Oracle® SuperCluster M6-32

Owner's Guide: Overview

Part No.: E41535-04
September 2015

Please
Recycle

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

Accès au support électronique

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 xii

Determining SuperCluster M6-32 Configurations 1

Configuration Flowchart 1

- Determine the Number of Compute Servers 3
- Determine the Number of DCUs in Each Compute Server 4
- Determine the Number of CMUs in Each DCU 4

CMU Overview 5

- Determine the Amount of Memory in Each DCU 6
- Determine the PDomain Configuration on Each Compute Server 7
- Determine the LDom Configuration for Each PDomain 7

Determining the Best Configuration for Your Situation 9

Understanding PDomain Configurations 9

Understanding High Redundancy 11

High Redundancy Flowchart 11

Number of Compute Servers 13

Type of Configuration PDomains 13

Allocating CPU Resources for LDoms 14

CPU Resources for LDoms Flowchart 14

Compute Server Level Considerations 16

LDom Level Considerations 16

Allocating Memory Resources for LDoms 17

Memory Resources for LDoms Flowchart 17
Compute Server Level Considerations 19
LDom Level Considerations 20
Understanding PCIe Cards and Slots for LDoms 22
PCIe Cards and Slots for LDoms Flowchart 22
Compute Server Level Considerations 24
LDom Level Considerations 24
Understanding Storage for LDoms 25
Storage for LDoms Flowchart 25
Compute Server Level Considerations 27
LDom Level Considerations 27

Understanding the Hardware 29
Understanding SuperCluster M6-32 29
SuperCluster M6-32 Overview 29
Spares Kit Components 31
Configuration Restrictions 31
Identifying SuperCluster M6-32 Components 33
Understanding the Compute Server 33
Compute Server Overview 34
Front Components 35
Rear Components 37
Storage Rack Overview 38

Understanding DCUs 39
DCU Overview 39
Understanding DCU Configurations 42
PCIe Device Root Complexes Overview 42
PCIe Communication and Paths Overview 43

Understanding DCU PCIe and EMS Slot Locations 43
DCU 0 PCIe and EMS Slot Locations 44
DCU 1 PCIe and EMS Slot Locations 45
DCU 2 PCIe and EMS Slot Locations 46
DCU 3 PCIe and EMS Slot Locations 47
Understanding Half-Populated DCU Root Complexes 48
Half-Populated DCU 0 PCIe Slot Root Complexes 48
Half-Populated DCU 1 PCIe Slot Root Complexes 49
Half-Populated DCU 2 PCIe Slot Root Complexes 51
Half-Populated DCU 3 PCIe Slot Root Complexes 52
Understanding Fully-Populated DCU Root Complexes 54
Fully-Populated DCU 0 PCIe Slot Root Complexes 54
Fully-Populated DCU 1 PCIe Slot Root Complexes 56
Fully-Populated DCU 2 PCIe Slot Root Complexes 57
Fully-Populated DCU 3 PCIe Slot Root Complexes 59

Understanding PDomains 61

PDomains Overview 61
PDomain Guidelines 62
Extended Configuration PDomain Overview 63
Understanding Extended Configuration PDomains 64
Understanding Four DCUs in One Compute Server (R1 Extended Configuration PDomains) 64
Four Fully-Populated DCUs (R1_1 PDomain Configuration) 65
Two Fully-Populated DCUs and Two Half-Populated DCUs (R1_2 PDomain Configuration) 66
Four Half-Populated DCUs (R1_3 PDomain Configuration) 66
Understanding Four DCUs Across Two Compute Servers (R2 Extended Configuration PDomains) 67
Four Fully-Populated DCUs (R2_1 PDomain Configuration) 68

Two Fully-Populated DCUs and Two Half-Populated DCUs (R2_2 PDomain Configuration) 69

Four Half-Populated DCUs (R2_3 PDomain Configuration) 69

Base Configuration Overview 70

Understanding Base Configuration PDomains 71

Understanding Four DCUs on One Compute Server (R3 Base Configuration PDomains) 72

Four Fully-Populated DCUs (R3_1 PDomain Configuration) 73

Two Fully-Populated DCUs and Two Half-Populated DCUs (R3_2 PDomain Configuration) 73

