usable mirrored capacity (ASM normal redundancy)	40 TB	50 TB
Usable triple-mirrored capacity (ASM high redundancy)	30 TB	37.5 TB

Related Information

- “Extreme Flash Storage” on page 16
- “Expansion Rack Components” on page 18

Expansion Rack Components

The expansion rack provides additional storage for SuperCluster M8 and SuperCluster M7. The additional storage can be used for backups, historical data, and unstructured data. Expansion racks can be used to add space to SuperCluster M8 and SuperCluster M7 as follows:

- Add new 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 an expansion rack.
- Split the expansion rack among multiple SuperCluster M8 or SuperCluster M7 systems.

The expansion rack is available as a quarter rack, with four storage servers. You can increase the number of storage servers in the expansion rack up to a maximum of 18 storage servers. The storage servers are available with either Extreme Flash or High Capacity storage.

Each expansion rack has the following components:

- 4 storage servers, with 8 Extreme Flash or 12 High Capacity drives in each storage server
- 2 IB switches
- Keyboard, video, and mouse (KVM) hardware
- 2 redundant 15 kVA PDUs (single-phase or three-phase, high voltage or low voltage)
- 1 Ethernet management switch

Related Information

- “Single Compute Server Components” on page 12
- “Dual Compute Server Components” on page 14
- “Compute Server Overview” on page 15
- “Understanding Storage Servers” on page 16

Asymmetric and Elastic Configurations

Asymmetric configurations allow for these configurations:

- Different number of CMIOUs in each compute server within the SuperCluster M8 or SuperCluster M7
- Different number of CMIOUs in each PDomain within each compute server
- Individual CMIOUs that can be added to PDomains in compute servers
- A second compute server that can be added to a single-compute server SuperCluster M8 or SuperCluster M7

Elastic configurations enable SuperCluster M8 or SuperCluster M7 to have the following customer-defined combinations of compute servers and Exadata Storage Servers:

- One compute server and three storage servers in a single system, expandable to eleven total storage servers
- Two compute servers and three storage servers in a single system, expandable to six total storage servers

See “SuperCluster M8 and SuperCluster M7 Rules and Restrictions” on page 20 for rules and restrictions on asymmetric and elastic configurations.

Also see “Asymmetric PDomain Configuration Overview” on page 26.

SuperCluster M8 and SuperCluster M7 Rules and Restrictions

The following rules and restrictions apply to hardware and software modifications to SuperCluster M8 and SuperCluster M7. Violating these restrictions can result in loss of warranty and support.

- These rules and restrictions apply to asymmetric configurations:
  - Adding a second compute server to a single compute server SuperCluster M8 or SuperCluster M7 after the initial installation of the system requires a software reset and reinstallation process by an Oracle installer. See "Single Compute Server Components" on page 12 for more information.
  - Within the entire SuperCluster M8 or SuperCluster M7, at least two PDomains must be populated, with a minimum of one CMIOU each. For a single compute server system, which has two PDomains total, both PDomains must be populated with at least one CMIOU. For a dual compute server system, which has four PDomains total, at least two of those four PDomains must be populated with at least one CMIOU. See "Understanding PDomain Configurations" on page 25 for more information.
  - You can have a different number of populated and unpopulated PDomains in each compute server. For example, you can have one compute server with two populated PDomains, and the second compute server with one populated and one unpopulated PDomain. See "Understanding PDomain Configurations" on page 25 for more information.
  - For populated PDomains, you can have a different number of CMIOUs in each PDomain in each compute server. For example, you can have one PDomain with one CMIOU and the second PDomain with two CMIOUs in the same compute server. See "Understanding PDomain Configurations" on page 25 for more information.

  Note - If you have a different number of CMIOUs in each populated PDomain, for configurations with only two PDomains, it is best practice to have an n+1 CMIOU layout for those PDomains (for example, one PDomain with one CMIOU and the second PDomain with two CMIOUs).

- The following restrictions apply to SuperCluster M8 and SuperCluster M7 elastic configurations:
  - You can have up to eleven total storage servers in a single compute server system or up to six total storage servers in a dual compute server system.
  - At least three storage servers must be installed in SuperCluster M8 or SuperCluster M7. The storage servers must all be the same type.

  - When adding storage servers, only these storage servers are supported within certain SuperCluster systems:
    - X7-2L Extreme Flash or High Capacity storage servers in SuperCluster M8
    - X5-2L or X6-2L Extreme Flash or High Capacity storage servers in SuperCluster M7.
  - Storage servers are installed in the rack in the following order:
    - Three storage servers are always installed in rack slots U2, U4, and U6.
    - Additional storage servers are installed starting at rack slot U41 and going down, ending at rack slot U37 in the dual compute server system or rack slot U27 in the single compute server system.
- SuperCluster M8 or SuperCluster M7 hardware cannot be modified or customized. There is one exception to this. The only allowed hardware modification to SuperCluster M8 or SuperCluster M7 is to the administrative Ethernet management switch included with SuperCluster M8 or SuperCluster M7. Customers may choose to do the following:
  - Replace the Ethernet management switch, at customer expense, with an equivalent Ethernet management 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 SuperCluster M8 or SuperCluster M7. 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 Ethernet management 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 Ethernet management switch in the rack can be turned off and unconnected to the data center network.
- The expansion rack can only be connected to SuperCluster M8, SuperCluster M7, or Oracle Exadata Database Machine. In SuperCluster M8 and SuperCluster M7, the expansion rack only supports databases running on the database domains.
- Standalone storage servers can only be connected to SuperCluster M8, SuperCluster M7 or Oracle Exadata Database Machine. In SuperCluster M8 or SuperCluster M7, the storage servers only support databases running on the database domains.
- Earlier Oracle Database releases can be run in Oracle Solaris 10 Branded Zones in Application Domains running Oracle Solaris 11. Refer to the Supported Virtualization matrix at https://www.oracle.com/database/technologies/virtualization-matrix.html for information about Oracle Database releases supported in Oracle Solaris 10 Branded Zones.

Non-Oracle databases can be run either natively in Application Domains running Oracle Solaris 11 or in Oracle Solaris 10 Branded Zones in Application Domains running 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 storage servers.
- Customers cannot update the firmware directly on the storage servers. The firmware is updated as part of a storage server patch.
- Customers may load additional software on the Database Domains on the compute 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 OS 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 OS other than by applying official patches and upgrades is not supported. IB-related packages should always be maintained at the officially supported release.
- SuperCluster M7 supports separate domains dedicated to applications, with high throughput/low latency access to the database domains through IB. 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 storage servers except as documented in the *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 compute servers can be used to connect to external nonstorage servers using iSCSI or NFS. However, the Fibre Channel Over Ethernet (FCoE) protocol is not supported.
- Only switches specified for use in SuperCluster M8, SuperCluster M7, Oracle Exadata, Oracle Exalogic Elastic Cloud, and Oracle Big Data Appliance may be connected to the SuperCluster M8 or SuperCluster M7 IB network. It is not supported to connect other IB switches, including third-party switches, to the SuperCluster M8 or SuperCluster M7 IB network. Only the IB networking topologies specified in SuperCluster M8 or SuperCluster M7 documentation are supported, and any other IB network topology is not supported.

  You may connect external servers that are not part of Oracle Engineered Systems to the IB switches in SuperCluster M8 or SuperCluster M7. However, it is your responsibility to upgrade and maintain the compatibility of the IB software of the external servers with the IB software release for SuperCluster M8 or SuperCluster M7. You should maintain the same release of IB software and operating system on the external server as on SuperCluster M8 or SuperCluster M7. If an IB fabric problem is encountered and an external server is connected, then you may be asked to remove the external server and reproduce the problem.

Related Information

- "Single Compute Server Components" on page 12
- "Dual Compute Server Components" on page 14
- "Compute Server Overview" on page 15
- "Understanding Storage Servers" on page 16
- "Expansion Rack Components" on page 18

Using Exalytics Software in SuperCluster M8 and SuperCluster M7

The Exalytics In-memory Machine T5-8 is no longer available for purchase. If you want to run the Exalytics software on a SPARC-based platform, you can install and run the Exalytics software in an Application Domain in SuperCluster M8 or SuperCluster M7 (either an Application Dedicated Domain or an Application I/O Domain). The following applies when running the Exalytics software in an Application Domain in SuperCluster M8 or SuperCluster M7:

- You can run the Exalytics software on any Application Domain that you have currently set up for other purposes. In addition, the Exalytics software, and therefore the Application Domain running the Exalytics software, does not have to be mirrored or clustered.

  As an example, consider a configuration where you have a SuperCluster M8 or SuperCluster M7 with a single compute server, with the following asymmetric configuration:

  - **PDomain 1** — Contains one CMIOU, with 32 cores and 512 GB of memory, configured with one Database Dedicated Domain
  - **PDomain 2** — Contains three CMIOUs, each with 32 cores and 512 GB of memory, configured with the following domains:
    - One Database Dedicated Domain, using resources from one CMIOU (32 cores and 512 GB of memory)
    - One Application Dedicated Domain, using resources from the remaining two CMIOUs (64 cores and 1 TB of memory total)

  In this example configuration, the Database Dedicated Domains on PDomains 1 and 2 run Oracle DB RAC (are part of a cluster), and the Application Dedicated Domain on PDomain 2 is available to run the Exalytics software.

  Similar configurations are supported for Application Domains running the Exalytics software, such as a second compute server that is set up only with Application Domains

running the Exalytics software, or Application I/O Domains that are created specifically to run Exalytics software.

- You can set up an entire SuperCluster specifically to run only Exalytics software. In this situation, only Application Domains are configured, with each Application Domain running Exalytics software. Every SuperCluster M8 and SuperCluster M7 has a minimum of three storage servers, which are accessed only by Database Domains. However, you are only required to license Exadata Storage Server disks when they are actually used, which is not be the case in this scenario.

  In practice, though, it is more sensible to include two or more Database Domains to take advantage of the included Exadata Storage Servers.
- You must install the following assembler package in order to use Exalytics software on an Application Domain in your SuperCluster M8 or SuperCluster M7:

  ```
  # pkg install pkg:/developer/assembler
  ```

Refer to the Exalytics documentation to install and set up the Exalytics software on an Application Domain, available here:

http://docs.oracle.com/cd/E41246_01/index.htm

Understanding PDomain Configurations

These topics describe SuperCluster PDomain configurations.

PDomain Overview

Each compute server is split into two PDomains, where the bottom four CMIOU slots are part of PDomain 0, and the top four CMIOU slots are part of the PDomain 1.

A PDomain operates like an independent server that has full hardware isolation from the other PDomain in the server. For example, you can reboot one PDomain while the other PDomain on a server continues to operate.

You can have from one to four CMIOUs in each PDomain, or you can have an empty PDomain that you can populate later.

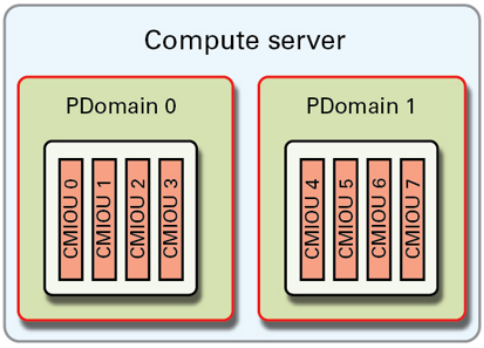

Related Information

- “Asymmetric PDomain Configuration Overview” on page 26
- “Understanding PDomain PCIe Card Configurations” on page 33
- “Compute Server Overview” on page 15

Asymmetric PDomain Configuration Overview

These asymmetric configurations are supported:

- Compute servers with asymmetric PDomains configurations. For example:
  - First compute server with two populated PDomains
  - Second compute server with one populated and one unpopulated PDomain

  Another example:
  - First compute server with eight CMIOUs
  - Second compute server with four CMIOUs
- PDomains with asymmetric CMIOU configurations. For example, within a compute server:
  - PDomain 0 with 1 CMIOU
  - PDomain 1 with 2 CMIOUs

However, when ordering a SuperCluster M8 or SuperCluster M7, you typically are provided with symmetric PDomain and CMIOU configurations. To create asymmetric configurations, order additional individual CMIOUs as part of your initial order. Those CMIOUs will be installed in the appropriate slots to create the asymmetric configuration.

You cannot mix SPARC M8 or SPARC M7 compute server CMIOUs, as described in “SuperCluster M8 and SuperCluster M7 Rules and Restrictions” on page 20.

For example, assume you want two compute servers, and you want these asymmetric configurations on those compute servers:

- Compute server 1:
  - PDomain 0 — 1 CMIOU
  - PDomain 1 — 2 CMIOUs
- Compute server 2:
  - PDomain 0 — 3 CMIOUs
  - PDomain 1 — 4 CMIOUs

To create those asymmetric configurations, you can order a SuperCluster M8 or SuperCluster M7 with the following symmetric configurations, and add the necessary CMIOUs to create the asymmetric configurations:

- Compute server 1:
  - PDomain 0 — 1 CMIOU
  - PDomain 1 — 1 CMIOU
  - 1 extra CMIOU to add to PDomain 1
- Compute server 2:
  - PDomain 0 — 3 CMIOUs
  - PDomain 1 — 3 CMIOUs
  - 1 extra CMIOU to add to PDomain 1

In addition, by having the additional CMIOUs installed as part of the initial installation, your Oracle installer sets up your domain configurations based on the total number of CMIOUs in each PDomain that are part of the final asymmetric configuration. If you order additional CMIOUs after your system has been installed, contact Oracle to request a software reset and reinstallation process, so that the domain configuration is changed to reflect the new CMIOUs.

Refer to the *Oracle SuperCluster M7 Series Upgrade Configuration Worksheets* for information on upgrading your SuperCluster.

Related Information

- "Understanding PDomain CMIOU Configurations" on page 27
- "Understanding PDomain PCIe Card Configurations" on page 33
- "Understanding Logical Domain Configurations" on page 43

Understanding PDomain CMIOU Configurations

CMIOU stands for CPU, memory, and I/O unit. Each CMIOU provides the CPU (1 CMP), memory (16 DIMM slots) resources, and an I/O hub. Each PDomain can include zero to four CMIOUs.

There are several PDomain configurations, based on these factors:

- Number of compute servers in SuperCluster M8 or SuperCluster M7
- Number of PDomains in each compute server
- Number of CMIOUs in each PDomain

These topics describe SuperCluster PDomain configurations:

Understanding Single Compute Server Configurations (R1 Configurations)

The R1 configurations are available for SuperClusters with a single compute server.

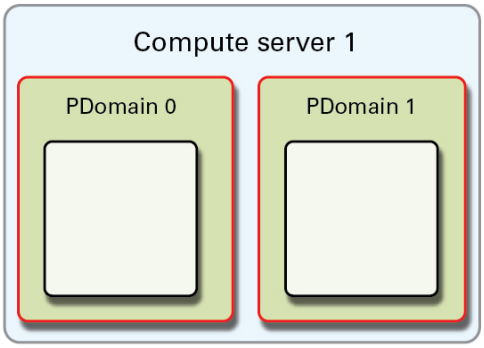

The R1-1 PDomain configuration is the only available configuration for the R1 PDomain configurations.

CMIOUs in Both PDomains (R1-1 Configuration)

The R1-1 PDomain configuration is the only available configuration for a R1 PDomain configuration.

The R1-1 PDomain configuration has these characteristics:

- Two populated PDomains in a single compute server
- One to four CMIOUs in each PDomain

This figure shows the CMIOU slots on each PDomain in this configuration.

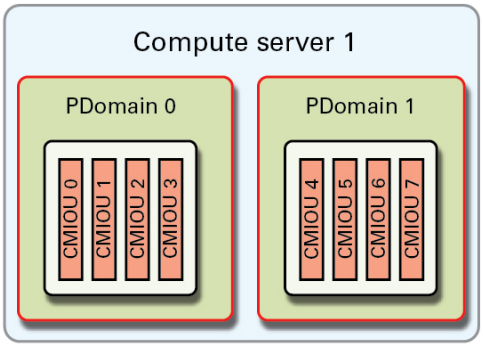

Related Information

- “Understanding PDomain PCIe Card Configurations” on page 33
- “Understanding Logical Domain Configurations” on page 43

Understanding Dual Compute Server Configurations (R2 Configurations)

The R2 configurations are available for SuperClusters with two compute servers.

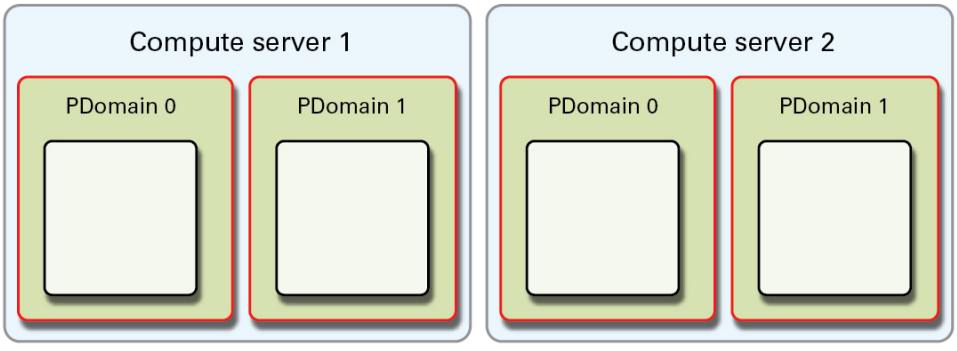

These choices are available for the R2 configuration, depending on which PDomains are populated with CMIOUs:

- “CMIOUs in Both PDomains (R2-1 Configuration)” on page 30
- “CMIOUs Only in PDomain 0 in Both Compute Servers (R2-2 Configuration)” on page 31
- “CMIOUs in PDomain 0 in Compute Server 1, and in PDomains 0 and 1 in Compute Server 2 (R2-3 Configuration)” on page 31
- “CMIOUs in PDomain 0 and 1 in Compute Server 1, and in PDomain 0 in Compute Server 2 (R2-4 Configuration)” on page 32

CMIOUs in Both PDomains (R2-1 Configuration)

The R2-1 PDomain configuration has these characteristics:

- Four populated PDomains across two compute servers
- One to four CMIOUs in each populated PDomain

This figure shows the CMIOU slots on each PDomain in this configuration.