Four Half-Populated DCUs (R3_3 PDomain Configuration) 74

Understanding Four DCUs Across Two Compute Servers (R4 Base Configuration PDomains) 74

Four Fully-Populated DCUs (R4_1 PDomain Configuration) 75

Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_2 PDomain Configuration) 76

Four Half-Populated DCUs (R4_3 PDomain Configuration) 76

Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_4 PDomain Configuration) 77

Understanding Two DCUs on One Compute Server (R5 Base Configuration PDomains) 77

Two Fully-Populated DCUs (R5_1 PDomain Configuration) 78

Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains) 78

Two Fully-Populated DCUs (R6_1 PDomain Configuration) 79

Understanding LDoms 81

Understanding Compute Server Hardware and Networks 82

CPU and Memory Resources Overview 82

LDoms and the PCIe Slots Overview 83

Management Network Overview 84

10GbE Client Access Network Overview 84

Understanding the IB Network 85

IB Network Overview 85

IB Network Data Paths for a Database Domain 85

IB Network Data Paths for an Application Domain 86

Understanding Domains 87

Dedicated Domains 87

Understanding SR-IOV Domain Types 88

Root Domains 89

I/O Domains 93

Understanding LDom Configurations for Extended Configuration PDomains 99

Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains) 99

LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains) 100

E4-1 LDom Configuration 101

E4-2 LDom Configuration 103

E4-3 LDom Configuration 104

E4-4 LDom Configuration 106

Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains) 108

LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains) 108

E2-1 LDom Configuration 109

E2-2 LDom Configuration 111

E2-3 LDom Configuration 112

E2-4 LDom Configuration 114

Understanding LDom Configurations for Base Configuration PDomains 116

Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains) 117

LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains) 118

B4-1 LDom Configuration 119

B4-2 LDom Configuration 120

B4-3 LDom Configuration 121

B4-4 LDom Configuration 123

Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains) 124

LDom Configurations for Half-Populated DCUs (Base Configuration PDomains) 125

B2-1 LDom Configuration 126

B2-2 LDom Configuration 127

B2-3 LDom Configuration 128

B2-4 LDom Configuration 130

Understanding Clustering Software 133

Cluster Software Overview 133

Cluster Software for the Database Domain 134

Cluster Software for the Oracle Solaris Application Domains 135

Understanding System Administration Resources 137

Oracle ILOM Overview 137

Understanding Platform-Specific Oracle ILOM Features 138

Server-Specific and New Oracle ILOM Features and Requirements 138

Unsupported Oracle ILOM Features 139

Oracle Solaris OS Overview 139

OpenBoot Overview 140

Oracle ILOM Remote Console Plus Overview 140

Oracle Hardware Management Pack Overview 141

Time Synchronization and NTP Service 141

SNMP Service 142

Multidomain Extensions to Oracle ILOM MIBs 142

LDAP/SSL 144

Active Directory 144

Glossary 145

Index 155

Using This Documentation

This document provides and overview of Oracle SuperCluster M6-32.

- **Overview** – Describes configurations, hardware components, LDom configurations, and clustering software.
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience in system administration
- "Product Documentation Library" on page xii
- "Feedback" on page xii

Product Documentation Library

Documentation and resources for this product and related products are available on the system. Access the documentation by using a browser to view this directory on the first compute server installed in SuperCluster:

`/opt/oracle/node/doc/E41531_01`

Feedback

Provide feedback about this documentation at:

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

Determining SuperCluster M6-32 Configurations

There are several configuration choices that are available to you for your SuperCluster M6-32, both hardware and software. The information in this document helps you decide the best configurations for your SuperCluster M6-32. As you decide on various aspects of the configuration, record them in the *Oracle SuperCluster M6-32 Configuration Worksheets*.

These topics describe how to determine a configuration.

- "Configuration Flowchart" on page 1
- "Determine the Number of Compute Servers" on page 3
- "Determine the Number of DCUs in Each Compute Server" on page 4
- "Determine the Number of CMUs in Each DCU" on page 4
- "CMU Overview" on page 5
- "Determine the Amount of Memory in Each DCU" on page 6
- "Determine the PDomain Configuration on Each Compute Server" on page 7
- "Determine the LDom Configuration for Each PDomain" on page 7

Related Information

- "Determining the Best Configuration for Your Situation" on page 9

Configuration Flowchart

The following flowchart describes the decision-making process for your SuperCluster M6-32.

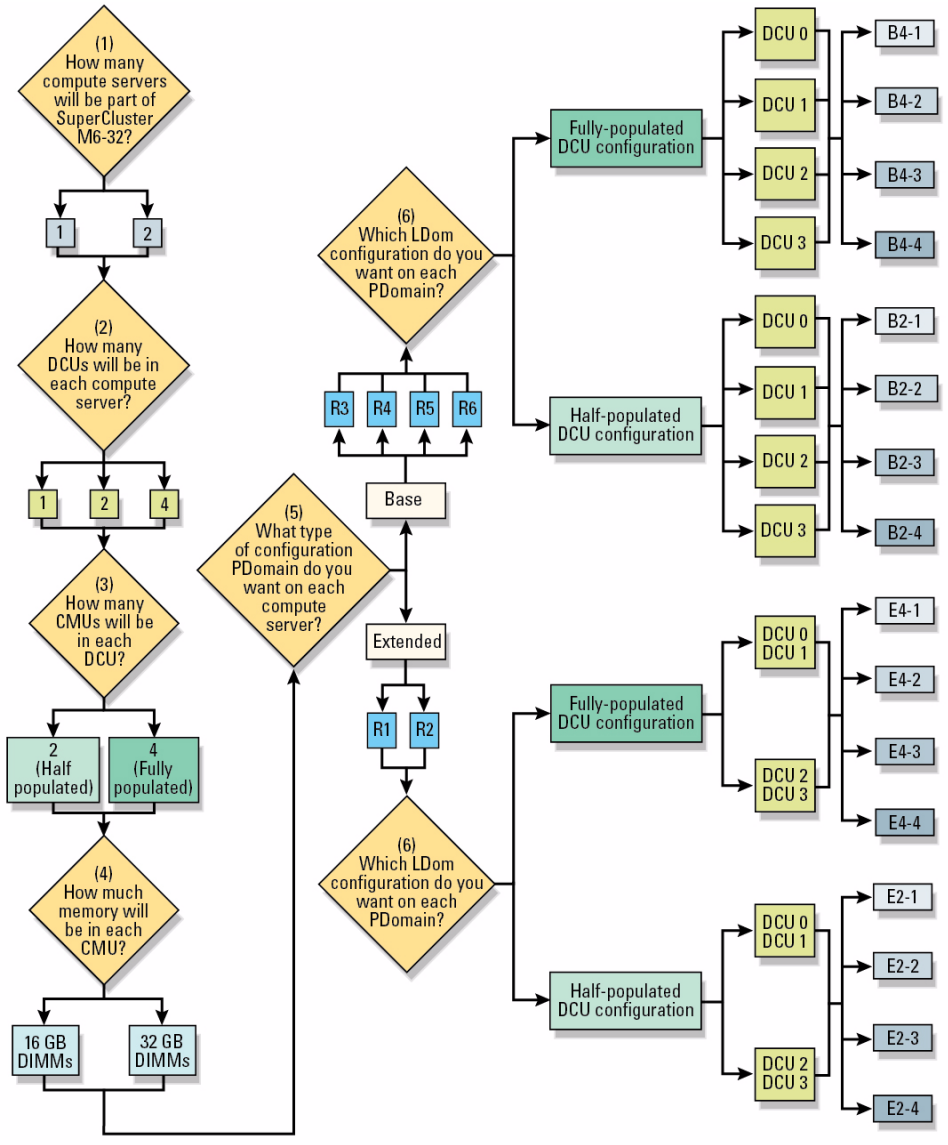

This table describes the steps in the flowchart.