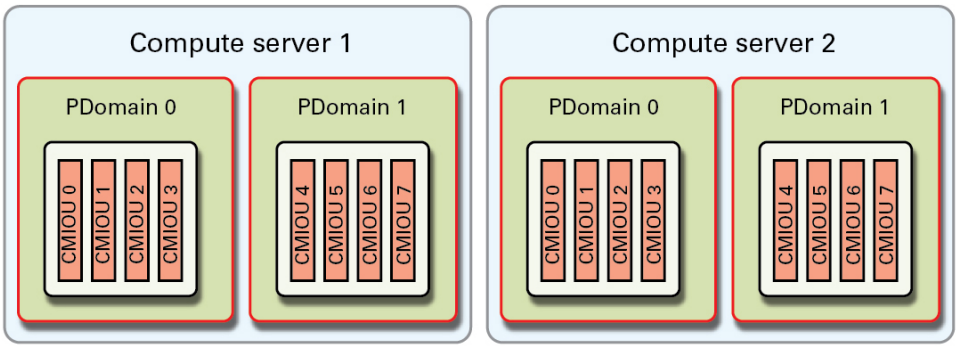

Related Information

- “Understanding PDomain PCIe Card Configurations” on page 33
- “Understanding Logical Domain Configurations” on page 43

CMIOUs Only in PDomain 0 in Both Compute Servers (R2-2 Configuration)

The R2-2 PDomain configuration has these characteristics:

- Two populated PDomains across two compute servers
- One to four CMIOUs in each populated PDomain

This figure shows the CMIOU slots on each PDomain in this configuration.

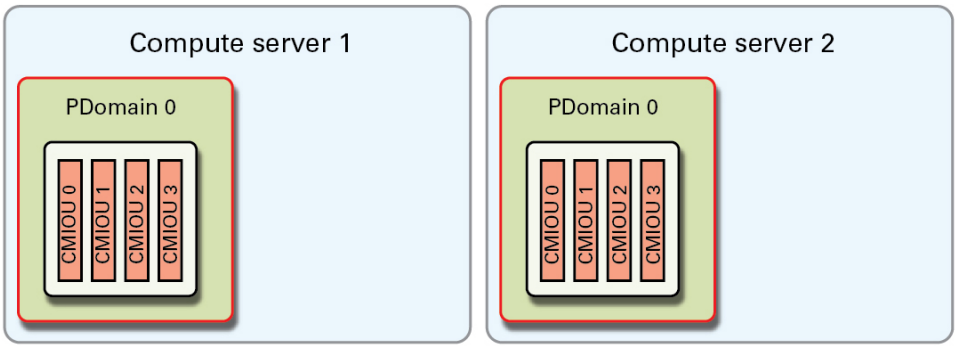

Related Information

- “Understanding PDomain PCIe Card Configurations” on page 33
- “Understanding Logical Domain Configurations” on page 43

CMIOUs in PDomain 0 in Compute Server 1, and in PDomains 0 and 1 in Compute Server 2 (R2-3 Configuration)

The R2-3 PDomain configuration has these characteristics:

- Populated PDomain 0 in compute server 1, and populated PDomains 0 and 1 in compute server 2

- One to four CMIOUs in each populated PDomain

This figure shows the CMIOU slots on each PDomain in this configuration.

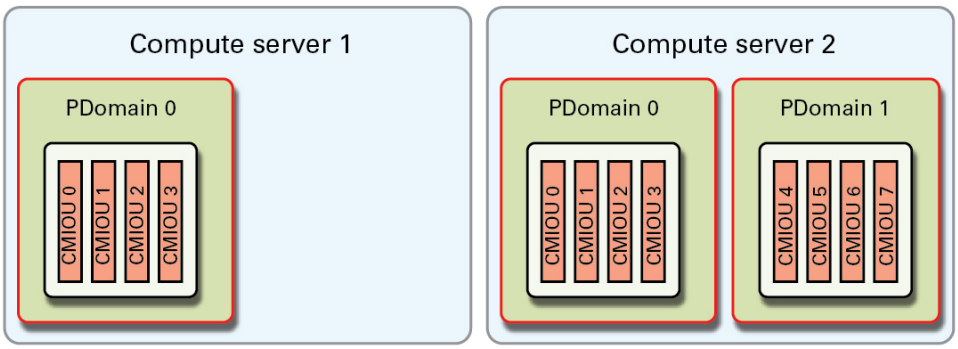

Related Information

- “Understanding PDomain PCIe Card Configurations” on page 33
- “Understanding Logical Domain Configurations” on page 43

CMIOUs in PDomain 0 and 1 in Compute Server 1, and in PDomain 0 in Compute Server 2 (R2-4 Configuration)

This configuration is one of the R2 PDomain configurations (see “Understanding Dual Compute Server Configurations (R2 Configurations)” on page 29).

The R2-4 PDomain configuration has these characteristics:

- Populated PDomains 0 and 1 in compute server 1, and populated PDomain 0 in compute server 2
- One to four CMIOUs in each populated PDomain

This figure shows the CMIOU slots on each PDomain in this configuration.

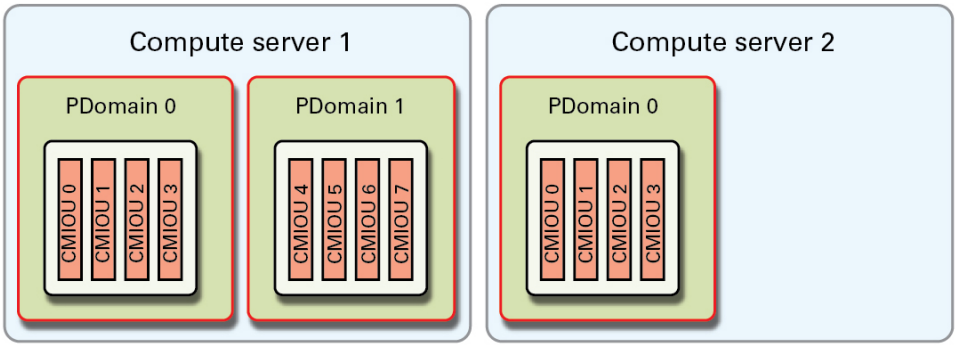

Related Information

- “Understanding PDomain PCIe Card Configurations” on page 33
- “Understanding Logical Domain Configurations” on page 43

Understanding PDomain PCIe Card Configurations

Each CMIOU has three PCIe slots. When present, the following cards are installed in certain PCIe slots and are used to connect to these networks:

- **1GbE NICs, installed in PCIe slot 1** — Connect to the 1GbE management network
- **10GbE NICs, installed in PCIe slot 2** — Connect to the 10GbE client access network
- **IB HCAs, installed in PCIe slot 3** — Connect to the private IB network

Optional Fibre Channel PCIe cards are also available to facilitate migration of data from legacy storage subsystems to the storage servers integrated with SuperCluster M8 or SuperCluster M7 for Database Domains, or to access SAN-based storage for the Application Domains. Fibre Channel PCIe cards can be installed in any open PCIe slot 1 in the CMIOUs installed in your system. Refer to the *Oracle SuperCluster M8 and SuperCluster M7 Installation Guide* for more information.

These topics show the factory default PCIe card locations based on how many CMIOUs are installed:

- “One CMIOU PCIe Card Configurations” on page 34
- “Two CMIOU PCIe Card Configurations (SuperCluster M8)” on page 35

One CMIOU PCIe Card Configurations

These diagrams show the PCIe card configurations for PDomains 0 and 1 with one CMIOU installed. Note that PDomains 0 and 1 can have a different number of CMIOUs, so one diagram or the other might not apply to your configuration.

For the possible domain configurations for this configuration, see "Domain Configurations for PDomains With One CMIOU" on page 54

FIGURE 1 PDomain 0

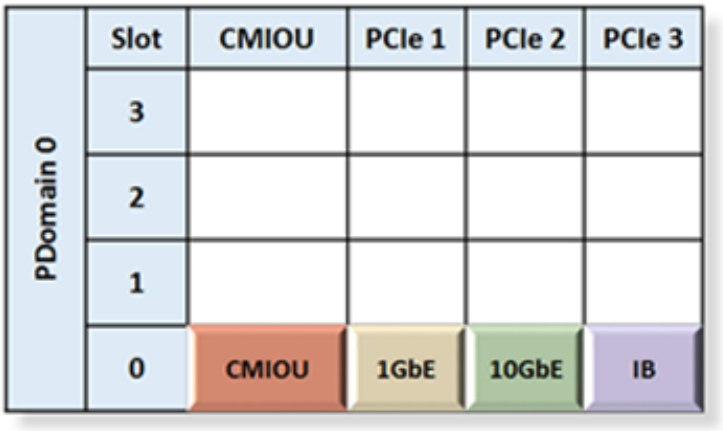

FIGURE 2 PDomain 1

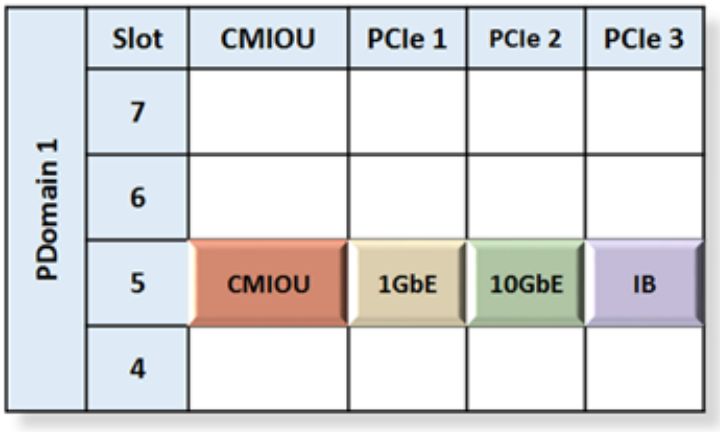

Related Information

- "Understanding SuperCluster Networks" on page 85
- "Domain Configurations for PDomains With One CMIOU" on page 54

Two CMIOU PCIe Card Configurations (SuperCluster M8)

These diagrams show the PCIe card configurations for PDomains 0 and 1 with two CMIOUs installed in SuperCluster M8. Note that PDomains 0 and 1 can have a different number of CMIOUs, so one diagram or the other might not apply to your configuration.

For the possible domain configurations for this configuration, see "Domain Configurations for PDomains With Two CMIOUs" on page 56.

FIGURE 3 PDomain 0

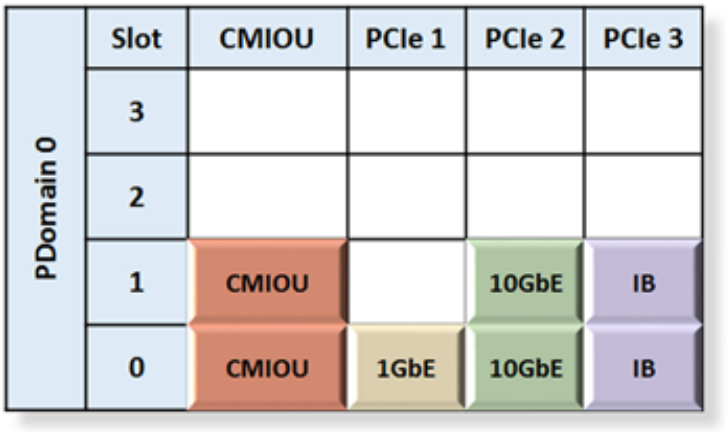

FIGURE 4 PDomain 1

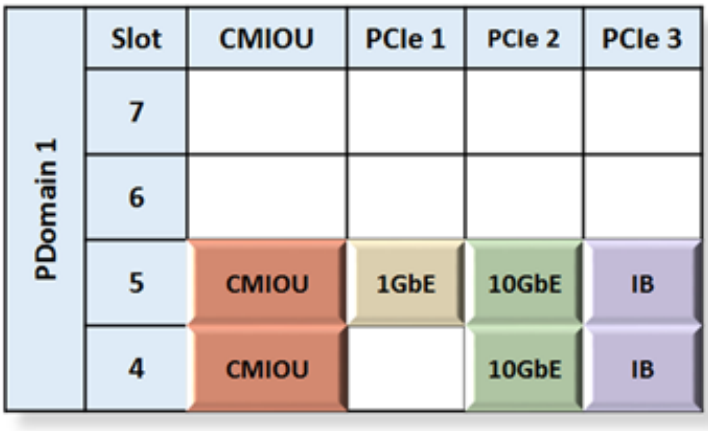

Related Information

- “Understanding SuperCluster Networks” on page 85
- “Domain Configurations for PDomains With Two CMIOUs” on page 56

Two CMIOUs PCIe Card Configurations (SuperCluster M7)

These diagrams show the PCIe card configurations for PDomains 0 and 1 with two CMIOUs installed in SuperCluster M7. Note that PDomains 0 and 1 can have a different number of CMIOUs, so one diagram or the other might not apply to your configuration.

For the possible domain configurations for this configuration, see “Domain Configurations for PDomains With Two CMIOUs” on page 56.

FIGURE 5 PDomain 0

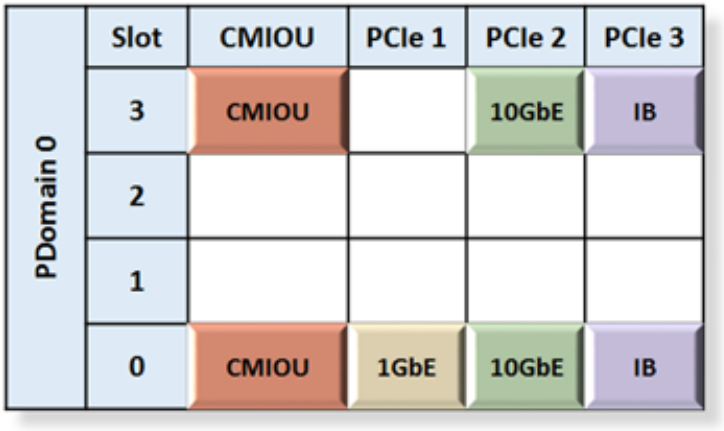

FIGURE 6 PDomain 1

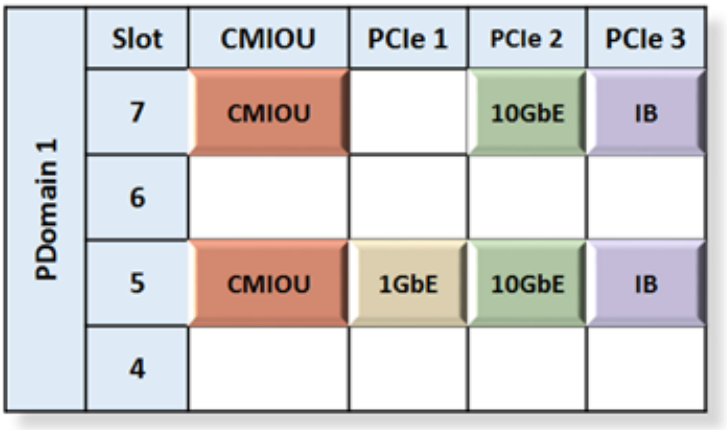

Related Information

- "Understanding SuperCluster Networks" on page 85
- "Domain Configurations for PDomains With Two CMIOUs" on page 56

Three CMIOU PCIe Card Configurations

These diagrams show the PCIe card configurations for PDomains 0 and 1 with three CMIOUs installed. Note that PDomains 0 and 1 can have a different number of CMIOUs, so one diagram or the other might not apply to your configuration.

For the possible domain configurations for this configuration, see "Domain Configurations for PDomains With Three CMIOUs" on page 60.

FIGURE 7 PDomain 0

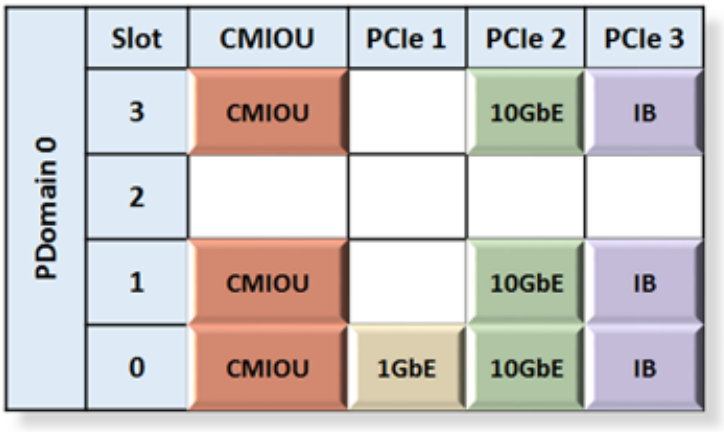

FIGURE 8 PDomain 1

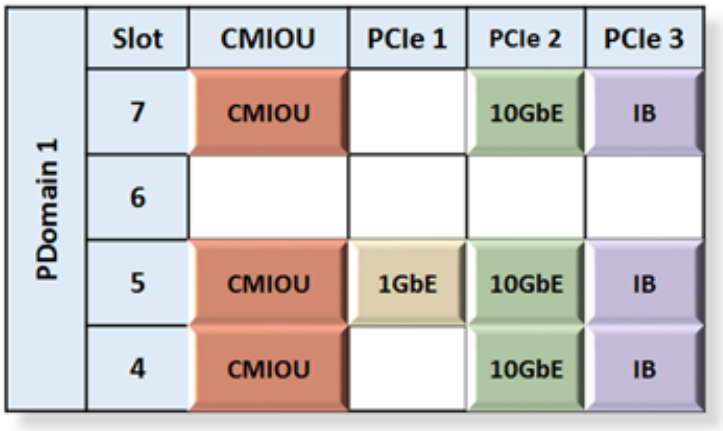

Related Information

- “Understanding SuperCluster Networks” on page 85
- “Domain Configurations for PDomains With Three CMIOUs” on page 60

Four CMIOUs PCIe Card Configurations

These diagrams show the PCIe card configurations for PDomains 0 and 1 with four CMIOUs installed. Note that PDomains 0 and 1 can have a different number of CMIOUs, so one diagram or the other might not apply to your configuration.

For the possible domain configurations for this configuration, see “Domain Configurations for PDomains With Four CMIOUs” on page 66 .

FIGURE 9 PDomain 0

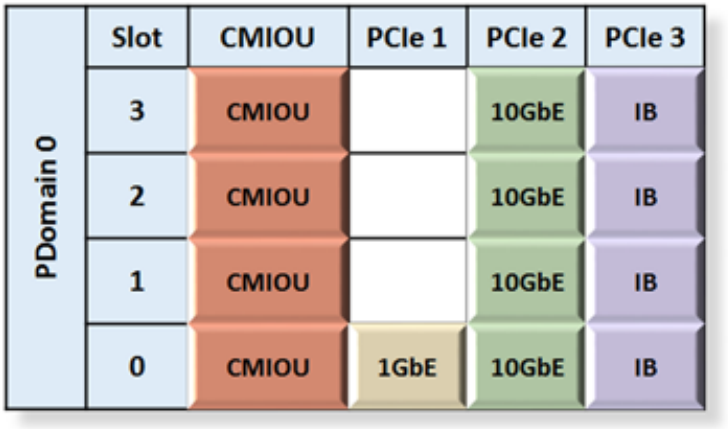

FIGURE 10 PDomain 1

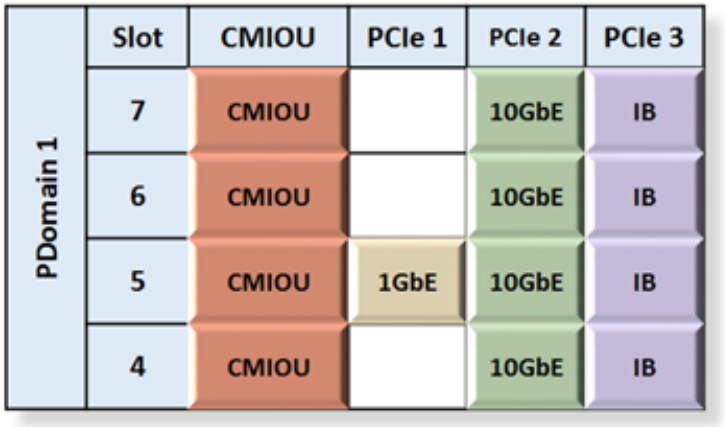

Related Information

- “Understanding SuperCluster Networks” on page 85
- “Domain Configurations for PDomains With Four CMIOUs” on page 66

Understanding Logical Domain Configurations

These topics describe SuperCluster logical domain configurations.