Step	Description	Links
1.	Determine the number of compute servers - 1 or 2.	"Determine the Number of Compute Servers" on page 3
2.	Determine the number of DCUs in each compute server - 1, 2, or 4.	"Determine the Number of DCUs in Each Compute Server" on page 4 "Understanding DCUs" on page 39
3.	Determine the number of CMUs in each DCU - 2 (half-populated) or 4 (fully populated).	"Determine the Number of CMUs in Each DCU" on page 4 "CMU Overview" on page 5
4.	Determine the amount of memory in each CMU - 16 or 32 GByte.	"Determine the Amount of Memory in Each DCU" on page 6
5.	Determine the type of PDomain configuration on each DCU.	"Determine the PDomain Configuration on Each Compute Server" on page 7 "Understanding PDomains" on page 61
6.	Determine the LDom configuration for each PDomain.	"Determine the LDom Configuration for Each PDomain" on page 7 "Understanding LDoms" on page 81

▼ Determine the Number of Compute Servers

1. **Determine whether you want one or two compute servers.**

   Consider that:

   - The compute server can contain up to 32 processors, 32 drives, and 64 PCIe cards. See "Understanding the Compute Server" on page 33.
   - The hardware architecture enables you to divide these hardware resources into smaller units called *PDomains* to address the demands of your applications. The PDomain configurations available for SuperCluster M6-32 provide a level of redundancy at the PDomain level. See "Understanding PDomains" on page 61 for more information.

     You can also select two compute servers for your SuperCluster M6-32 to increase redundancy at the compute server level. See "Understanding High Redundancy" on page 11.

2. **Record the decision in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "Understanding High Redundancy" on page 11

▼ Determine the Number of DCUs in Each Compute Server

You can select one, two, or four DCUs in each compute server. See "Understanding DCUs" on page 39.

1. **Determine how many DCUs you want in each compute server.**

   A number of factors affect this decision:

   - The number of CMUS in each DCU. See "Determine the Number of CMUs in Each DCU" on page 4.
   - The amount of memory that is in each CMU. See "Determine the Amount of Memory in Each DCU" on page 6.
   - The configuration PDomain that you choose. See "Determine the PDomain Configuration on Each Compute Server" on page 7.

2. **Record the decision in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "Understanding DCUs" on page 39
- "Allocating CPU Resources for LDoms" on page 14
- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22
- "Understanding Storage for LDoms" on page 25

▼ Determine the Number of CMUs in Each DCU

Each DCU can contain two (half-populated) or four (fully-populated) CMUs. See "CMU Overview" on page 5.

1. **Determine how many CMUs you want in each DCU.**

   The amount of memory in each CMU affects the number of CMUs in each DCU. See "Determine the Amount of Memory in Each DCU" on page 6.

2. **Record the decisions in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "CMU Overview" on page 5
- "Allocating CPU Resources for LDoms" on page 14
- "Allocating Memory Resources for LDoms" on page 17

CMU Overview

A fully-populated DCU contains four CMUs, and a half-populated DCU contains two CMUs. Each DCU can contain the following components.

DCU	Components
DCU 0	• **Fully-populated:** CMU 0, CMU 1, CMU 2, CMU 3 • **Half-populated:** CMU 0 and CMU 3 • SPP 0 • IOU 0
DCU 1	• **Fully-populated:** CMU 4, CMU 5, CMU 6, CMU 7 • **Half-populated:** CMU 4 and CMU 7 • SPP 1 • IOU 1
DCU 2	• **Fully-populated:** CMU 8, CMU 9, CMU 10, CMU 11 • **Half-populated:** CMU 8 and CMU 11 • SPP 2 • IOU 2
DCU 3	• **Fully-populated:** CMU 12, CMU 13, CMU 14, CMU 15 • **Half-populated:** CMU 12 and CMU 15 • SPP 3 • IOU 3

The number of CMUs that you can have in each DCU is determined by the number of DCUs that you have in each compute server:

- **Two DCUs in a compute server:** Both DCUs are fully-populated (four CMUs in each DCU, or eight CMUs total).
- **Four DCUs in a compute server:**
  - All four DCUs are half-populated (two CMUs in each DCU, or eight CMUs total).
  - All four DCUs are fully-populated (four CMUs in each DCU, or sixteen CMUs total).
  - Two DCUs are half-populated (two CMUs in both DCUs) and two DCUs are fully-populated (four CMUs in both DCUs, or twelve CMUs total).

Note – Certain PDomain configurations have two DCUs associated with each PDomain. For those PDomain configurations, both DCUs must have the same configuration, where both DCUs are either half-populated or fully-populated. See "Understanding PDomains" on page 61 for more information.

▼ Determine the Amount of Memory in Each DCU

Each DCU can contain 16 or 32 GByte of memory. There are 64 DIMM slots in each CMU, or 32 DIMMs slots in each CMP. See "Understanding DCUs" on page 39.