Logical Domain Configuration Overview

Using Oracle VM Server for SPARC (logical domains) technology, SuperCluster provides flexible virtualization with both dedicated domains, established during initial install, and dynamically created SR-IOV domains created any time as required. In SuperCluster terminology, the term *domain* refers to a logical domain. Additionally, there are different types of domains which are described in subsequent sections.

Note - Oracle Solaris Zones are also supported for an additional layer of virtualization.

The number of domains supported on each compute server depends on the number of CMIOUs that are in each PDomain:

- **PDomains with one CMIOU** — One domain
- **PDomains with two CMIOUs** — One or two domains
- **PDomains with three CMIOUs** — One to three domains
- **PDomains with four CMIOUs** — One to four domains

For SuperCluster, there are two main types of domains:

Dedicated Domains

Dedicated domains have dedicated CPU and memory resources for each domain, and one or more dedicated InfiniBand host channel adapters (HCAs), 10 GbE NICs, and optionally fibre channel HBAs. Because the key CPU, memory, and I/O resources are not shared, these domains are referred to as dedicated domains.

Two types of dedicated domains are available to run workloads:

- **Application Domains** – Run the Oracle Solaris 11 OS. and any applications that are supported on Oracle's SPARC/Oracle Solaris platforms. You can have Oracle Solaris 10 branded zones installed in Oracle Solaris 11 domains..
- **Database Domains** – Run the Oracle Database software, and directly leverage the Oracle Exadata Storage Servers. A Database Domain can also include Database zones.

Each domain has access to the 10GbE NICs and IB HCAs, with connections to those networks occurring in the following manner:

- To the 10GbE client access network through the physical ports on each 10GbE NIC
- To the IB network through the physical ports on each IB HCA

This diagram shows this concept on a SuperCluster with four domains.

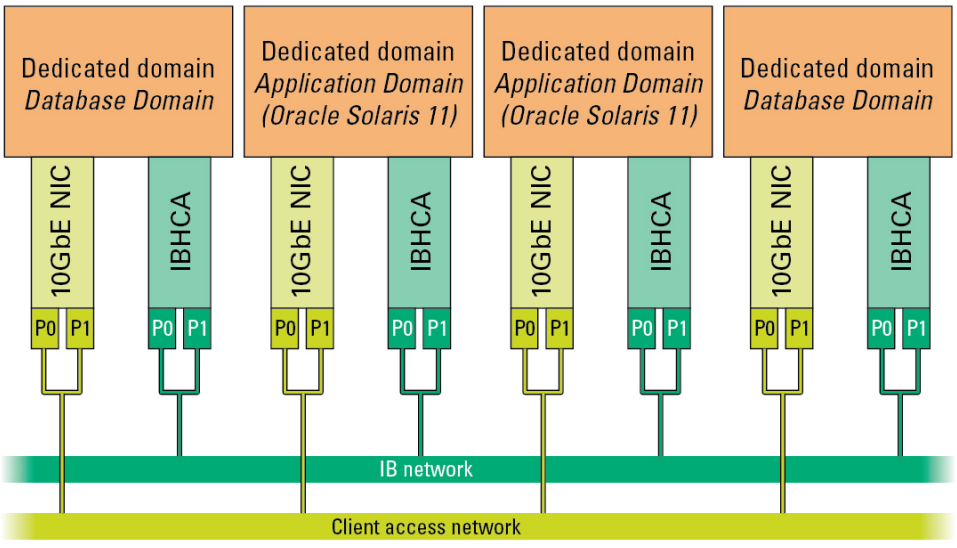

In addition, the first domain (the control domain) in each PDomain has direct access to the management network through the physical port on 1GbE NICs, and the other domains in each PDomain connect to the management network through VNETs. See "Understanding PDomain PCIe Card Configurations" on page 33 and "SuperCluster Networks Overview" on page 85.

For dedicated domains, the domain configuration (the number of domains and the SuperCluster-specific types assigned to each) are set at the time of the initial installation, and can only be changed by an Oracle representative.

Related Information

- "SR-IOV Domain Types" on page 45

SR-IOV Domain Types

SR-IOV (Single-Root I/O Virtualization) technology enables you to dynamically deploy or deleted SR-IOV I/O Domains as needed over the life cycle of the system. This is unlike dedicated domains which are created during the installation and cannot be easily changed.

These SR-IOV domain types are supported:

- "Root Domains" on page 45
- "I/O Domains" on page 49

Root Domains

The sole purpose of Root Domains is to provide shared I/O resources to user-created domains (I/O Domains). Like dedicated domains, root domains are created during the installation by Oracle personnel. Root Domains own InfiniBand HCA and 10 GbE NIC resources (one of each per root domain), and in some cases an optional fibre channel HBA. Unlike dedicated domains, Root Domains do not host user applications. For that reason they are assigned only a small allocation of CPU and memory. Additional HCAs and NICs cannot be added to Root Domains.

When deciding which domains will be a Root Domain, the last domain must always be the first Root Domain, and the remaining domains can be any combination of Root Domains or dedicated domains. However, a domain can only be a Root Domain if it has either one or two IB HCAs associated with it. A domain cannot be a Root Domain if it has more than two IB HCAs associated with it. If a domain has more than two IB HCAs associated with it.

Note - For SuperCluster M8, only one HCA is supported in a Root Domain.

Note - Even though a domain with two IB HCAs is valid for a Root Domain, domains with only one IB HCA should be used as Root Domains. When a Root Domain has a single IB HCA, fewer I/O Domains have dependencies on the I/O devices provided by that Root Domain. Flexibility around high availability also increases with Root Domains with one IB HCA.

A certain amount of CPU and memory resources are reserved for each Root Domain, depending on which domain is being used as a Root Domain in the domain configuration and the number of IB HCAs and 10GbE NICs that are associated with that Root Domain:

- The last domain in a domain configuration:
  - Two cores and 32 GB of memory reserved for a Root Domain with one IB HCA and 10GbE NIC
  - Four cores and 64 GB of memory reserved for a Root Domain with two IB HCAs (SuperCluster M7 only) and 10GbE NICs
- Any other domain in a domain configuration — One core and 16 GB of memory reserved for any remaining Root Domains with one IB HCA and 10GbE NIC

Note - The amount of CPU and memory reserved for Root Domains is sufficient to support only the physical functions in each Root Domain. There is insufficient CPU or memory resources to support zones or applications in Root Domains, so zones and applications are supported only in the I/O Domains.

The remaining CPU and memory resources associated with each Root Domain are parked in CPU and memory repositories, as shown in this diagram.

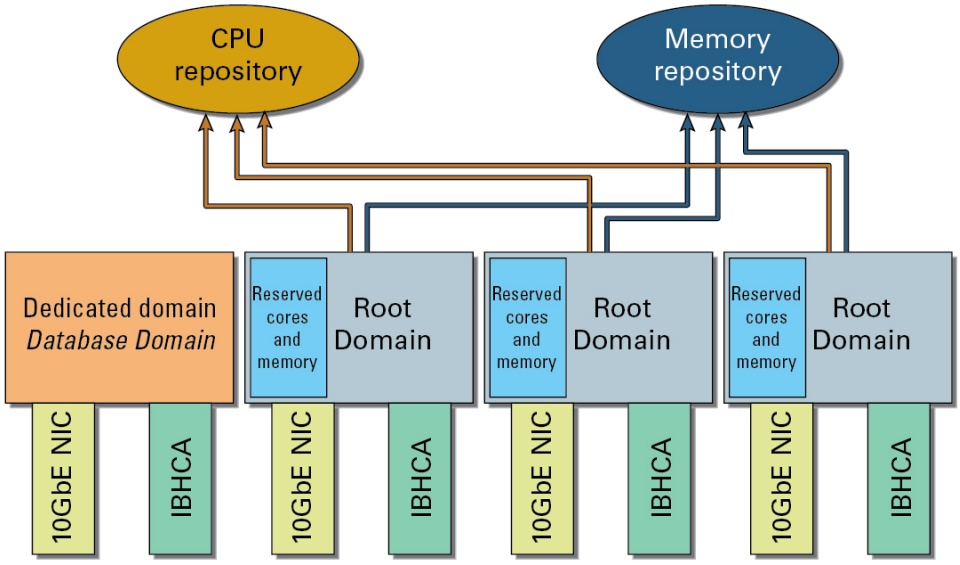

CPU and memory repositories contain resources not only from the Root Domains, but also any parked resources from the dedicated domains. Whether CPU and memory resources originated from dedicated domains or from Root Domains, once those resources have been parked in the CPU and memory repositories, those resources are no longer associated with their originating domain. These resources become available to I/O Domains.

In addition, CPU and memory repositories contain parked resources only from the compute server that contains the domains providing those parked resources. In other words, if you have two compute servers and both compute servers have Root Domains, there are two sets of CPU and memory repositories, where each compute server has its own CPU and memory repositories with parked resources.

For example, assume you have four domains on your compute server, with three of the four domains as Root Domains, as shown in the previous diagram. Assume each domain has the following IB HCAs and 10GbE NICs, and the following CPU and memory resources:

- One IB HCA and one 10GbE NIC
- 32 cores
- 512 GB of memory

In this situation, the following CPU and memory resources are reserved for each Root Domain, with the remaining resources available for the CPU and memory repositories:

- Two cores and 32 GB of memory reserved for the last Root Domains in this configuration. 30 cores and 480 GB of memory available from this Root Domain for the CPU and memory repositories.
- One core and 16 GB of memory reserved for the second and third Root Domains in this configuration.
  - 31 cores and 496 GB of memory available from each of these Root Domains for the CPU and memory repositories.
  - A total of 62 cores (31 x 2) and 992 GB of memory (496 GB x 2) available for the CPU and memory repositories from these two Root Domains.

A total of 92 cores (30 + 62 cores) are therefore parked in the CPU repository, and 1472 GB of memory (480 + 992 GB of memory) are parked in the memory repository and are available for the I/O Domains.

With Root Domains, connections to the three networks (client access, IB, and management) go through the physical ports on NIC and HCA, just as they did with dedicated domains. However, the 10GbE NICs and IB HCAs used with Root Domains must also be SR-IOV compliant. SR-IOV compliant cards enable VFs to be created on each card, where the virtualization occurs in the card itself. VFs are not created on the 1GbE NIC for the management network.

The VFs from each Root Domain are parked in the IB VF and 10GbE VF repositories, similar to the CPU and memory repositories, as shown in the following graphic.

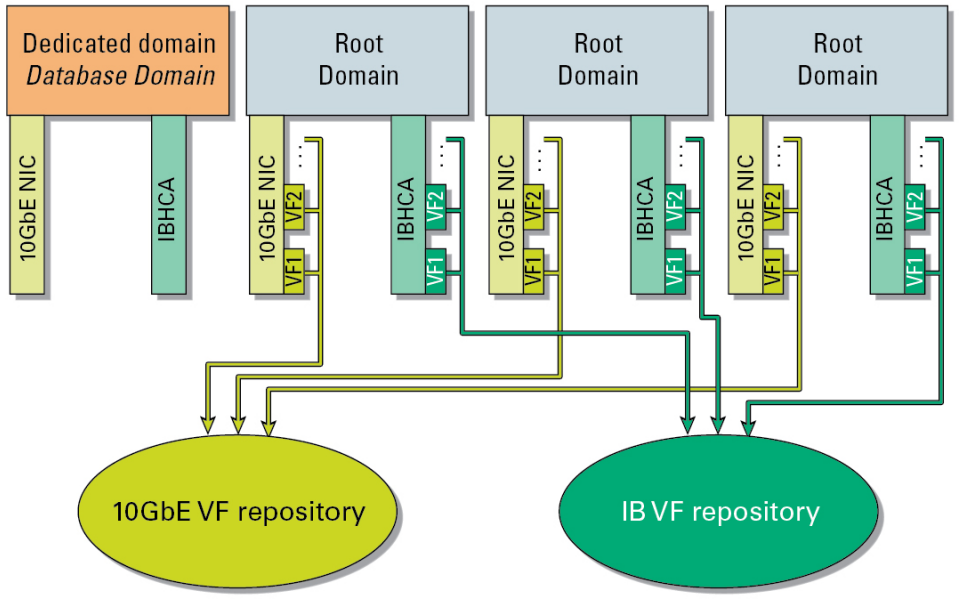

Even though the VFs from each Root Domain are parked in the VF repositories, the VFs are created on each 10GbE NIC and IB HCA, so those VFs are associated with the Root Domain that contains those specific 10GbE NIC and IB HCA cards. For example, looking at the example configuration in the previous graphic, the VFs created on the last (rightmost) 10GbE NIC and IB HCA are associated with the last Root Domain.

Related Information

- "I/O Domains" on page 49
- "Dedicated Domains" on page 44

I/O Domains

I/O Domains are SR-IOV type domains that you can dynamically deploy or delete any time over the life cycle of the system. You can easily administer I/O Domains using the browser-based SuperCluster Virtual Assistant. For additional details about administering I/O Domains, refer to the *Oracle I/O Domain Administration Guide*.

I/O Domains use virtual functions (VFs) that are provided by the Root Domain PFs.

When I/O Domains are created, they are specified as one of these I/O Domain types:

- Application I/O Domain running Oracle Solaris 11
- Database I/O Domain
- Database Zone I/O Domain

The CPU and memory resources owned by an I/O Domain are assigned from the CPU and memory repositories (the cores and memory released from Root Domains on the system) when an I/O Domain is created, as shown in this diagram.

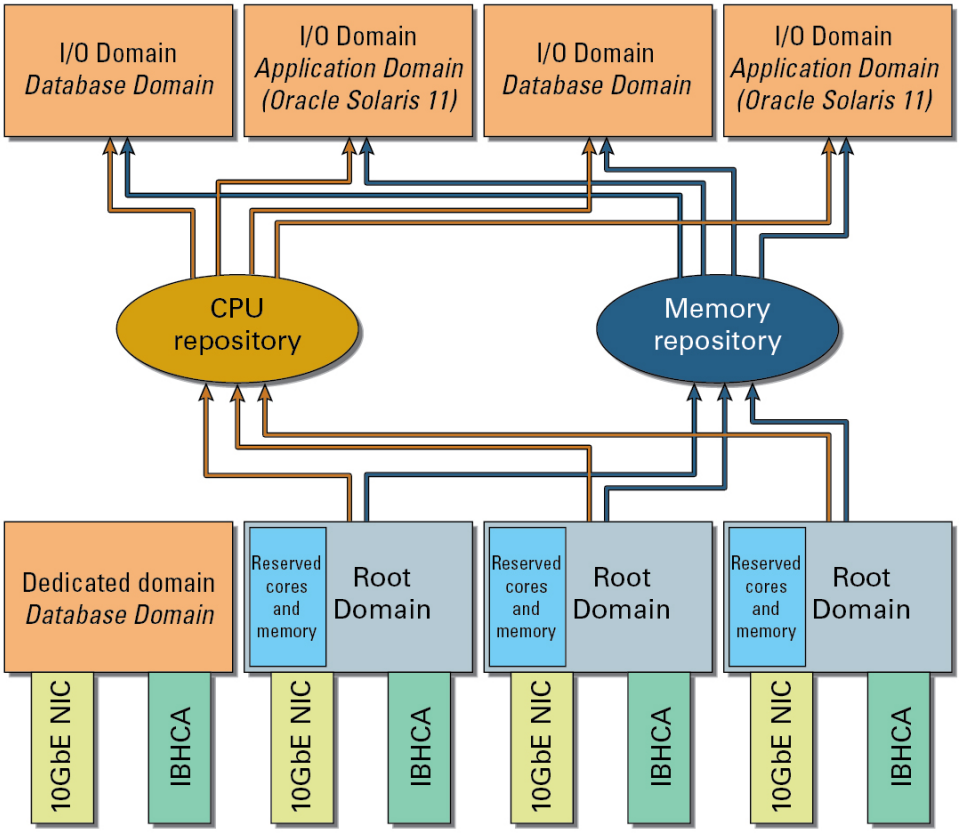

You use the SuperCluster Virtual Assistant to assign the CPU and memory resources to the I/O Domains, based on the amount of CPU and memory resources that you want to assign to each I/O Domain and the total amount of CPU and memory resources available in the CPU and memory repositories. Refer to the *I/O Domain Administration Guide* for more information.

Similarly, the IB VFs and 10GbE VFs owned by the I/O Domains come from the IB VF and 10GbE VF repositories (the IB VFs and 10GbE VFs released from Root Domains on the system), as shown in the following graphic.