1. **Determine how much memory you want in each DCU.**

   A number of factors affect this decision:

   - The size of the DIMMs in each CMU.
   - The number of CMUs in each DCU. See "Determine the Number of CMUs in Each DCU" on page 4.
   - The LDom configuration on each DCU. See "Determine the LDom Configuration for Each PDomain" on page 7.
   - The configuration PDomain on each server. See "Determine the PDomain Configuration on Each Compute Server" on page 7.

2. **Record the decisions in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "Understanding DCUs" on page 39
- "Allocating Memory Resources for LDoms" on page 17

▼ Determine the PDomain Configuration on Each Compute Server

1. **Determine the PDomain configuration for each compute server.**

   There are two types of PDomain configurations available for the compute servers:

   - Base configuration PDomain – Two or four PDomains, on either one or two severs, where one DCU is associated with each PDomain
   - Extended configuration PDomain – Two PDomains, on either one or two servers, where two DCUs are associated with each PDomain

   See "Understanding PDomains" on page 61.

2. **Record the decisions in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "Understanding PDomains" on page 61
- "Understanding High Redundancy" on page 11
- "Allocating CPU Resources for LDoms" on page 14
- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22
- "Understanding Storage for LDoms" on page 25

▼ Determine the LDom Configuration for Each PDomain

1. **Determine the LDom configuration for each PDomain.**

   There are four types of LDom configurations. See "Understanding LDoms" on page 81.

2. **Record the decisions in the** *Oracle SuperCluster M6-32 Configuration Worksheets***.**

Related Information

- "Understanding LDoms" on page 81
- "Allocating CPU Resources for LDoms" on page 14

- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22
- "Understanding Storage for LDoms" on page 25

Determining the Best Configuration for Your Situation

There are several factors that you should keep in mind when deciding which configuration is correct for you:

- "Understanding PDomain Configurations" on page 9
- "Understanding High Redundancy" on page 11
- "Allocating CPU Resources for LDoms" on page 14
- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22
- "Understanding Storage for LDoms" on page 25

Understanding PDomain Configurations

There are two types of PDomain configurations available for the compute servers:

- Base configuration PDomain – Two or four PDomains, on either one or two severs, where one DCU is associated with each PDomain. Base configuration PDomains allow for greater flexibility around availability.
- Extended configuration PDomain – Two PDomains, on either one or two servers, where two DCUs are associated with each PDomain. Extended configuration PDomains allow for more resources (CPU, memory, IB HCAs, 10GbE ports and internal hard drives) to be concentrated in one of the LDoms in that configuration.

Both configuration types support from 1 to 4 LDoms each. See "Understanding PDomains" on page 61 for more information.

The following table provides information on the amount of resources that's available at the overall PDomain level for the different configurations. Keep the following information in mind when reviewing the information in the following table:

- A fully-populated DCU contains four CMUs, and a half-populated DCU contains two CMUs
- Each CMU contains two processors
- Each CMU contains 64 DIMM slots

PDomain/DCU Configuration	Total Number of Processors	Total Amount of Memory	Total Number of IB HCAs
Extended Configuration PDomain with Fully-Populated DCUs	16	• 8 TB (16 GB DIMMs) • 16 TB (32 GB DIMMs)	8
Extended Configuration PDomain with Half-Populated DCUs	8	• 4 TB (16 GB DIMMs) • 8 TB (32 GB DIMMs)	8
Base Configuration PDomain with Fully-Populated DCUs	8	• 4 TB (16 GB DIMMs) • 8 TB (32 GB DIMMs)	4
Base Configuration PDomain with Half-Populated DCUs	4	• 2 TB (16 GB DIMMs) • 4 TB (32 GB DIMMs)	4

How these resources are divided between LDoms within these PDomains depends on the type of LDom configuration you choose. See "Understanding LDoms" on page 81 for more information.

When determining which PDomain configuration that you want for your system, consider if you need a configuration where one of the LDoms would have one or more of the following characteristics:

- More than eight processors
- More than 4 TB of memory (16 GB DIMMs) or 8 TB of memory (32 GB DIMMs)
- More than four IB HCAs

If you want an LDom with any of the characteristics listed above, then you should choose one of the extended configuration PDomains. If none of the LDoms in your system needs to have any of the characteristics listed above, then you should choose one of the base configuration PDomains.

Related Information

- "Understanding PDomains" on page 61
- "Understanding LDoms" on page 81
- "Understanding High Redundancy" on page 11
- "Allocating CPU Resources for LDoms" on page 14

- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22
- "Understanding Storage for LDoms" on page 25

Understanding High Redundancy

- "High Redundancy Flowchart" on page 11
- "Number of Compute Servers" on page 13
- "Type of Configuration PDomains" on page 13

High Redundancy Flowchart

The following figure shows the decision points for high redundancy in the overall decision-making flowchart.

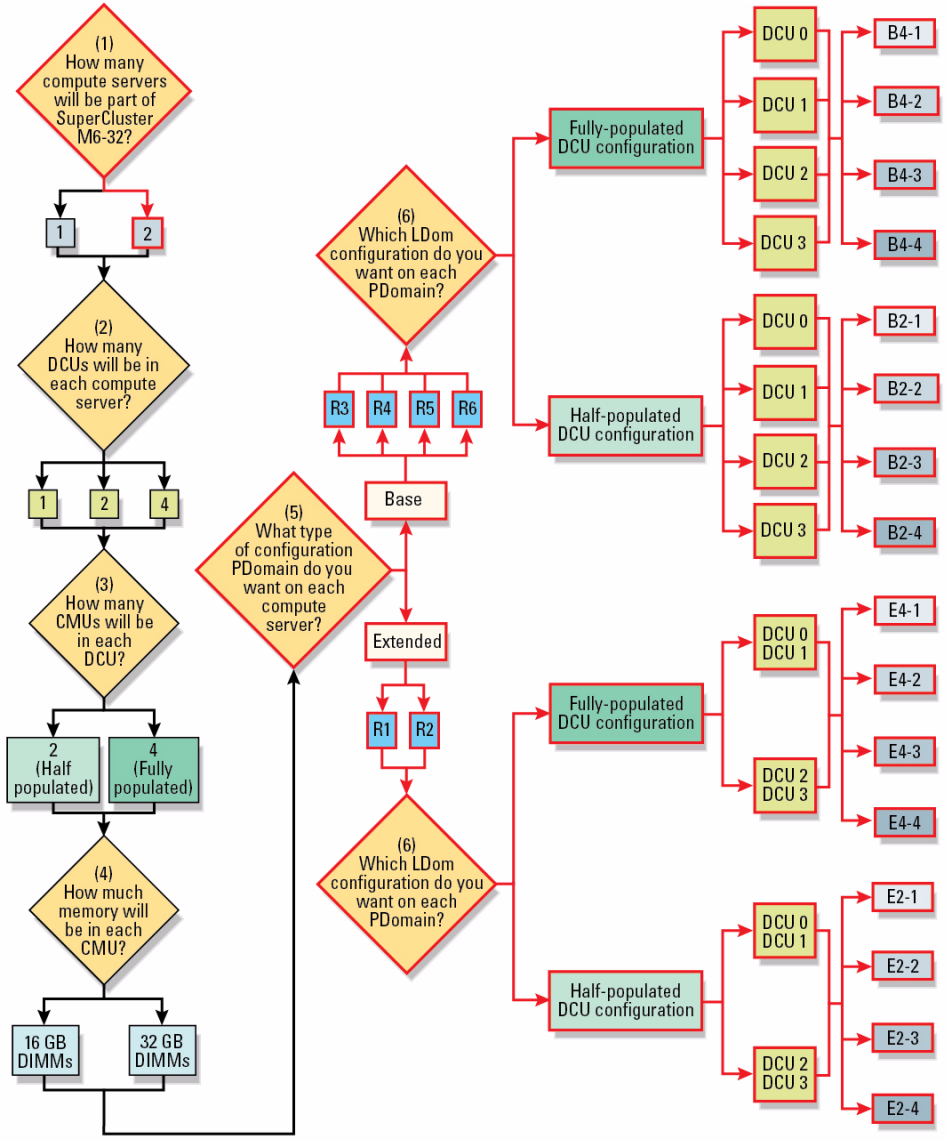

There are two areas where you must make decisions if you want a high-redundancy configuration for your SuperCluster M6-32:

- "Number of Compute Servers" on page 13
- "Type of Configuration PDomains" on page 13

Number of Compute Servers

You can choose either one or two compute servers as part of your SuperCluster M6-32 configuration. If you want high redundancy at the compute server level, you would choose two compute servers. If you have the configuration from one compute server mirrored to the other compute server, you would have optimal high redundancy because if one compute server goes down for any reason, the system would remain available through the second compute server.