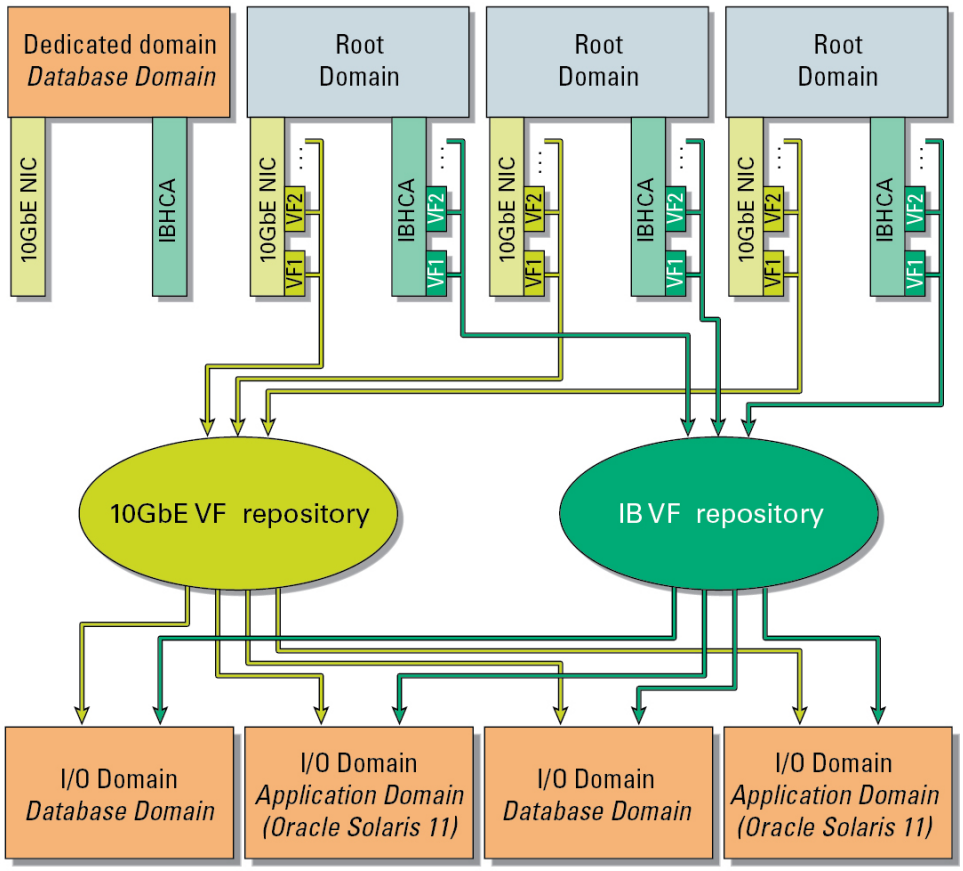

The number and size of the I/O Domains that you can create depends on several factors, including the amount of CPU and memory resources that are available in the CPU and memory repositories. However, while it is useful to know the total amount of resources that are parked in the repositories, it does not necessarily translate into the maximum number of I/O Domains that you can create for your system. In addition, you should not create an I/O Domain that uses more than one socket's worth of resources.

The SuperCluster Virtual Assistant keeps track of available and used resources, simplifying the administration of I/O Domains and resources.

Note - The following examples describe how resources can be divided up between domains using percentages to make the conceptual information easier to understand. However, you actually divide CPU and memory resources between domains at a socket granularity or core granularity level. Refer to the *Oracle SuperCluster M8 and SuperCluster M7 Administration Guide* for more information.

As an example configuration showing how you might assign CPU and memory resources to each domain, assume that you have a domain configuration where one of the domains is a Root Domain, and the other three domains are dedicated domains, as shown in the following figure.

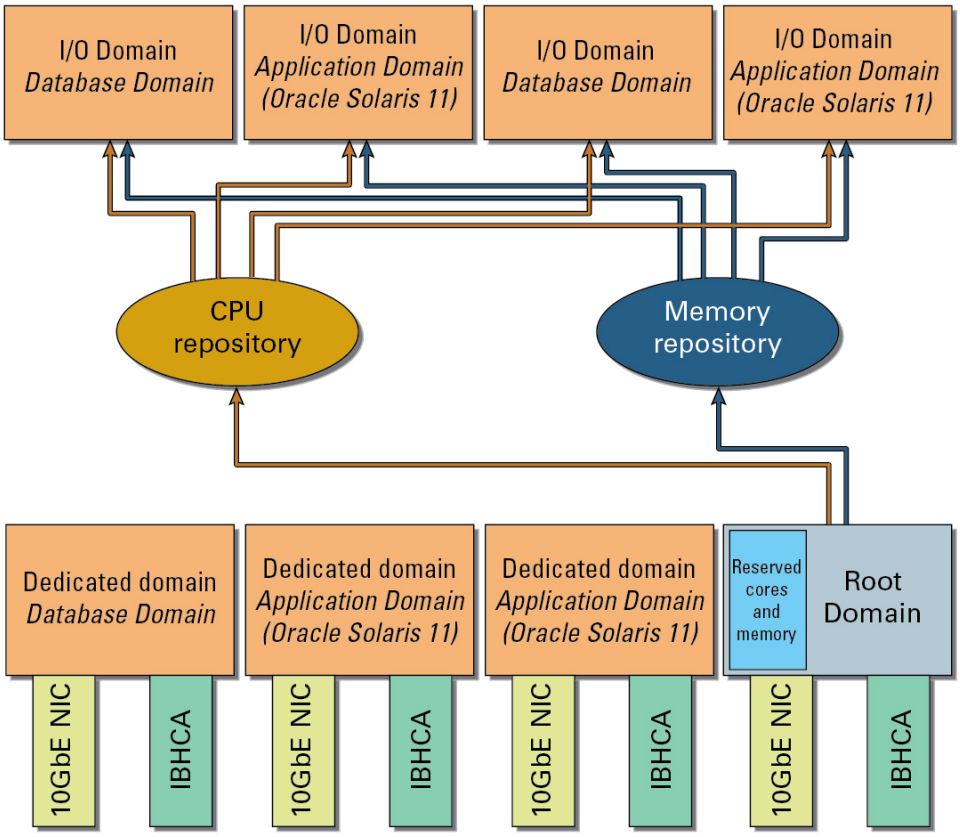

Even though dedicated domains and Root Domains are all shown as equal-sized domains in the preceding figure, that does not mean that CPU and memory resources must be split evenly across all four domains (where each domain is allocated 25% of the CPU and memory resources). Using information that you provide in the Configuration Worksheets, you can

request different sizes of CPU and memory resources for each domain when your SuperCluster M8 or SuperCluster M7 is initially installed.

For example, you can request that each dedicated domain have 30% of the CPU and memory resources (for a total of 90% of the CPU cores and memory resources allocated to the three dedicated domains), and the remaining 10% allocated to the single Root Domain. Having this configuration means that only 10% of the CPU core and memory resources are available for I/O Domains to pull from the CPU and memory repositories. However, you can also request that some of the resources from the dedicated domains be parked at the time of the initial installation of your system, which further increases the amount of CPU and memory resources available for I/O Domains to pull from the repositories.

You can also use the CPU/Memory tool after the initial installation to resize the amount of CPU and memory resources used by the existing domains, depending on the configuration that you chose at the time of your initial installation:

- If all of the domains on your compute server are dedicated domains, you can use the CPU/Memory tool to resize the amount of CPU and memory resources used by those domains.
- If you have a mixture of dedicated domains and Root Domains on your compute server:
  - For the dedicated domains, you can use the CPU/Memory tool to resize the amount of CPU and memory resources used by those dedicated domains. You can also use the tool to park some of the CPU and memory resources from the dedicated domains, which parks those resources in the CPU and Memory repositories, making them available for the I/O Domains.
  - For the Root Domains, you cannot resize the amount of CPU and memory resources for any of the Root Domains after the initial installation. Whatever resources that you asked to have assigned to the Root Domains at the time of initial installation are set and cannot be changed unless you have the Oracle installer come back out to your site to reconfigure your system.

Refer to the *Oracle SuperCluster M8 and SuperCluster M7 Administration Guide* for more information.

Assume you have a mixture of dedicated domains and Root Domains as mentioned earlier, where each dedicated domain has 30% of the CPU and memory resources (total of 90% resources allocated to dedicated domains), and the remaining 10% allocated to the Root Domain. You can then make the following changes to the resource allocation, depending on your situation:

- If you are satisfied with the amount of CPU and memory resources allocated to the Root Domain, but you find that one dedicated domain needs more resources while another needs less, you can reallocate the resources between the three dedicated domains (for example, having 40% for the first dedicated domain, 30% for the second, and 20% for the third), as long as the total amount of resources add up to the total amount available for all the dedicated domains (in this case, 90% of the resources).

- If you find that the amount of CPU and memory resources allocated to the Root Domain is insufficient, you can park resources from the dedicated domains, which parks those resources in the CPU and Memory repositories, making them available for I/O Domains. For example, if you find that you need 20% of the resources for I/O Domains created through the Root Domain, you can park 10% of the resources from one or more of the dedicated domains, which increases the amount of resources in the CPU and Memory repositories by that amount for the I/O Domains.

Related Information

- "Root Domains" on page 45
- "Dedicated Domains" on page 44

Domain Configurations for PDomains With One CMIOU

These topics describe the domain configurations available for PDomains with one CMIOU.

U1 – One-CMIOU Domain Configurations

This diagram shows the only available domain configuration for PDomains with one CMIOU.

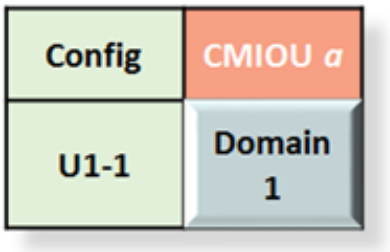

Also see “One CMIOU PCIe Card Configurations” on page 34.

From an overall PDomain level, the configuration with one CMIOU has the following characteristics:

- One processor, with 32 cores (8 hardware threads per core)
- 16 DIMM slots, providing:
  - SuperCluster M8 – A total of 1 TB (64 GB DIMMs) of memory, 960 GB available after DIMM sparing
  - SuperCluster M7 – A total of 512 GB (32 GB DIMMs) of memory, 480 GB available after DIMM sparing
- One IB HCA, one 10GbE NIC and one 1GbE NIC available for each PDomain

Related Information

- “U1-1 Domain Specifications” on page 55

U1-1 Domain Specifications

These tables provide information about the U1-1 domain configuration for the PDomains with one CMIOU.

TABLE 3 PCIe Slots and Cards, and CPU/Memory Resources (U1-1 Domain Configuration)

Item	Domain 1
1GbE NIC	PCIe slot 1
10GbE NIC	PCIe slot 2
IB HCA	PCIe slot 3
Empty (free) PCIe Slots	N/A
Default CPU Resources	100% (32 cores)
Default Memory Resources	100%: ■ 960 GB in SuperCluster M8 ■ 512 GB in SuperCluster M7

TABLE 4 Networks (U1-1 Domain Configuration)

		Domain 1
Management Network	Active	NET0, using P0 in 1GbE NIC

		Domain 1
	Standby	NET3, using P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC
	Standby	P1 in 10GbE NIC
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA
	Standby	P0 in IB HCA
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA
	Standby	P1 in IB HCA
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA
	Standby	P1 in IB HCA

Related Information

- “U1 – One-CMIOU Domain Configurations” on page 54

Domain Configurations for PDomains With Two CMIOUs

These topics describe the domain configurations available for PDomains with two CMIOUs.

- “U2 – Two-CMIOU Domain Configurations” on page 56
- “U2-1 Domain Specifications” on page 58
- “U2-2 Domain Specifications” on page 59

U2 – Two-CMIOU Domain Configurations

This diagram provides information on the available domain configurations for PDomains with two CMIOUs.

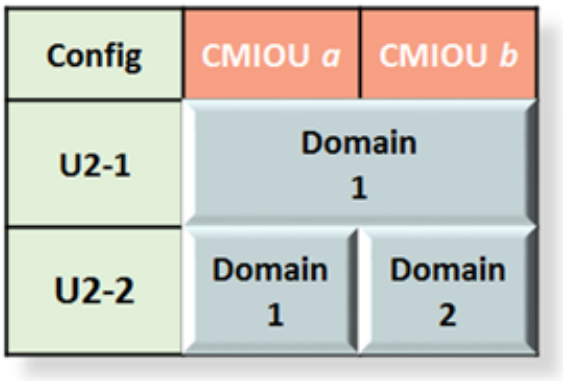

The specific CMIOU number varies, depending on the PDomain and SuperCluster model as shown in this table.

CMIOU No.	PDomain 0	PDomain 1
CMIOU *a*	CMIOU 0	SuperCluster M8 – CMIOU 4 SuperCluster M7 –CMIOU 5
CMIOU *b*	SuperCluster M8 – CMIOU 1 SuperCluster M7 –CMIOU 3	SuperCluster M8 – CMIOU 5 SuperCluster M7 – CMIOU 7

Also see "Two CMIOU PCIe Card Configurations (SuperCluster M8)" on page 35 and "Two CMIOUs PCIe Card Configurations (SuperCluster M7)" on page 37.

From an overall PDomain level, the configuration with two CMIOUs has the following characteristics:

- Two processors (one processor per CMIOU), each processor with 32 cores (8 hardware threads per core), for a total of 64 cores
- 32 DIMM slots (16 DIMM slots per CMIOU), providing:
  - SuperCluster M8 – A total of 2 TB (64 GB DIMMs) of total memory, 1.9 TB available after DIMM sparing
  - SuperCluster M7 – A total of 1 TB (32 GB DIMMs) of total memory, 960 GB available after DiMM sparing
- Two IB HCAs and two 10GbE NICs (one in each CMIOU) available for each PDomain

- One 1GbE NIC available for each PDomain, installed in the lowest-numbered CMIOU in that PDomain

How these resources are divided between domains within this PDomain depends on the type of domain configuration you choose.

Related Information

- "U2-1 Domain Specifications" on page 58
- "U2-2 Domain Specifications" on page 59

U2-1 Domain Specifications

These tables provide information on the U2-1 domain configuration for the PDomains with two CMIOUs.

TABLE 5 PCIe Slots and Cards, and CPU/Memory Resources (U2-1 Configuration)

Item	Domain 1
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in both CMIOUs in PDomain
IB HCAs	PCIe slot 3 in both CMIOUs in PDomain
Empty (free) PCIe Slots	■ PCIe slot 1 in CMIOU 1 or 5 in PDomain in SuperCluster M8 ■ PCIe slot 1 in CMIOU 3 or 7 in PDomain in SuperCluster M7
Default CPU Resources	100% (64 cores)
Default Memory Resources	100%: ■ 1.9 TB in SuperCluster M8 ■ 960 GB in SuperCluster M7

TABLE 6 Networks (U2-1 Configuration)

		Domain 1
Management Network	Active	NET0, using P0 in 1GbE NIC
	Standby	NET3, using P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain
	Standby	P1 in 10GbE NIC in second CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain
	Standby	P0 in IB HCA in second CMIOU in PDomain

		Domain 1
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in both CMIOUs in PDomain
	Standby	P1 in IB HCAs in both CMIOUs in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain
	Standby	P1 in IB HCA in second CMIOU in PDomain

Related Information

- "U2 – Two-CMIOU Domain Configurations" on page 56
- "U2-2 Domain Specifications" on page 59

U2-2 Domain Specifications

These tables provide information on the U2-2 PDomain configuration for the PDomains with two CMIOUs.

TABLE 7 PCIe Slots and Cards, and CPU/Memory Resources (U2-2 Configuration)

Item	Domain 1	Domain 2
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in first CMIOU in PDomain	PCIe slot 2 in second CMIOU in PDomain
IB HCAs	PCIe slot 3 in first CMIOU in PDomain	PCIe slot 3 in second CMIOU in PDomain
Empty (free) PCIe Slots	No free PCIe slots	■ PCIe slot 1 in CMIOU 1 or 5 in PDomain in SuperCluster M8 ■ PCIe slot 1 in CMIOU 3 or 7 in PDomain in SuperCluster M7
Default CPU Resources	50% (32 cores)	50% (32 cores)
Default Memory Resources	50%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	50%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7

TABLE 8 Networks (U2-2 Configuration)

		Domain 1	Domain 2
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC

		Domain 1	Domain 2
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in second CMIOU in PDomain
	Standby	P1 in 10GbE NIC in first CMIOU in PDomain	P1 in 10GbE NIC in second CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain
	Standby	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain
	Standby	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain
	Standby	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain

Related Information

- “U2 – Two-CMIOU Domain Configurations” on page 56
- “U2-1 Domain Specifications” on page 58

Domain Configurations for PDomains With Three CMIOUs

These topics describe the domain configurations available for PDomains with three CMIOUs.

- “U3 – Three-CMIOU Domain Configurations” on page 60
- “U3-1 Domain Specifications” on page 62
- “U3-2 Domain Specifications” on page 63
- “U3-3 Domain Specifications” on page 64

U3 – Three-CMIOU Domain Configurations

This diagram provides information on the available domain configurations for PDomains with three CMIOUs.

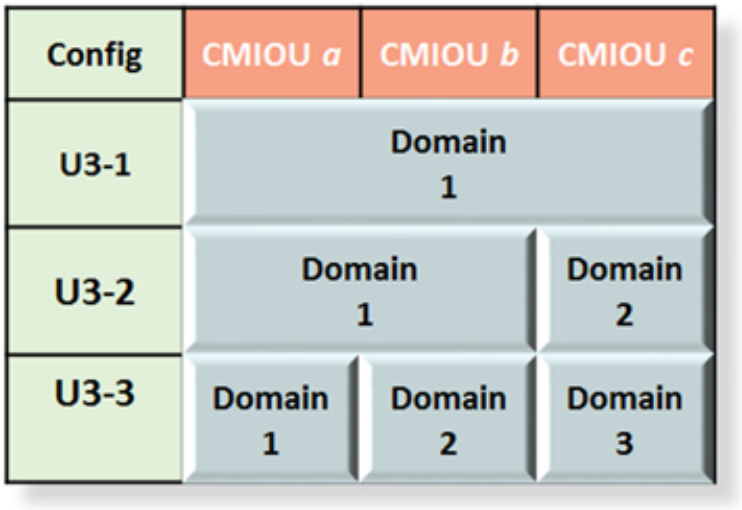

The specific CMIOU number varies, depending on which PDomain is being used as shown in this table.

CMIOU No.	PDomain 0	PDomain 1
CMIOU *a*	CMIOU 0	CMIOU 4
CMIOU *b*	CMIOU 1	CMIOU 5
CMIOU *c*	CMIOU 3	CMIOU 7

Also see "Three CMIOU PCIe Card Configurations" on page 38.

From an overall PDomain level, the configuration with three CMIOUs has the following characteristics:

- Three processors (one processor per CMIOU), each processor with 32 cores (8 hardware threads per core), for a total of 96 cores
- 48 DIMM slots (16 DIMM slots per CMIOU), providing:
  - SuperCluster M8 – A total of 3 TB (64 GB DIMMs) of total memory, 2.88 TB after DIMM sparing
  - SuperCluster M7 – A total of 1.5 TB (32 GB DIMMs) of total memory, 1.44 TB after DIMM sparing

- Three IB HCAs and three 10GbE NICs (one in each CMIOU) available for each PDomain
- One 1GbE NIC available for each PDomain, installed in the lowest-numbered CMIOU in that PDomain

How these resources are divided between domains within this PDomain depends on the type of domain configuration you choose.

Related Information

- "U3-1 Domain Specifications" on page 62
- "U3-2 Domain Specifications" on page 63
- "U3-3 Domain Specifications" on page 64

U3-1 Domain Specifications

These tables provide information on the U3-1 domain configuration for the PDomains with three CMIOUs.

TABLE 9 PCIe Slots and Cards, and CPU/Memory Resources (U3-1 Configuration)

Item	Domain 1
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in all CMIOUs in PDomain
IB HCAs	PCIe slot 3 in all CMIOUs in PDomain
Empty (free) PCIe Slots	■ PCIe slot 1 in CMIOU 1 and 3 in PDomain 0 ■ PCIe slot 1 in CMIOU 4 and 7 in PDomain 1
Default CPU Resources	100% (96 cores)
Default Memory Resources	100%: ■ 2..88 TB in SuperCluster M8 ■ 1.4 TB in SuperCluster M7

TABLE 10 Networks (U3-1 Configuration)

		Domain 1
Management Network	Active	NET0, using P0 in 1GbE NIC
	Standby	NET3, using P3 in 1GbE NIC

		Domain 1
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain
	Standby	P1 in 10GbE NIC in last CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain
	Standby	P0 in IB HCA in first CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in all CMIOUs in PDomain
	Standby	P1 in IB HCAs in all CMIOUs in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in second CMIOU in PDomain
	Standby	P1 in IB HCA in third CMIOU in PDomain

Related Information

- “U3 – Three-CMIOU Domain Configurations” on page 60
- “U3-2 Domain Specifications” on page 63
- “U3-3 Domain Specifications” on page 64

U3-2 Domain Specifications

These tables provide information on the U3-2 PDomain configuration for the PDomains with three CMIOUs.

TABLE 11 PCIe Slots and Cards, and CPU/Memory Resources (U3-2 Configuration)

Item	Domain 1	Domain 2
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in first and second CMIOU in PDomain	PCIe slot 2 in third CMIOU in PDomain
IB HCAs	PCIe slot 3 in first and second CMIOU in PDomain	PCIe slot 3 in third CMIOU in PDomain
Empty (free) PCIe Slots	PCIe slot 1 in CMIOU 1 or 4 in PDomain	PCIe slot 1 in CMIOU 3 or 7 in PDomain
Default CPU Resources	66% (64 cores)	33% (32 cores)
Default Memory Resources	66%: ■ 1.92 GB in SuperCluster M8 ■ 960 GB in SuperCluster M7	33%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7

TABLE 12 Networks (U3-2 Configuration)

		Domain 1	Domain 2
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in third CMIOU in PDomain
	Standby	P1 in 10GbE NIC in second CMIOU in PDomain	P1 in 10GbE NIC in third CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain
	Standby	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in first and second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
	Standby	P1 in IB HCAs in first and second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
	Standby	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain

Related Information

- “U3 – Three-CMIOU Domain Configurations” on page 60
- “U3-1 Domain Specifications” on page 62
- “U3-3 Domain Specifications” on page 64

U3-3 Domain Specifications

These tables provide information on the U3-3 PDomain configuration for the PDomains with three CMIOUs.

TABLE 13 PCIe Slots and Cards, and CPU/Memory Resources (U3-3 Configuration)

Item	Domain 1	Domain 2	Domain 3
1GbE NIC	■ PCIe slot 1 in CMIOU 0 in PDomain 0	■ Using VNET through 1GbE NIC in PCIe slot 1	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain

Item	Domain 1	Domain 2	Domain 3
	■ Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 5 in PDomain 1	in CMIOU 0 in PDomain 0 ■ PCIe slot 1 in CMIOU 5 in PDomain 1	
10GbE NICs	PCIe slot 2 in first CMIOU in PDomain	PCIe slot 2 in second CMIOU in PDomain	PCIe slot 2 in third CMIOU in PDomain
IB HCAs	PCIe slot 3 in first CMIOU in PDomain	PCIe slot 3 in second CMIOU in PDomain	PCIe slot 3 in third CMIOU in PDomain
Empty (free) PCIe slots	■ No free PCIe slots in PDomain 0 ■ PCIe slot 1 in CMIOU 4 in PDomain 1	■ PCIe slot 1 in CMIOU 1 in PDomain 0 ■ No free PCIe slots in PDomain 1	PCIe slot 1 in CMIOU 3 or 7 in PDomain
Default CPU Resources	33% (32 cores)	33% (32 cores)	33% (32 cores)
Default Memory Resources	33%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	33%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	33%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7

TABLE 14 Networks (U3-3 Configuration)

		Domain 1	Domain 2	Domain 3
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in second CMIOU in PDomain	P0 in 10GbE NIC in third CMIOU in PDomain
	Standby	P1 in 10GbE NIC in first CMIOU in PDomain	P1 in 10GbE NIC in second CMIOU in PDomain	P1 in 10GbE NIC in third CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain
	Standby	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
	Standby	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
	Standby	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain

Related Information

- “U3 – Three-CMIOU Domain Configurations” on page 60
- “U3-1 Domain Specifications” on page 62
- “U3-2 Domain Specifications” on page 63

Domain Configurations for PDomains With Four CMIOUs

These topics describe the domain configurations available for PDomains with four CMIOUs.

- “U4 – Four-CMIOU Domain Configurations” on page 66
- “U4-1 Domain Specifications” on page 68
- “U4-2 Domain Specifications” on page 69
- “U4-3 Domain Specifications” on page 71
- “U4-4 Domain Specifications” on page 72

U4 – Four-CMIOU Domain Configurations

This diagram provides information on the available domain configurations for PDomains with four CMIOUs.

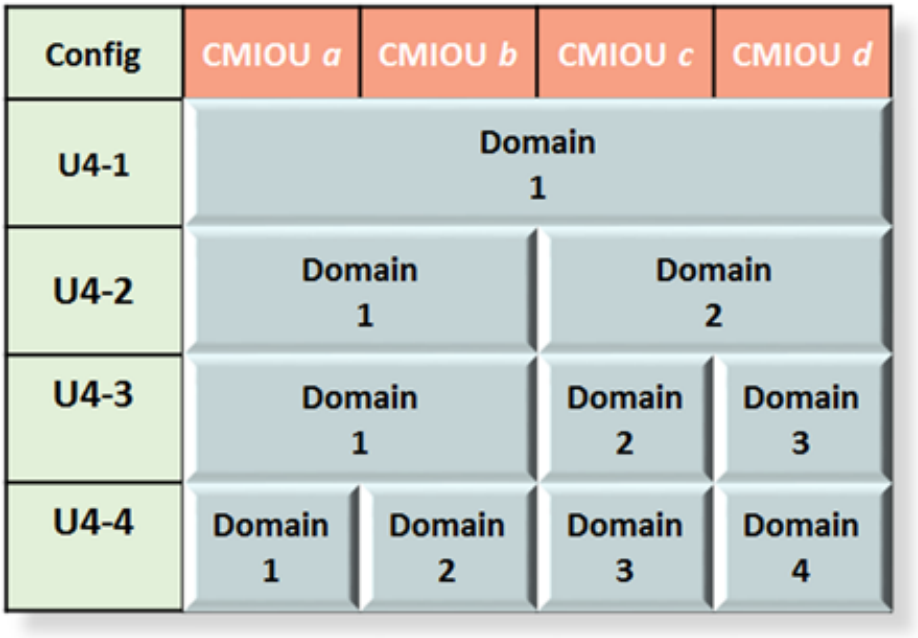

The specific CMIOU number varies, depending on which PDomain is being used as shown in this table.

CMIOU No.	PDomain 0	PDomain 1
CMIOU *a*	CMIOU 0	CMIOU 4
CMIOU *b*	CMIOU 1	CMIOU 5
CMIOU *c*	CMIOU 2	CMIOU 6
CMIOU *d*	CMIOU 3	CMIOU 7

Also see “Four CMIOUs PCIe Card Configurations” on page 40.

From an overall PDomain level, the configuration with four CMIOUs has the following characteristics:

- Four processors (one processor per CMIOU), each processor with 32 cores (8 hardware threads per core), for a total of 128 cores
- 64 DIMM slots (16 DIMM slots per CMIOU), providing:

  - SuperCluster M8 – A total of 4 TB (64 GB DIMMs) of total memory, 3.84 TB after DIMM sparing
  - SuperCluster M8 – A total of 2 TB (32 GB DIMMs) of total memory, 1.92 TB after DIMM sparing
- Four IB HCAs and four 10GbE NICs (one in each CMIOU) available for each PDomain
- One 1GbE NIC available for each PDomain, installed in the lowest-numbered CMIOU in that PDomain

How these resources are divided between domains within this PDomain depends on the type of domain configuration you choose.

Related Information

- "U4-1 Domain Specifications" on page 68
- "U4-2 Domain Specifications" on page 69
- "U4-3 Domain Specifications" on page 71
- "U4-4 Domain Specifications" on page 72

U4-1 Domain Specifications

These tables provide information on the U4-1 domain configuration for the PDomains with four CMIOUs.

TABLE 15 PCIe Slots and Cards, and CPU/Memory Resources (U4-1 Configuration)

Item	Domain 1
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in all CMIOUs in PDomain
IB HCAs	PCIe slot 3 in all CMIOUs in PDomain
Empty (free) PCIe Slots	■ PCIe slot 1 in CMIOU 1, 2, and 3 in PDomain 0 ■ PCIe slot 1 in CMIOU 4, 6, and 7 in PDomain 1
Default CPU Resources	100% (128 cores)
Default Memory Resources	100%: ■ 3.84 TB in SuperCluster M8 ■ 1.92 TB in SuperCluster M7

TABLE 16 Networks (U4-1 Configuration)

		Domain 1
Management Network	Active	NET0, using P0 in 1GbE NIC
	Standby	NET3, using P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain
	Standby	P1 in 10GbE NIC in last CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain
	Standby	P0 in IB HCA in first CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in all CMIOUs in PDomain
	Standby	P1 in IB HCAs in all CMIOUs in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in second CMIOU in PDomain
	Standby	P1 in IB HCA in third CMIOU in PDomain

Related Information

- “U4 – Four-CMIOU Domain Configurations” on page 66
- “U4-2 Domain Specifications” on page 69
- “U4-3 Domain Specifications” on page 71
- “U4-4 Domain Specifications” on page 72

U4-2 Domain Specifications

These tables provide information on the U4-2 PDomain configuration for the PDomains with four CMIOUs.

TABLE 17 PCIe Slots and Cards, and CPU/Memory Resources (U4-2 Configuration)

Item	Domain 1	Domain 2
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in first and second CMIOU in PDomain	PCIe slot 2 in third and fourth CMIOU in PDomain
IB HCAs	PCIe slot 3 in first and second CMIOU in PDomain	PCIe slot 3 in third and fourth CMIOU in PDomain

Item	Domain 1	Domain 2
Empty (free) PCIe Slots	PCIe slot 1 in CMIOU 1 or 4 in PDomain	■ PCIe slot 1 in CMIOU 2 and 3 in PDomain 0 ■ PCIe slot 1 in CMIOU 6 and 7 in PDomain 1
Default CPU Resources	50% (64 cores)	50% (64 cores)
Default Memory Resources	50%: ■ 1.92 TB in SuperCluster M8 ■ 960 GB in SuperCluster M7	50%: ■ 1.92 TB in SuperCluster M8 ■ 960 GB in SuperCluster M7

TABLE 18 Networks (U4-2 Configuration)

		Domain 1	Domain 2
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in third CMIOU in PDomain
	Standby	P1 in 10GbE NIC in second CMIOU in PDomain	P1 in 10GbE NIC in fourth CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain
	Standby	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA in first and second CMIOU in PDomain	P0 in IB HCA in third and fourth CMIOU in PDomain
	Standby	P1 in IB HCA in first and second CMIOU in PDomain	P1 in IB HCA in third and fourth CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain
	Standby	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain

Related Information

- "U4 – Four-CMIOU Domain Configurations" on page 66
- "U4-1 Domain Specifications" on page 68
- "U4-3 Domain Specifications" on page 71
- "U4-4 Domain Specifications" on page 72

U4-3 Domain Specifications

These tables provide information on the U4-3 PDomain configuration for the PDomains with four CMIOUs.

TABLE 19 PCIe Slots and Cards, and CPU/Memory Resources (U4-3 Configuration)

Item	Domain 1	Domain 2	Domain 3
1GbE NIC	PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in first and second CMIOU in PDomain	PCIe slot 2 in third CMIOU in PDomain	PCIe slot 2 in fourth CMIOU in PDomain
IB HCAs	PCIe slot 3 in first and second CMIOU in PDomain	PCIe slot 3 in third CMIOU in PDomain	PCIe slot 3 in fourth CMIOU in PDomain
Empty (free) PCIe slots	PCIe slot 1 in CMIOU 1 or 4 in PDomain	PCIe slot 1 in CMIOU 2 or 6 in PDomain	PCIe slot 1 in CMIOU 3 or 7 in PDomain
Default CPU Resources	50% (64 cores)	25% (32 cores)	25% (32 cores)
Default Memory Resources	50%: ■ 1.92 TB in SuperCluster M8 ■ 960 GB in SuperCluster M7	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7

TABLE 20 Networks (U4-3 Configuration)

		Domain 1	Domain 2	Domain 3
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in third CMIOU in PDomain	P0 in 10GbE NIC in fourth CMIOU in PDomain
	Standby	P1 in 10GbE NIC in second CMIOU in PDomain	P1 in 10GbE NIC in third CMIOU in PDomain	P1 in 10GbE NIC in fourth CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain
	Standby	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in first and second CMIOUs in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain

		Domain 1	Domain 2	Domain 3
	Standby	P1 in IB HCAs in first and second CMIOUs in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
	Standby	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain

Related Information

- “U4 – Four-CMIOU Domain Configurations” on page 66
- “U4-1 Domain Specifications” on page 68
- “U4-2 Domain Specifications” on page 69
- “U4-4 Domain Specifications” on page 72

U4-4 Domain Specifications

These tables provide information on the U4-4 PDomain configuration for the PDomains with four CMIOUs.

TABLE 21 PCIe Slots and Cards, and CPU/Memory Resources (U4-4 Configuration)

Item	Domain 1	Domain 2	Domain 3	Domain 4
1GbE NIC	■ PCIe slot 1 in CMIOU 0 in PDomain 0 ■ Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 5 in PDomain 1	■ Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 in PDomain 0 ■ PCIe slot 1 in CMIOU 5 in PDomain 1	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain	Using VNET through 1GbE NIC in PCIe slot 1 in CMIOU 0 or 5 in PDomain
10GbE NICs	PCIe slot 2 in first CMIOU in PDomain	PCIe slot 2 in second CMIOU in PDomain	PCIe slot 2 in third CMIOU in PDomain	PCIe slot 2 in fourth CMIOU in PDomain
IB HCAs	PCIe slot 3 in first CMIOU in PDomain	PCIe slot 3 in second CMIOU in PDomain	PCIe slot 3 in third CMIOU in PDomain	PCIe slot 3 in fourth CMIOU in PDomain
Empty (free) PCIe slots	■ No free PCIe slots in PDomain 0 ■ PCIe slot 1 in CMIOU 4 in PDomain 1	■ PCIe slot 1 in CMIOU 1 in PDomain 0 ■ No free PCIe slots in PDomain 1	PCIe slot 1 in CMIOU 2 or 6 in PDomain	PCIe slot 1 in CMIOU 3 or 7 in PDomain

Item	Domain 1	Domain 2	Domain 3	Domain 4
Default CPU Resources	25% (32 cores)	25% (32 cores)	25% (32 cores)	25% (32 cores)
Default Memory Resources	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7	25%: ■ 960 GB in SuperCluster M8 ■ 480 GB in SuperCluster M7

TABLE 22 Networks (U4-4 Configuration)

		Domain 1	Domain 2	Domain 3	Domain 4
Management Network	Active	NET0, using P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC	NET0, using VNET through P0 in 1GbE NIC	NET0, using VNET through P2 in 1GbE NIC
	Standby	NET1, using P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC	NET1, using VNET through P1 in 1GbE NIC	NET1, using VNET through P3 in 1GbE NIC
10GbE Client Access Network	Active	P0 in 10GbE NIC in first CMIOU in PDomain	P0 in 10GbE NIC in second CMIOU in PDomain	P0 in 10GbE NIC in third CMIOU in PDomain	P0 in 10GbE NIC in fourth CMIOU in PDomain
	Standby	P1 in 10GbE NIC in first CMIOU in PDomain	P1 in 10GbE NIC in second CMIOU in PDomain	P1 in 10GbE NIC in third CMIOU in PDomain	P1 in 10GbE NIC in fourth CMIOU in PDomain
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain
	Standby	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in first CMIOU in PDomain	P0 in IB HCAs in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
	Standby	P1 in IB HCAs in first CMIOU in PDomain	P1 in IB HCAs in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in first CMIOU in PDomain	P0 in IB HCA in second CMIOU in PDomain	P0 in IB HCA in third CMIOU in PDomain	P0 in IB HCA in fourth CMIOU in PDomain
	Standby	P1 in IB HCA in first CMIOU in PDomain	P1 in IB HCA in second CMIOU in PDomain	P1 in IB HCA in third CMIOU in PDomain	P1 in IB HCA in fourth CMIOU in PDomain

Related Information

- "U4 – Four-CMIOU Domain Configurations" on page 66

Understanding CPU and Memory Resources

These topics provide information about SuperCluster CPU and memory resources:

Description	Links
Find out what kinds of resource allocations can be made and when such activities can be performed.	"Supported CPU and Memory Reallocations" on page 75
Learn about SuperCluster CPU resources.	"CPU Resource Overview" on page 76 "CPU Cores Available for Database Zones" on page 77 "CPU Cores Available for I/O Domains" on page 77
Learn about SuperCluster memory resources.	"Memory Resources Overview" on page 78 "Memory Resources Available for Database Zones" on page 79 "Memory Resources Available for I/O Domains" on page 80
Prior to the SuperCluster installation, find instructions for completing resource specifications in Configuration Worksheets.	"Specifying CPU and Memory Resources in Configuration Worksheets" on page 81

Supported CPU and Memory Reallocations

SuperCluster compute server CPU and memory resources are initially allocated during the installation. However, the amount of resources that are allocated to domains can be modified by Oracle personnel during the installation, or you can change the resources later using utilities provided in the SuperCluster software.

If you want to specify non-default resource allocations for the installation, enter your specifications in the Configuration Worksheets prior to the installation. See "Specifying CPU and Memory Resources in Configuration Worksheets" on page 81.

This table lists the supported resource allocation activities.