Related Information

- "High Redundancy Flowchart" on page 11
- "Type of Configuration PDomains" on page 13

Type of Configuration PDomains

All configuration PDomains contain either two or four DCUs. This allows for redundancy at the PDomain level, where you can have LDom configurations mirrored across the two or four DCUs in each PDomain. See "Understanding PDomains" on page 61 and "Understanding LDoms" on page 81 for more information.

If you choose two compute servers for your SuperCluster M6-32, you would also want to choose a configuration PDomain that takes advantage of the two compute servers that you would be using to provide high redundancy.

The following configuration PDomains are designed specifically for two compute servers:

- "Understanding Four DCUs Across Two Compute Servers (R2 Extended Configuration PDomains)" on page 67
- "Understanding Four DCUs Across Two Compute Servers (R4 Base Configuration PDomains)" on page 74
- "Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains)" on page 78

Note that you have choices on the number of DCUs and the type of configuration PDomain (base or extended) that you would use for the compute servers. Information on the factors you would consider in making those choices are covered in the following sections:

- "Allocating CPU Resources for LDoms" on page 14
- "Allocating Memory Resources for LDoms" on page 17
- "Understanding PCIe Cards and Slots for LDoms" on page 22

- "Understanding Storage for LDoms" on page 25

Related Information

- "High Redundancy Flowchart" on page 11
- "Number of Compute Servers" on page 13

Allocating CPU Resources for LDoms

- "CPU Resources for LDoms Flowchart" on page 14
- "Compute Server Level Considerations" on page 16
- "LDom Level Considerations" on page 16

CPU Resources for LDoms Flowchart

The following figure shows the decision points in the overall decision-making flowchart for choosing the amount of CPU resources that you would want for the LDoms in your SuperCluster M6-32.