Domain Configuration	Supported Allocation Activities	Utility Used	For More Information
All domains are dedicated domains	Reallocate all of the resources across domains at the socket or core level (a reboot might be required if primary domain resources are changed).	`osc-setcoremem`	For `osc-setcoremem` instructions refer to:Configuring CPU and Memory Resources (osc-setcoremem) in the *Oracle SuperCluster M8 and SuperCluster M7 Administration Guide*
	Remove (park) resources from dedicated domains for licensing purposes.	`osc-setcoremem`	
Mixed domains – some are dedicated, some are Root Domains	**Before any I/O Domains are Created:** Reallocate all of the resources across domains at the socket or core level (a reboot is required if primary domain resources are changed).	`osc-setcoremem`	
	Anytime: Remove resources from dedicated domains so that the resources are available to I/O Domains.	`osc-setcoremem`	
	Move resources between dedicated domains.	`osc-setcoremem`	
	Configure resources for I/O Domains.	SuperCluster Virtual Assistant	*Oracle I/O Domain Administration Guide*

CPU and memory resources that have been released to CPU and memory repositories in support of I/O domains, cannot be unparked for use in dedicated domains

CPU Resource Overview

Every CMIOU has 32 cores of CPU resources. The amount of CPU resources available for each domain varies, depending on the following factors:

- The number of PDomains in the system
- The number of CMIOUs in each PDomain
- The domain configuration on each PDomain

To determine the CPU resources for a particular SuperCluster configuration, refer to the *Domain Specifications* in these sections:

- "Domain Configurations for PDomains With One CMIOU" on page 54
- "Domain Configurations for PDomains With Two CMIOUs" on page 56
- "Domain Configurations for PDomains With Three CMIOUs" on page 60
- "Domain Configurations for PDomains With Four CMIOUs" on page 66

CPU Cores Available for Database Zones

Note - Database zones can only be created on dedicated Database Domains. The information in this section applies to database zones and Database Domains that are dedicated domains.

When the OS is installed on a domain, that domain is automatically designated as the global zone. When creating zones on Database Domains, the Database Domain is designated as the global zone, and the zones created on that Database Domain are designated as nonglobal zones.

A certain number of cores are always set aside for the global zone (the Database Domain). The remaining cores in the Database Domain are available for the nonglobal zones (the zones in the Database Domain). The number of cores that are set aside for the global zone varies, depending on the number of CMIOUs that are associated with the domain:

- One CMIOU associated with a domain – 2 cores are reserved for the global zone, and the remaining cores are available for the nonglobal zones.
- Two or more CMIOUs associated with a domain – 4 cores are reserved for the global zone, and the remaining cores are available for the nonglobal zones.
- A Database Domain with no zones – All the cores are available for that Database Domain.

For each zone that you create, use a minimum of one core per zone. However, depending on the workload that you expect on a zone, a larger number of cores per zone might be preferable, thereby reducing the total number of zones on each compute server. Carefully consider the expected workload on each zone that you create, so that you allot the appropriate number of cores to those zones.

CPU Cores Available for I/O Domains

If you want I/O Domains set up on your Oracle SuperCluster, either at the time of the initial installation or afterwards, you must have at least one Root Domain set up at the time of the initial installation. I/O Domains can then be created from the Root Domains.

A certain number of CPU cores are always reserved for each Root Domain, depending on which domain is being used as a Root Domain and the number of IB HCAs and 10GbE NICs that are associated with that Root Domain:

- The last domain in a domain configuration:
  - Two cores reserved for a Root Domain with one IB HCA and 10GbE NIC
  - Four cores reserved for a Root Domain with two IB HCAs and 10GbE NICs
- Any other domain in a domain configuration:
  - One core reserved for a Root Domain with one IB HCA and 10GbE NIC
  - Two cores reserved for a Root Domain with two IB HCAs and 10GbE NICs

All remaining Root Domain CPU cores are parked in the CPU repository, and can be allocated to I/O Domains.

The CPU repository contains resources not only from the Root Domains, but also any parked resources from the dedicated domains. Whether the CPU core resources originated from dedicated domains or from Root Domains, once those resources have been parked in the CPU repository, those resources are no longer associated with their originating domain. These resources can also be allocated to I/O Domains.

In addition, the CPU repository contains parked resources only from the compute server that contains the domains providing those parked resources. In other words, if you have two compute servers and both compute servers have Root Domains, there are two sets of CPU repositories, where each compute server has its own CPU repository with parked resources.

For more information about parked resources, see "Root Domains" on page 45 and "I/O Domains" on page 49.

For more information about planning I/O Domain resource allocations, refer to the *Oracle I/O Domain Administration Guide*.

Memory Resources Overview

SuperCluster M8 and SuperCluster M7 employ DIMM sparing technology. DIMM sparing provides a mechanism to automatically unconfigure a failed DIMM without downtime. DIMM sparing reserves the space of one DIMM for use in the event of a failed DIMM. Due to the requirements of DIMM sparing, the amount of available memory reported is slightly less than the amount of memory that is calculated based on the quantity and capacity of the DIMMs.

Every CMIOU has 16 memory slots.

- • SuperCluster M8 – Each CMIOU is fully populated with 64 GB DIMMs, for a total of 1 TB (1024 GB) of memory in each CMIOU, with 960 GB available after DIMM sparing.
- • SuperCluster M7 – Each CMIOU is fully populated with 32 GB DIMMs, for a total of 512 GB of memory in each CMIOU, with 480 available after DIMM sparing.

The amount of memory resources that you have available for domains depends on these factors:

- The number of PDomains in the system
- The number of CMIOUs in each PDomain
- The domain configuration on each PDomain

To determine the memory resources for a particular SuperCluster configuration, refer to the applicable *Domain Specifications* in these sections:

- "Domain Configurations for PDomains With One CMIOU" on page 54
- "Domain Configurations for PDomains With Two CMIOUs" on page 56
- "Domain Configurations for PDomains With Three CMIOUs" on page 60
- "Domain Configurations for PDomains With Four CMIOUs" on page 66

Memory Resources Available for Database Zones

Note – Database zones can only be created on Database Domains that are dedicated domains. The information in this section applies to database zones and Database Domains that are dedicated domains.

The amount of memory resources available for database zones depends on the amount of memory resources that you have assigned to the Database Domain, and then how you want to divide those memory resources up for the database zones within that Database Domain.

For example, assume you have a Database Domain that has two CMIOUs. By default, 960 GB of memory is available to that Database Domain. You can therefore have four equal-sized database zones within that Database Domain, where each database zone has 200 GB of memory assigned to it, for a total of 800 GB of memory for all database zones. The remaining 160 GB of memory in this Database Domain can be saved for future database zones that you might want to create on this Database Domain.

Memory Resources Available for I/O Domains

If you want I/O Domains set up on your Oracle SuperCluster, either at the time of the initial installation or afterwards, you must have at least one Root Domain set up at the time of the initial installation. I/O Domains can then be created from these Root Domains.

A certain amount of memory resources are always reserved for each Root Domain, depending on which domain is being used as a Root Domain in the domain configuration and the number of IB HCAs and 10GbE NICs that are associated with that Root Domain:

- • The last domain in a domain configuration:
  - • 32 GB of memory reserved for a Root Domain with one IB HCA and 10GbE NIC
  - • 64 GB of memory reserved for a Root Domain with two IB HCAs and 10GbE NICs
- • Any other domain in a domain configuration:
  - • 16 GB of memory reserved for a Root Domain with one IB HCA and 10GbE NIC
  - • 32 GB of memory reserved for a Root Domain with two IB HCAs and 10GbE NICs

The remaining memory resources allocated with each Root Domain are parked in the memory repository, which can then be used by I/O Domains.

Note – For more information on the number of IB HCAs and 10GbE NICs associated with each domain, refer to the Oracle SuperCluster M8 and SuperCluster M7 Overview Guide.

The memory repository contains resources not only from the Root Domains, but also any parked resources from the dedicated domains. Whether memory resources originated from dedicated domains or from Root Domains, once those resources have been parked in the memory repository, those resources are no longer associated with their originating domain. These resources become equally available to I/O Domains.

In addition, the memory repository contains parked resources only from the compute server that contains the domains providing those parked resources. In other words, if you have two compute servers and both compute servers have Root Domains, there are two sets of memory repositories, where each compute server has its own memory repository with parked resources.

For example, assume you have four domains on your compute server, with three of the four domains as Root Domains. Assume each domain has the following:

- • One IB HCA and one 10GbE NIC
- • 480 GB of memory

In this situation, the following memory resources are reserved for each Root Domain, with the remaining resources available for the memory repository:

- • 32 GB of memory reserved for the last Root Domain in this configuration. 448 GB of memory available from this Root Domain for the memory repository.
- • 16 GB of memory reserved for the second and third Root Domains in this configuration.
  - • 464 GB of memory available from each of these Root Domains for the memory repository.
  - • A total of 928 GB of memory (464 x 2) available for the memory repository from these two Root Domains.

A total of 1376 GB of memory (448 + 928) are therefore parked in the memory repository and are available for the I/O Domains.

Specifying CPU and Memory Resources in Configuration Worksheets

Prior to the delivery of your SuperCluster M8 or SuperCluster M7, you are asked to provide site specific SuperCluster configuration information.

These tasks are intended to assist in competing the Configuration Worksheets for CPU and memory resource allocations for the domains that will be configured:

- “Specify CPU Resources in Configuration Worksheets (Prior to the Installation)” on page 81
- “Specify Memory Resources in Configuration Worksheets (Prior to the Installation)” on page 83

▼ Specify CPU Resources in Configuration Worksheets (Prior to the Installation)

Use this task when completing your Configuration Worksheets prior to the SuperCluster installation.

Every CMIOU has 32 cores of CPU resources. The amount of CPU resources available for each domain varies, depending on the following factors:

- The number of CMIOUs in each server
- The domain configuration on each PDomain

1. **Determine the default CPU resources for the domains in your SuperCluster configuration.**

   Refer to the applicable *Domain Specifications* in these sections:

   - "Domain Configurations for PDomains With One CMIOU" on page 54
   - "Domain Configurations for PDomains With Two CMIOUs" on page 56
   - "Domain Configurations for PDomains With Three CMIOUs" on page 60
   - "Domain Configurations for PDomains With Four CMIOUs" on page 66

2. **If the default CPU allocations are not ideal, determine and specify how many CPU resources should be assigned to each domain.**

   Enter the number of CPU cores for each domain in the Configuration Worksheet.

   Example

   For a U2-2 domain configuration (two domains in a PDomain). by default, 50% of the CPU resources are assigned to each domain. Alternatively, you can specify different values for each domain, such as 60% of the CPU resources to the first domain and 40% to the second domain.

   In this example, the total CPU resources for U2-2 are 64 cores. A 60% / 40% configuration allocates these CPU resources (values are rounded):

   - Domain 1 – 38 cores
   - Domain 2 – 26 cores

3. **If any Database Domains will include zones, enter that information into the Configuration Worksheets.**

   Enter the CPU resources for each zone in the Configuration Worksheet.

   Example

   The first domain is a Database Domain that will contain four equal-sized zones.

   Domain 1 has 38 cores. Subtract 2 cores (minimum required) for the Database Domain itself. That leaves 36 cores available for zones. 36 / 4 = 9 cores per zone.

4. **For Root Domains, specify the CPU resources for the Root Domain and I/O Domains.**

   Enter the CPU resources for each I/O Domain in the Configuration Worksheet.

   Example:

   The second domain is a Root Domain that will support four equal-sized I/O Domains.

   Domain 2 has 26 cores. Subtract 2 cores (minimum required) for the Root Domain itself. That leaves 24 cores for I/O Domains. 24 / 4 = 6 cores per I/O Domain.

Related Information

- “CPU Cores Available for Database Zones” on page 77
- “CPU Cores Available for I/O Domains” on page 77

▼ Specify Memory Resources in Configuration Worksheets (Prior to the Installation)

Use this task when completing your Configuration Worksheets prior to the SuperCluster installation.

Every CMIOU has 16 memory slots.

- • SuperCluster M8 – Each CMIOU is fully populated with 64 GB DIMMs, for a total of 1 TB (1024 GB) of memory in each CMIOU, with 960 GB available after DIMM sparing.
- • SuperCluster M7 – Each CMIOU is fully populated with 32 GB DIMMs, for a total of 512 GB of memory in each CMIOU, with 480 available after DIMM sparing.

1. **Determine the default memory resources for the domains in your SuperCluster configuration.**

   Refer to the applicable *Domain Specifications* in these sections:

   - “Domain Configurations for PDomains With One CMIOU” on page 54
   - “Domain Configurations for PDomains With Two CMIOUs” on page 56
   - “Domain Configurations for PDomains With Three CMIOUs” on page 60
   - “Domain Configurations for PDomains With Four CMIOUs” on page 66

2. **If the default memory allocations are not ideal, determine and specify the memory resources that should be assigned to each domain.**

   Enter the amount of memory for each domain in the Configuration Worksheet.

   Example

   For a U2-2 domain configuration (two domains in a PDomain). by default, 50% of the memory resources are assigned to each domain. Alternatively, you can specify different values for each domain, such as 60% of the memory resources to the first domain and 40% to the second domain.

   In this example, the total memory resources for SuperCluster M8 with a U2-2 configuration is 1.92 TB. A 60% / 40% configuration allocates these CPU resources (values are rounded):

   - Domain 1 – 1.15 TB

- Domain 2 – 768 GB

3. **If any Database Domains will include zones, enter memory resource information into the Configuration Worksheets.**

   Enter the memory resources for each zone in the Configuration Worksheet.

   Example

   The first domain is a Database Domain that will contain four equal-sized zones.

   Domain 1 has 1.15 TB of memory. 1.15 / 4 = 480 GB per zone.

4. **For Root Domains, specify the memory resources for I/O Domains.**

   Enter the memory resources for each I/O Domain in the Configuration Worksheet.

   Example:

   The second domain is a Root Domain that will support four equal-sized I/O Domains.

   Domain 2 has 768 GB of memory. Subtract 32 GB (minimum required) for the Root Domain itself. That leaves 736 GB for I/O Domains. 736 / 4 = 184 GB of memory per I/O Domain.

Related Information

- “Memory Resources Available for Database Zones ” on page 79
- “Memory Resources Available for I/O Domains” on page 80

Understanding SuperCluster Networks

These topics describe the network requirements for SuperCluster M8 and SuperCluster M7.

SuperCluster Networks Overview

Before deploying SuperCluster in your environment, review the networking information in this section, and ensure that your site provides the minimum network requirements listed in "Network Connectivity Requirements" on page 89.

Networking Interfaces

SuperCluster M8 and SuperCluster M7 include compute servers, storage servers, and the ZFS storage appliance, as well as equipment to connect the compute servers to your network. The network connections enable the servers to be administered remotely and enable clients to connect to the compute servers.

Each compute server consists of the following network components and interfaces:

- Several 10GbE NICs for connection to the 10GbE client access network:
  - SuperCluster M8 – Oracle Quad Port 10GbE NIC (the Oracle Quad 10 Gb or Dual 40 Gb Ethernet Adapter is used in the 2x2x10GbE mode, where each port is split into two physical functions that operate at 10Gbps.)
  - SuperCluster M7 – Sun Dual Port 10GbE SFP+ NIC
- Four 1GbE ports (NET 0, NET 1, NET 2, and NET 3) provided by 1GbE NICs for connections to the host management network
- One Ethernet port (NET MGT) for Oracle ILOM remote management
- Several dual-ported IB HCAs for connection to the IB private network

Each storage server consists of the following network components and interfaces:

- One embedded Gigabit Ethernet port (NET 0) for connection to the host management network
- One dual-ported Sun QDR IB PCIe Low Profile HCA for connection to the IB private network
- One Ethernet port (NET MGT) for Oracle ILOM remote management

Each storage controller consists of the following network components and interfaces:

- One embedded Gigabit Ethernet port for connection to the host management network:
  - NET 0 on the first storage controller (installed in slot 25 in the rack)
  - NET 1 on the second storage controller (installed in slot 26 in the rack)
- One dual-port QDR IB HCA for connection to the IB private network
- One Ethernet port (NET 0) for Oracle ILOM remote management using sideband management. The dedicate Oracle ILOM port is not used due to sideband.

The Ethernet management switch supplied with SuperCluster M8 and SuperCluster M7 is minimally configured during installation. The minimal configuration disables IP routing, and sets these parameters:

- 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, your network administrator must perform the necessary configuration steps during installation of SuperCluster.

There are three networks for SuperCluster. Each network must be on a distinct and separate subnet from the others.

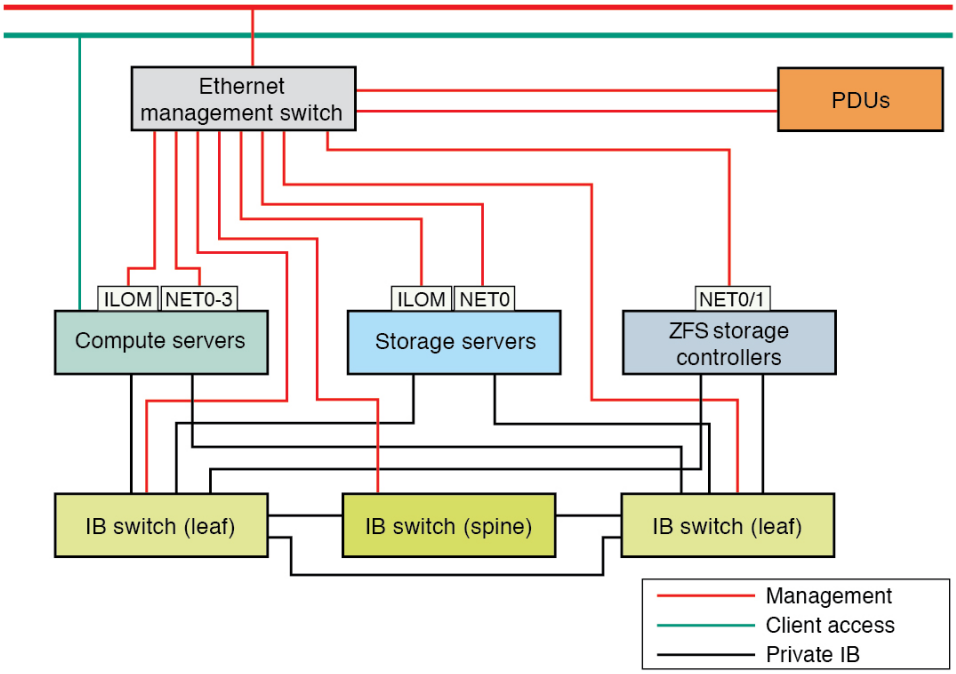

The network descriptions are as follows:

- **Client access network —** This required 10GbE network connects the compute 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.
- **Management network** — This required network connects to your existing management network, and is used for administrative work for all SuperCluster components. This network

connects the servers and IB switches to the Ethernet switch. There is one uplink from the Ethernet management switch to your existing management network.

Note - Management network connectivity to the PDUs is optional, and only needed if the PDU power will be monitored remotely.

Each compute server and storage server use two network interfaces for management. One provides management access to the operating system through the 1GbE host management interfaces, and the other provides access to the Oracle Integrated Lights Out Manager through the NET MGT Ethernet interfaces.

The method used to connect the storage controllers to the management network varies depending on the controller:

- **Storage controller 1** — NET 0 is used to provide access to the Oracle ILOM network using sideband management, as well as access to the 1GbE host management network.
- **Storage controller 2** — NET 0 is used to provide access to the Oracle ILOM network using sideband management, and NET1 is used to provide access to the 1GbE host management network.

The 1GbE host management interfaces on the compute servers should not be used for client or application network traffic. Cabling or configuration changes to these interfaces is not permitted.

- **IB private network** — This network connects the compute servers, ZFS storage appliance, and storage servers using the IB switches. This nonroutable network is fully contained in SuperCluster, and does not connect to your existing network. This network is automatically configured during installation.

  For systems with Database Domains, Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on storage servers and the ZFS storage appliance.

  For systems with Application Domains, Oracle Solaris Cluster uses this network for cluster interconnect traffic and to access data on the ZFS storage appliance.

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

Supported IPMP Network Features

SuperCluster network interfaces use link-based IP network multipathing (IPMP) for the IB switches, the 10GbE Client Access network, and the GbE management network. Datalink multipathing (DLMP) and probe-based IPMP (including transitive IPMP) are not supported.

IPMP ports are always configured active-standby, with the exception of IB HCA ports in the global zone of Database Dedicated Domains, where ports are configured active-active for the database pkey 0xFFFF partition (and therefore consume two IP addresses per HCA rather than one). All other IB ports are configured active-standby, including for non-global Database Zones in Database Dedicated Domains, and for both global zones and non-global zones in Database I/O domains.

VLAN tagging and trunking are supported on SuperCluster both in domains and zones.

Aggregation (LACP) is supported in Dedicated Domains (some manual configuration is required), but it is not supported in I/O Domains.

Related Information

- “Network Connectivity Requirements” on page 89
- “Default IP Addresses” on page 94
- “Understanding Default Host Names and IP Addresses (Single-Server Version)” on page 95
- “Understanding Default Host Names and IP Addresses (Dual-Server Version)” on page 99

Network Connectivity Requirements

These are the mandatory network connectivity requirements that must be provided for the SuperCluster installation:

- A 10GbE switch for connectivity from the compute server 10GbE ports to your client access network.

  You can use the provided transceivers and optical cables, or use your own transceivers and cables.

  Note - At the time of installation, if you do not have a 10GbE client access network infrastructure set up at the site, you must provide a 10GbE network switch that SuperCluster can connect to, even if the network speed drops from 10Gb to 1Gb on the other side of the switch.

- Cable connections listed in this table:

Connection Type	Number of cable connections	Comments
Management network	1 for Ethernet management switch	Connect to the existing management network
Client access network	2 per logical domain For SuperCluster M8, you can connect another 2 connections (second port pair), This port pair provides connectivity to another network, or serves as a second link to the same client network.	Connect to the client access network. **Note -** You will not have redundancy through IPMP if there is only one connection per logical domain.

For specific hardware connection options and requirements, refer to Network Infrastructure Requirements in the *Oracle SuperCluster M8 and SuperCluster M7 Installation Guide*.

Related Information

- "SuperCluster Networks Overview" on page 85
- "Default IP Addresses" on page 94
- "Understanding Default Host Names and IP Addresses (Single-Server Version)" on page 95
- "Understanding Default Host Names and IP Addresses (Dual-Server Version)" on page 99

Multiple Client Access Networks

As of SuperCluster software version 2.4, you can have up to eight different client access networks configured on your SuperCluster. These networks are used by dedicated domains, and are configured by Oracle personnel during the installation based on information you provide in the SuperCluster Configuration Worksheets.

You are asked to provide the starting IP address, gateway address and subnet mask for each client access network. After the installation you can use different client access networks for the Oracle RAC groups, dedicated application and database domains, and zones in your SuperCluster.

Note - For I/O Domains, any time after the installation, you can configure networks using the SuperCluster Virtual Assistant.

Network Recipes

Network recopies are predefined network parameters that are used to simplify the network configuration process.

As of SuperCluster software version 2.4, you define up to eight network recopies in the Configuration Worksheets. The network recopies are available to be used when configuring domain networks.

Each network recipe defines these parameters:

Parameter	Example
Recipe name	Western_US
Domain name	example.com
IP addresses of name servers	192.0.2.11 198.51.100.5
IP addresses of time servers	192.0.2.12 198.51.100.8
Time zone	America/Los_Angeles

VLAN Tags

VLAN tagging segregates traffic between domains so that a domain only sees the traffic on its virtual network. The VLAN tags are used for Oracle RAC groups, application domains, and dedicated database domains containing zones.

As of SuperCluster software version 2.4, you can have up to sixteen virtual LAN (VLAN) tags. You specify the number of VLAN tags and the VLAN tag IDs in the SuperCluster Configuration Worksheets. The VLAN tags are configured in SuperCluster by Oracle personnel.

Network Features Applied to RAC Clusters

With the SuperCluster software version 2.4 , you can apply new features (multiple client access networks, network recipes and VLAN tagging) to RAC clusters.

For example, to take advantage of all three of the new features, provide this information in the Configuration Worksheets:

- • Multiple client access networks (for example, four separate client access networks, labeled client access networks 1 - 4)
- • Multiple network recipes (for example, four network recipes, labeled network recipes 1 - 4)
- • Multiple VLAN tags (for example, four VLAN tags, labeled VLANs 101 – 104)

If you are using four RAC clusters in your SuperCluster. You can assign a different client access network, network recipe, and VLAN tag to each of these four RACs, such as the following configurations:

- • RAC 1:
  - o Client access network 1
  - o Network recipe 1
  - o VLAN tag 101
- • RAC 2:
  - o Client access network 2
  - o Network recipe 2
  - o VLAN tag 102

And so on. You can also have multiple RAC clusters using the same configurations, where, for example, RACs 3 and 4 are both be on the same client access network, use the same network recipe, and have the same VLAN tags.

Client Access Network IP Addresses

As of SuperCluster software version 2.4, you can have up to eight different client access networks for dedicated domains. Provide this information in the Configuration Worksheets so that Oracle personnel configure SuperCluster with the desired number of networks during the installation.

Keep in mind that after the installation, you might want to configure additional database zones and I/O Domains , therefore additional IP addresses will be needed for the 10GbE client access network. Plan accordingly for the total number of IP addresses needed for the 10GbE client access network for the future.

Minimum Required Client Network IP Addresses Required for the Installation:

- One 10GbE client access IP address for certain domains:
  - One 10GbE client access IP address for every dedicated domain (Database Domain or Application Domain) in each PDomain.

    - For Root Domains, one 10GbE client access IP address only if the Root Domain is the first domain in the first PDomain.

      Note that no other Root Domains in your SuperCluster M8 or SuperCluster M7 need a 10GbE client access IP address.
- One 10GbE client access IP address for every database zone in a Database Domain that will be set up by your Oracle installer .
- One 10GbE client access IP address for every I/O Domain that will be set up by your Oracle installer.
- Two 10GbE client access IP addresses for the ZFS storage controllers. Note that these IP addresses are automatically pulled from the first (default) client access network.
- 10GbE client access IP addresses for Oracle RAC VIP and SCAN for every Database Domain (either dedicated domain or Database I/O Domain) and database zone that are part of a RAC:
    - One Oracle RAC VIP address for each Database Domain (either dedicated domain or Database I/O Domain) that is part of a RAC
    - One Oracle RAC VIP address for every database zone within a Database Domain (dedicated domain) that is part of a RAC
    - Three SCAN IP addresses for each Oracle RAC in your SuperCluster M8 or SuperCluster M7

Management Network IP Addresses

You need management network IP addresses for these components in SuperCluster M8 or SuperCluster M7:

- • One 1GbE host management IP address for every dedicated domain (Database Domain or Application Domain) and Root Domain in each PDomain.
- • One 1GbE host management IP address for every database zone in a Database Domain that will be set up by your Oracle installer .
- • One 1GbE host management IP address for every I/O Domain that will be set up by your Oracle installer6.
- • Two 1GbE host management IP addresses for every PDomain in your system (two for each service processor in your system)
- • One 1GbE host management IP address for every SPARC M8 or SPARC M7 server in your SuperCluster (one for each floating service processor alias in each server)
- • One 1GbE host management IP address for each of these components:
    - • Cisco Catalyst switch

- • IB switches (3)
- • PDUs (2)
- • Storage servers
- • ZFS storage controllers (2)
- • ISCSI for ZFS controllers (2)

- • One Oracle ILOM IP address for each of these components:
  - • Storage servers
  - • ZFS storage controllers (2)

Note - It is best practice to have all the IP addresses on this network in sequential order. If you cannot set aside the appropriate number of sequential IP addresses for this network, and you must break the IP addresses into nonsequential addresses, the Oracle installer can break the IP addresses on this network into nonsequential blocks. However, this makes the information in the Installation Template more complex, and requires additional communication between you and your Oracle representative to ensure that the nonsequential IP addresses are correctly assigned to the appropriate components or domains in the system.

Keep in mind that you might want to configure additional database zones and I/O Domains after the initial installation, therefore additional IP addresses will be required. Plan accordingly for the total number of IP addresses needed for the SuperCluster networks in the future.

Default IP Addresses

Four sets of default IP addresses are assigned at manufacturing:

- **Management IP addresses** — IP addresses used by Oracle ILOM for the compute servers, storage servers, and the storage controllers.
- **Host IP addresses** — Host IP addresses used by the compute servers, storage servers, storage controllers, and switches.
- **IB IP addresses** — IB interfaces are the default channel of communication among compute servers, storage servers, and the storage controllers. If you are connecting SuperCluster to another SuperCluster or to an Oracle Exadata or Exalogic machine on the same IB fabric, the IB interface enables communication between the compute servers and storage server heads in one SuperCluster and the other SuperCluster or Oracle Exadata or Exalogic machine.
- **10GbE IP addresses** — The IP addresses used by the 10GbE client access network interfaces.

Related Information

- “SuperCluster Networks Overview” on page 85
- “Network Connectivity Requirements” on page 89
- “Understanding Default Host Names and IP Addresses (Single-Server Version)” on page 95
- “Understanding Default Host Names and IP Addresses (Dual-Server Version)” on page 99

Understanding Default Host Names and IP Addresses (Single-Server Version)

These topics describe the default IP addresses used in SuperCluster M8or SuperCluster M7 when one compute server is installed in the rack:

- “Default Host Names and IP Addresses for the Management Network (Single-Server Version)” on page 95
- “Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Single-Server Version)” on page 97

Default Host Names and IP Addresses for the Management Network (Single-Server Version)

TABLE 23 Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Single-Server Version)

		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	N/A
	PCU-B (right from rear view)	sscpdub	192.168.1.211	N/A	N/A
42	Storage Server 4	ssces4-sp	192.168.1.104	cell04	192.168.1.4
41					
40	Storage Server 5	ssces5-sp	192.168.1.105	cell05	192.168.1.5
39					
38	Storage Server 6	ssces6-sp	192.168.1.106	cell06	192.168.1.6
37					

		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**
36	Storage Server 7	ssces7-sp	192.168.1.107	cell07	192.168.1.70
35					
34	Storage Server 8	ssces8-sp	192.168.1.108	cell08	192.168.1.71
33					
32	Storage Server 9	ssces9-sp	192.168.1.109	cell09	192.168.1.72
31					
30	Storage Server 10	ssces10-sp	192.168.1.110	cell10	192.168.1.73
29					
28	Storage Server 11	ssces11-sp	192.168.1.111	cell11	192.168.1.74
27					
26	Storage Controller 2	sscsn2-sp	192.168.1.116	sscsn2	192.168.1.16
25	Storage Controller 1	sscsn1-sp	192.168.1.115	sscsn1	192.168.1.15
24	IB Switch (Leaf 2)	sscnm3	192.168.1.203	N/A	N/A
23	Disk Shelf for the ZFS Storage Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Ethernet Management Switch	ssc4948	192.168.1.200	N/A	N/A
18	IB Switch (Leaf 1)	sscnm2	192.168.1.202	N/A	N/A
17	Compute Server 1: Top Half (CMIOU slots 4-7)			ssccn2	192.168.1.10
16					
15					
14					
13					
12	Compute Server 1: Bottom Half (CMIOU slots 0-3)	sscch1-sp	**SuperCluster M7:** 192.168.1.122 **SuperCluster M8:** 192.168.1.120	ssccn1	192.168.1.9
11					
10		sscch1-sp1	**SuperCluster M7:**192.168.1.121 **SuperCluster M8:** 192.168.1.122		
9					

		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**
8		sscch1-sp0	**SuperCluster M7:**192.168.1.120 **SuperCluster M8:** 192.168.1.121		
7	Storage Server 3	ssces3-sp	192.168.1.103	cell03	192.168.1.3
6					
5	Storage Server 2	ssces2-sp	192.168.1.102	cell02	192.168.1.2
4					
3	Storage Server 1	ssces1-sp	192.168.1.101	cell01	192.168.1.1
2					
1	IB Switch (Spine)	sscnm1	192.168.1.201	N/A	N/A

Related Information

- "Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Single-Server Version)" on page 97

Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Single-Server Version)

The IB network is a private nonroutable network. The IB network IP addresses and host names assigned to the components and domains for the IB network should not be registered in the DNS.

If there are conflicts with the default IP addresses for the IB network and existing IP addresses already on your network, or if this is another SuperCluster system that is being monitored through the same Enterprise Controller host (see "IP Addresses and Oracle Enterprise Manager Ops Center" on page 104), you can change the default IP addresses. The addresses for the components associated with the Database Domains must remain on a different subnet from the addresses for the components associated with the ZFS storage appliance.

TABLE 24 Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Single-Server Version)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**IB Host Names**	**IB IP Addresses**	**10GbE Client Access Host Names**	**10GbE 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	Storage Server 4	ssces4-stor	192.168.10.107	N/A	N/A
41					
40	Storage Server 5	ssces5-stor	192.168.10.109	N/A	N/A
39					
38	Storage Server 6	ssces6-stor	192.168.10.111	N/A	N/A
37					
36	Storage Server 7	ssces7-stor	192.168.10.113	N/A	N/A
35					
34	Storage Server 8	ssces8-stor	192.168.10.115	N/A	N/A
33					
32	Storage Server 9	ssces9-stor	192.168.10.117	N/A	N/A
31					
30	Storage Server 10	ssces10-stor	192.168.10.119	N/A	N/A
29					
28	Storage Server 11	ssces11-stor	192.168.10.121	N/A	N/A
27					
26	Storage Controller 2	sscsn2-stor	N/A (clustered)	N/A	N/A
25	Storage Controller 1	sscsn1-stor	192.168.10.15	N/A	N/A
24	IB Switch (Leaf 2)	N/A	N/A	N/A	N/A
23	Disk Shelf for the ZFS Storage Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Ethernet Management Switch	N/A	N/A	N/A	N/A
18	IB Switch (Leaf 1)	N/A	N/A	N/A	N/A
17	Compute Server 1: Top Half (CMIOU slots 4-7)	ssccn2-ib4	192.168.10.40	ssccn2-tg8 ssccn2-tg7	192.168.40.24 192.168.40.23
16		ssccn2-ib3	192.168.10.30	ssccn2-tg6 ssccn2-tg5	192.168.40.22 192.168.40.21
15		ssccn2-ib2	192.168.10.20	ssccn2-tg4	192.168.40.20

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**IB Host Names**	**IB IP Addresses**	**10GbE Client Access Host Names**	**10GbE Client Access IP Addresses**
				ssccn2-tg3	192.168.40.19
14		ssccn2-ib1	192.168.10.10	ssccn2-tg2	192.168.40.18
				ssccn2-tg1	192.168.40.17
13					
12	Compute Server 1: Bottom Half (CMIOU slots 0-3)	ssccn1-ib4	192.168.10.39	ssccn1-tg8	192.168.40.8
				ssccn1-tg7	192.168.40.7
11		ssccn1-ib3	192.168.10.29	ssccn1-tg6	192.168.40.6
				ssccn1-tg5	192.168.40.5
10		ssccn1-ib2	192.168.10.19	ssccn1-tg4	192.168.40.4
				ssccn1-tg3	192.168.40.3
9		ssccn1-ib1	192.168.10.9	ssccn1-tg2	192.168.40.2
				ssccn1-tg1	192.168.40.1
8					
7	Storage Server 3	ssces3-stor	192.168.10.105	N/A	N/A
6					
5	Storage Server 2	ssces2-stor	192.168.10.103	N/A	N/A
4					
3	Storage Server 1	ssces1-stor	192.168.10.101	N/A	N/A
2					
1	IB Switch (Spine)	N/A	N/A	N/A	N/A

Related Information

- “Default Host Names and IP Addresses for the Management Network (Single-Server Version)” on page 95

Understanding Default Host Names and IP Addresses (Dual-Server Version)

Refer to the following topics for the default IP addresses used in SuperCluster M8 and SuperCluster M7 when two compute servers are installed in the rack:

- "Default Host Names and IP Addresses for the Management Network (Dual-Server Version)" on page 100
- "Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Dual-Server Version)" on page 102

Default Host Names and IP Addresses for the Management Network (Dual-Server Version)

TABLE 25 Default Host Names and IP Addresses for the Oracle ILOM and Host Management Networks (Dual-Server Version)

		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	N/A
	PCU-B (right from rear view)	sscpdub	192.168.1.211	N/A	N/A
42	Storage Server 4	ssces4-sp	192.168.1.104	cell04	192.168.1.4
41					
40	Storage Server 5	ssces5-sp	192.168.1.105	cell05	192.168.1.5
39					
38	Storage Server 6	ssces6-sp	192.168.1.106	cell06	192.168.1.6
37					
36	Compute Server 2: Top Half (CMIOU slots 4-7)			ssccn4	192.168.1.12
35					
34					
33					
32					
31	Compute Server 2: Bottom Half (CMIOU slots 0-3)	sscch2-sp	**SuperCluster M7:** 192.168.1.127 **SuperCluster M8:** 192.168.1.125	ssccn3	192.168.1.11
30					
29		sscch2-sp1	**SuperCluster M7:** 192.168.1.126 **SuperCluster M8:** 192.168.1.127		
28					

		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**
27		sscch2-sp0	**SuperCluster M7:** 192.168.1.125 **SuperCluster M8:** 192.168.1.126		
26	Storage Controller 2	sscsn2-sp	192.168.1.116	sscsn2	192.168.1.16
25	Storage Controller 1	sscsn1-sp	192.168.1.115	sscsn1	192.168.1.15
24	IB Switch (Leaf 2)	sscnm3	192.168.1.203	N/A	N/A
23	Disk Shelf for the ZFS Storage Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Ethernet Management Switch	ssc4948	192.168.1.200	N/A	N/A
18	IB Switch (Leaf 1)	sscnm2	192.168.1.202	N/A	N/A
17	Compute Server 1: Top Half (CMIOU slots 4-7)			ssccn2	192.168.1.10
16					
15					
14					
13					
12	Compute Server 1: Bottom Half (CMIOU slots 0-3)	sscch1-sp	**SuperCluster M7:** 192.168.1.122 **SuperCluster M8:** 192.168.1.120	ssccn1	192.168.1.9
11					
10		sscch1-sp1	**SuperCluster M7:** 192.168.1.121 **SuperCluster M8:** 192.168.1.122		
9					
8		sscch1-sp0	**SuperCluster M7:** 192.168.1.120 **SuperCluster M8:** 192.168.1.121		
7	Storage Server 3	ssces3-sp	192.168.1.103	cell03	192.168.1.3
6					
5	Storage Server 2	ssces2-sp	192.168.1.102	cell02	192.168.1.2

		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**
4					
3	Storage Server 1	ssces1-sp	192.168.1.101	cell01	192.168.1.1
2					
1	IB Switch (Spine)	sscnm1	192.168.1.201	N/A	N/A

Related Information

- "Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Dual-Server Version)" on page 102

Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Dual-Server Version)

The IB network is a private nonroutable network. The IB network IP addresses and host names assigned to the components and domains for the IB network should not be registered in the DNS.

If there are conflicts with the default IP addresses for the IB network and existing IP addresses already on your network, or if this is another SuperCluster system that is being monitored through the same Enterprise Controller host.

TABLE 26 Default Host Names and IP Addresses for the IB and 10GbE Client Access Networks (Dual-Server Version)

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**IB Host Names**	**IB IP Addresses**	**10GbE Client Access Host Names**	**10GbE 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	Storage Server 4	ssces4-stor	192.168.10.107	N/A	N/A
41					
40	Storage Server 5	ssces5-stor	192.168.10.109	N/A	N/A
39					
38	Storage Server 6	ssces6-stor	192.168.10.111	N/A	N/A
37					

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**IB Host Names**	**IB IP Addresses**	**10GbE Client Access Host Names**	**10GbE Client Access IP Addresses**
36	Compute Server 2: Top Half (CMIOU slots 4-7)	ssccn4-ib4	192.168.10.160	ssccn4-tg8	192.168.40.56
				ssccn4-tg7	192.168.40.55
35		ssccn4-ib3	192.168.10.150	ssccn4-tg6	192.168.40.54
				ssccn4-tg5	192.168.40.53
34		ssccn4-ib2	192.168.10.140	ssccn4-tg4	192.168.40.52
				ssccn4-tg3	192.168.40.51
33		ssccn4-ib1	192.168.10.130	ssccn4-tg2	192.168.40.50
				ssccn4-tg1	192.168.40.49
32					
31	Compute Server 2: Bottom Half (CMIOU slots 0-3)	ssccn3-ib4	192.168.10.120	ssccn3-tg8	192.168.40.40
				ssccn3-tg7	192.168.40.39
30		ssccn3-ib3	192.168.10.115	ssccn3-tg6	192.168.40.38
				ssccn3-tg5	192.168.40.37
29		ssccn3-ib2	192.168.10.110	ssccn3-tg4	192.168.40.36
				ssccn3-tg3	192.168.40.35
28		ssccn3-ib1	192.168.10.90	ssccn3-tg2	192.168.40.34
				ssccn3-tg1	192.168.40.33
27					
26	Storage Controller 2	sscsn2-stor	N/A (clustered)	N/A	N/A
25	Storage Controller 1	sscsn1-stor	192.168.10.15	N/A	N/A
24	IB Switch (Leaf 2)	N/A	N/A	N/A	N/A
23	Disk Shelf for the ZFS Storage Appliance	N/A	N/A	N/A	N/A
22					
21					
20					
19	Ethernet Management Switch	N/A	N/A	N/A	N/A
18	IB Switch (Leaf 1)	N/A	N/A	N/A	N/A
17	Compute Server 1: Top Half (CMIOU slots 4-7)	ssccn2-ib4	192.168.10.40	ssccn2-tg8	192.168.40.24
				ssccn2-tg7	192.168.40.23
16		ssccn2-ib3	192.168.10.30	ssccn2-tg6	192.168.40.22
				ssccn2-tg5	192.168.40.21

		Information Assigned at Manufacturing			
Unit Number	**Rack Component (Front View)**	**IB Host Names**	**IB IP Addresses**	**10GbE Client Access Host Names**	**10GbE Client Access IP Addresses**
15		ssccn2-ib2	192.168.10.20	ssccn2-tg4	192.168.40.20
				ssccn2-tg3	192.168.40.19
14		ssccn2-ib1	192.168.10.10	ssccn2-tg2	192.168.40.18
				ssccn2-tg1	192.168.40.17
13					
12	Compute Server 1: Bottom Half (CMIOU slots 0-3)	ssccn1-ib4	192.168.10.39	ssccn1-tg8	192.168.40.8
				ssccn1-tg7	192.168.40.7
11		ssccn1-ib3	192.168.10.29	ssccn1-tg6	192.168.40.6
				ssccn1-tg5	192.168.40.5
10		ssccn1-ib2	192.168.10.19	ssccn1-tg4	192.168.40.4
				ssccn1-tg3	192.168.40.3
9		ssccn1-ib1	192.168.10.9	ssccn1-tg2	192.168.40.2
				ssccn1-tg1	192.168.40.1
8					
7	Storage Server 3	ssces3-stor	192.168.10.105	N/A	N/A
6					
5	Storage Server 2	ssces2-stor	192.168.10.103	N/A	N/A
4					
3	Storage Server 1	ssces1-stor	192.168.10.101	N/A	N/A
2					
1	IB Switch (Spine)	N/A	N/A	N/A	N/A

IP Addresses and Oracle Enterprise Manager Ops Center

For previous versions of Oracle Enterprise Manager Ops Center, the Ops Center software was installed and run from the SuperCluster system. Beginning with the Oracle Enterprise Manager Ops Center 12c Release 2 (12.2.0.0.0) release, the Ops Center software must now run on a system (Enterprise Controller host) outside of the SuperCluster system.

The following conditions apply to Oracle engineered systems, such as SuperCluster systems.

One or more Oracle engineered systems can be discovered and managed by a single Oracle Enterprise Manager Ops Center instance based on these conditions:

- • None of the Oracle engineered system instances have overlapping private networks connected through IB, that is, networks that have the same CIDR (Classless Inter-Domain Routing) or networks that are subblocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping.
- • None of the Oracle engineered system instances or generic datacenter assets have overlapping management or client access networks connected through Ethernet, that is, networks that have the same CIDR or networks that are subblocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping. As an exception, you can use the same CIDR (not subblock) for multiple systems. For example, you can use 192.0.2.1/22 as a CIDR for Ethernet network on one or more engineered systems or generic datacenter assets.
- • None of the Oracle engineered system instances have overlapping public networks connected through EoIB, that is, networks that have the same CIDR or networks that are subblocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping. As an exception, you can use the same CIDR (not subblock) for multiple systems. For example, you can use 192.2.0.0/22 as a CIDR for public EoIB network on multiple engineered systems.
- • None of the networks configured in Oracle Enterprise Manager Ops Center overlaps with any network, that is, overlapping networks are not supported by Oracle Enterprise Manager Ops Center.

Note - To manage two or more Oracle engineered systems that have overlapping networks or any networks already present in Oracle Enterprise Manager Ops Center, reconfigure one of the conflicting systems before it is discovered and managed by the same Oracle Enterprise Manager Ops Center.

Example

The following are example SuperCluster network configurations that you can use when configuring the network to discover and manage SuperCluster systems. Status OK indicates a valid configuration, and status Fail indicates an invalid configuration.

TABLE 27 Example Configuration 1

	1 GbE	10 GbE	IB
First SuperCluster System	192.0.251.0/21	192.4.251.0/24	192.168.8.0/22
Second SuperCluster System	192.0.251.0/21	192.4.251.0/24	192.168.12.0/22
Status	OK	OK	OK

Result:

OK – First SuperCluster system 1GbE and second SuperCluster system 1GbE share the same network.

OK – First SuperCluster system 10GbE and second SuperCluster system 10GbE share the same network.

OK – First SuperCluster system IB network does not overlap with second SuperCluster system IB.

TABLE 28 Example Configuration 2

	1 GbE	10 GbE	IB
First SuperCluster System	192.0.251.0/21	192.4.251.0/24	192.168.8.0/22
Second SuperCluster System	192.0.0.128/25	192.0.7.0/24	192.168.8.0/22
Status	FAIL	OK	FAIL

Result:

FAIL – First SuperCluster system 1GbE and second SuperCluster system 1GbE have overlapping networks.

OK – First SuperCluster system 10GbE and second SuperCluster system 10GbE have different non-overlapping networks.

FAIL – First SuperCluster system IB and second SuperCluster system IB networks do not have unique private networks (racks are not interconnected).

Glossary

A

Application Domain — A domain that runs Oracle Solaris and client applications.

ASMM — Automatic shared memory management.

ASR — Auto Service Request. A feature of Oracle or Sun hardware that automatically opens service requests when specific hardware faults occur. ASR is integrated with MOS and requires a support agreement. See also MOS.

C

CFM — Cubic feet per minute.

Cisco Catalyst Ethernet switch — Provides the SuperCluster M7 management network. Referred to in this documentation using the shortened name “Ethernet management switch.” See also Ethernet management switch.

CMIOU — CPU, memory, and I/O unit. Each CMIOU contains 1 CMP, 16 DIMM slots, and 1 I/O hub chip. Each CMIOU also hosts an eUSB device.

COD — Capacity on Demand.

compute server — Shortened name for the SPARC M7 server, a major component of SuperCluster M7.

D

Database Domain — The domain that contains the SuperCluster M7 database.

DB Oracle Database.

DCM Domain configuration management. The reconfiguration of boards in PDomains for Enterprise-class systems. See also PDomain.

dedicated domain A SuperCluster LDom category that includes the domains configured at installation time as either a Database Domain or an Application Domain (running the Oracle Solaris 11 OS). Dedicated domains have direct access to the 10GbE NICs and IB HCAs (and Fibre Channel cards, if present). See also Database Domain and Application Domain.

DHCP Dynamic Host Configuration Protocol. Software that automatically assigns IP addresses to clients on a TCP/IP network. See also TCP.

DIMM Dual in-line memory module.

DISM Dynamic intimate shared memory.

E

EECS Oracle Exalogic Elastic Cloud software.

EPO switch Emergency power-off switch.

ESD Electrostatic discharge.

Ethernet management switch Shortened name for the Cisco Catalyst Ethernet switch. See also Cisco Catalyst Ethernet switch.

eUSB Embedded USB. A flash-based drive designed specifically to be used as a boot device. An eUSB does not provide storage for applications or customer data.

expansion rack Shortened name for optional Oracle Exadata Storage Expansion Racks (up to 17) that can be added to SuperCluster M7. See also Oracle Exadata Storage Expansion Rack.

F

FAN Fast application notification event.

FCoE Fibre Channel over Ethernet.

FM Fan module.

FMA Fault management architecture. A feature of Oracle Solaris servers that includes error handlers, structured error telemetry, automated diagnostic software, response agents, and messaging.

FRU Field-replaceable unit.

G

GB Gigabyte. 1 gigabyte = 1024 megabytes.

GbE Gigabit Ethernet.

GNS Grid Naming Service.

H

HCA Host channel adapter.

HDD Hard disk drive. In Oracle Solaris OS output, HDD can refer to hard disk drives or SSDs.

I

I/O Domain If you have Root Domains, you create I/O Domains with your choice of resources at the time of your choosing. The SuperCluster Virtual Assistant enables you to assign resources to I/O Domains from the CPU and memory repositories, and from virtual functions hosted by Root Domains. When you create an I/O Domain, you assign it as a Database Domain or Application Domain running the Oracle Solaris 11 OS. See also Root Domain.

IB InfiniBand.

IB switch Shortened name for the Sun Datacenter InfiniBand Switch 36. See also leaf switch, spine switch, and Sun Datacenter InfiniBand Switch 36.

ILOM See Oracle ILOM.

IPMI Intelligent Platform Management Interface.

IPMP IP network multipathing.

iSCSI Internet Small Computer System Interface.

K

KVMS Keyboard video mouse storage.

L

LDom Logical domain. A virtual machine comprising a discrete logical grouping of resources that has its own operating system and identity within a single computer system. LDoms are created using Oracle VM Server for SPARC software. See also Oracle VM Server for SPARC.

leaf switch Two of the IB switches are configured as leaf switches, the third is configured as a spine switch. See also IB switch.

M

MIB Management information base.

MOS My Oracle Support.

N

NET MGT The network management port on an SP. See also SP.

NIC Network interface card.

NUMA Nonuniform memory access.

O

OBP OpenBoot PROM. Firmware on SPARC servers that enables the server to load platform-independent drivers directly from devices, and provides an interface through which you can boot the compute server and run low-level diagnostics.

OCM Oracle Configuration Manager.

ONS Oracle Notification Service.

Oracle ASM Oracle Automatic Storage Management. A volume manager and a file system that supports Oracle databases.

Oracle Exadata Storage Expansion Rack Optional expansion racks that can be added to SuperCluster M7 systems that require additional storage. Referred to in this documentation using the shortened name “expansion rack.” See also expansion rack.

Oracle ILOM Oracle Integrated Lights Out Manager. Software on the SP that enables you to manage a server independently from the operating system. See also SP.

Oracle Solaris OS Oracle Solaris operating system.

Oracle SuperCluster Refers to all Oracle SuperCluster models.

Oracle SuperCluster M7 The full name of the SuperCluster M7 systems. Referred to in this documentation using the shortened name “SuperCluster M7.” See also SuperCluster M7.

Oracle VM Server for SPARC SPARC server virtualization and partitioning technology. See also LDom.

Oracle VTS Oracle Validation Test Suite. An application, preinstalled with Oracle Solaris, that exercises the system, provides hardware validation, and identifies possible faulty components.

Oracle XA Oracle's implementation of the X/Open distributed transaction processing XA interface that is included in Oracle DB software.

Oracle ZFS ZS3-ES storage appliance Provides SuperCluster M7 with shared storage capabilities. Referred to in this documentation using the shortened name “ZFS storage appliance.” See also ZFS storage appliance.

OS Operating system.

P

parked resources CPU and memory resources that are set aside in the CPU and memory repositories. You assign parked resources to I/O Domains with the SuperCluster Virtual Assistant.

PCIe Peripheral Component Interconnect Express.

PDomain Physical domain. Each PDomain on the compute server is an independently configurable and bootable entity with full hardware domain isolation for fault isolation and security purposes. See also compute server and SSB.

PDomain-SPP The lead SPP of a PDomain. The PDomain-SPP on the compute server manages tasks and provides rKVMS service for that PDomain. See also PDomain.

PDU Power distribution unit.

PF Physical function. Functions provided by physical I/O devices, such as the IB HCAs, 10GbE NICs, and any Fibre Channel cards installed in the PCIe slots. Logical devices, or virtual functions (VFs), are created from PFs, with each PF hosting 16 VFs.

POST Power-on self-test. A diagnostic that runs when the compute server is powered on.

PS Power supply.

PSDB Power system distribution board.

PSH Predictive self healing. An Oracle Solaris OS technology that continuously monitors the health of the compute server and works with Oracle ILOM to take a faulty component offline if needed.

Q

QMU Quarterly maintenance update.

QSFP Quad small form-factor, pluggable. A transceiver specification for 10GbE technology.

R

RAC Real Application Cluster.

RCLB Runtime connection load balancing.

rKVMS Remote keyboard video mouse and storage.

root complex CMP circuitry that provides the base to a PCIe I/O fabric. Each PCIe I/O fabric consists of the PCIe switches, PCIe slots, and leaf devices associated with the root complex.

Root Domain A logical domain that is configured at installation time. Root Domains are required if you plan to configure I/O Domains. Root Domains host PFs from which I/O Domains derive VFs. The majority of Root Domain CPU and memory resources are parked for later use by I/O Domains.

S

SAS Serial attached SCSI.

SATA Serial advance technology attachment.

scalability The ability to increase (or scale up) processing power in a compute server by combining the server's physical configurable hardware into one or more logical groups (see also PDomain).

SCAN Single Client Access Name. A feature used in RAC environments that provides a single name for clients to access any Oracle Database running in a cluster. See also RAC.

SDP Session Description Protocol.

SER MGT The serial management port on an SP. See also SP.

SFP+ Small form-factor pluggable standard. SFP+ is a specification for a transceiver for 10GbE technology.

SGA System global area.

SMF Service Management Facility.

SNEEP Serial number in EEPROM.

SNMP Simple Management Network Protocol.

SP Service processor. A processor, separate from the host, that monitors and manages the host no matter what state the host is in. The SP runs Oracle ILOM, which provides remote lights out management. In SuperCluster M7, SPs are located on the compute servers, storage servers, ZFS storage appliance controllers, and IB switches. See also Oracle ILOM.

SPARC M7-8 server A major component of SuperCluster M7 that provides the main compute resources. Referred to in this documentation using the shortened name “compute server.” See also compute server.

spine switch One of the SuperCluster M7 IB switches that is configured as a spine switch. See also IB switch and leaf switch.

SPP Service processor proxy. One SPP in the compute server is assigned to manage each PDomain. SPPs monitor environmental sensors and manage the CMIOUs, memory controllers, and DIMMs. See also PDomain-SPP.

SR-IOV Domain	Single-Root I/O Virtualization Domain. A SuperCluster logical domain category that includes Root Domains and I/O Domains. This category of domains support single-root I/O virtualization. See also I/O Domain and Root Domain.
SSB	Scalability switch board in the compute server.
SSD	Solid state drive.
STB	Oracle Services Tool Bundle.
storage server	Storage servers in SuperCluster M7.
Sun Datacenter InfiniBand Switch 36	Interconnects SuperCluster M7 components on a private network. Referred to in this documentation using the shortened name “IB switch.” See also IB switch, leaf switch, and spine switch.
SuperCluster M7	Shortened name for Oracle SuperCluster M7 systems. See also Oracle SuperCluster M7.

T

TCP	Transmission Control Protocol.
TNS	Transparent Network Substrate.
TPM	Trusted platform module.

U

UPS	Uninterruptible power supply.

V

VAC	Voltage alternating current.
VF	Virtual function. Logical I/O devices that are created from PFs, with each PF hosting 16 VFs.
VIP	Virtual IP.

VLAN Virtual local area network.

VNET Virtual network.

W

WWN World Wide Name.

X

XA See Oracle XA.

Z

ZFS A file system with added volume management capabilities. ZFS is the default file system in Oracle Solaris 11.

ZFS disk shelf A component of the ZFS storage appliance that contains the storage. The ZFS disk shelf is controlled by the ZFS storage controllers. See also ZFS storage appliance and ZFS storage controller.

ZFS storage appliance Shortened name for Oracle ZFS Storage ZS3-ES storage appliance. See also Oracle ZFS ZS3-ES storage appliance.

ZFS storage controller Servers in the Oracle ZFS ZS3-ES storage appliance that manage the storage appliance. See also ZFS storage appliance.

Index

S

U

V