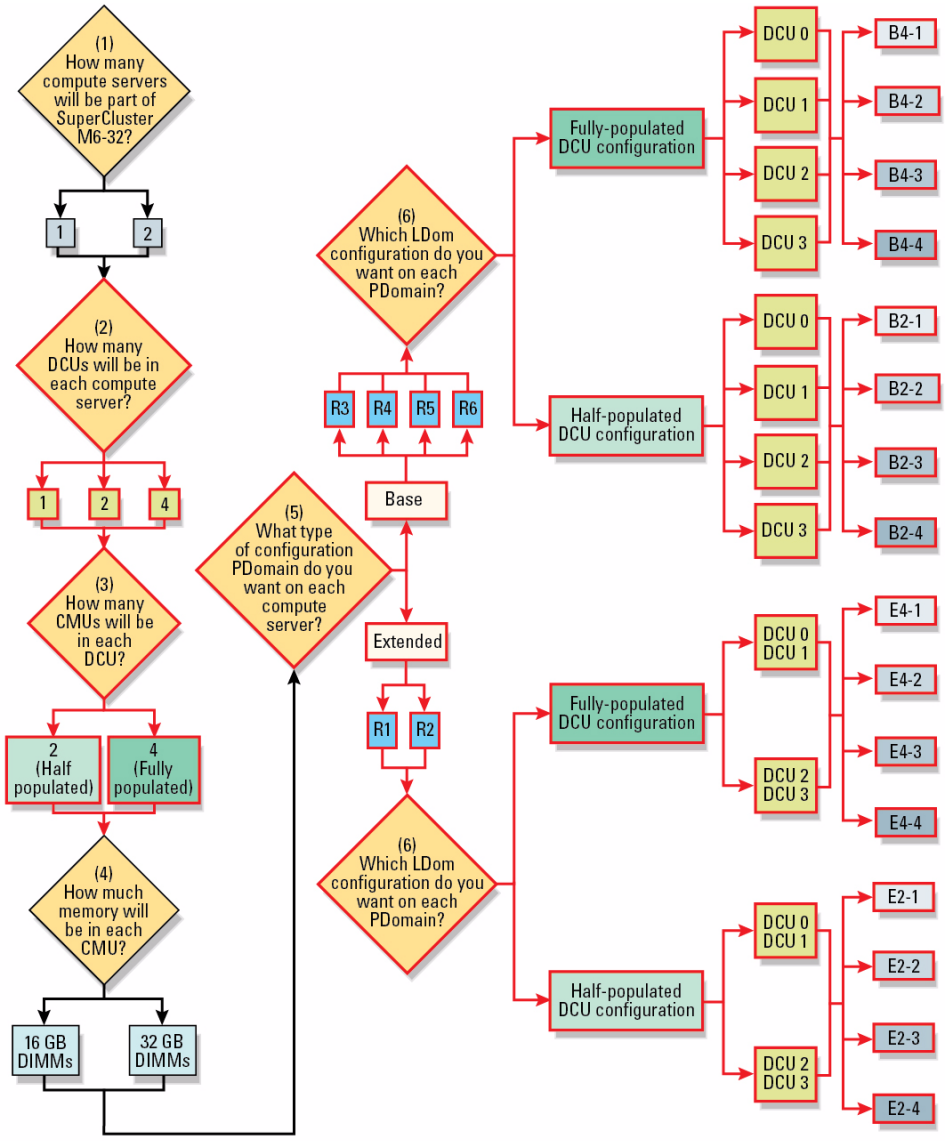

There are four areas that will affect your decisions on how to allocate a specific amount of CPU resources for the LDoms in your SuperCluster M6-32:

- The number of DCUs in each compute server
- The number of CMUs in each DCU
- The type of configuration PDomain that you want on each compute server
- The type of LDom configuration that you want on the PDomain

There are two levels that you need to consider when making the appropriate decision on how to allocate CPU resources:

- "Compute Server Level Considerations" on page 16
- "LDom Level Considerations" on page 16

Compute Server Level Considerations

At the hardware level, you should make a decision on the number of DCUs that would be in each compute server and the number of CMUs that would be in each DCU based on whether you would want future expandability at the DCU or the CMU level.

For example, assume you wanted 8 CMPs worth of CPU resources, where each CMU contains two CMPs. The following table explains the different ways that would allow you to get that amount of CPU resources in a compute server and how each decision would affect future expandability in some way.

No. of DCUs in Compute Server	No. of CMUs in DCUs	Notes
4	2 (half-populated)	• Allows for future expansion of CMUs in the DCUs • Cannot expand on the number of DCUs in each compute server
2	4 (fully-populated)	• Allows for future expansion of DCUs in each compute server • Cannot expand on the number of CMUs in each DCU

Related Information

- "CPU Resources for LDoms Flowchart" on page 14
- "LDom Level Considerations" on page 16
- "Understanding the Hardware" on page 29

LDom Level Considerations

At the LDom level, you should make a decision on the type of configuration PDomain that would be on each compute server (base or extended) and how much CPU resources would be allocated to a specific LDom within that configuration PDomain.

For example, assume that you want 8 CMPs worth of CPU resources to be allocated to a specific LDom. The following table explains how certain factors such as the amount of available storage and PCIe slots would differ depending on the overall configuration PDomain that would be used for this LDom.

Configuration PDomain Type	LDom	No. of EMSs	No. of Used Hard Drives	No. of Available Hard Drives	No. of Used PCIe Slots	No. of Available PCIe Slots
Base	LDom in B4-1	4	2	6	4	11
Extended	LDom in E2-1	8	2	14	8 (4 in each DCU)	22 (11 in each DCU)

Related Information

- "CPU Resources for LDoms Flowchart" on page 14
- "Compute Server Level Considerations" on page 16
- "Understanding PDomains" on page 61
- "Understanding LDoms" on page 81

Allocating Memory Resources for LDoms

- "Memory Resources for LDoms Flowchart" on page 17
- "Compute Server Level Considerations" on page 19
- "LDom Level Considerations" on page 20

Memory Resources for LDoms Flowchart

The following figure shows the decision points in the overall decision-making flowchart for choosing the amount of memory resources that you would want for the LDoms in your SuperCluster M6-32.

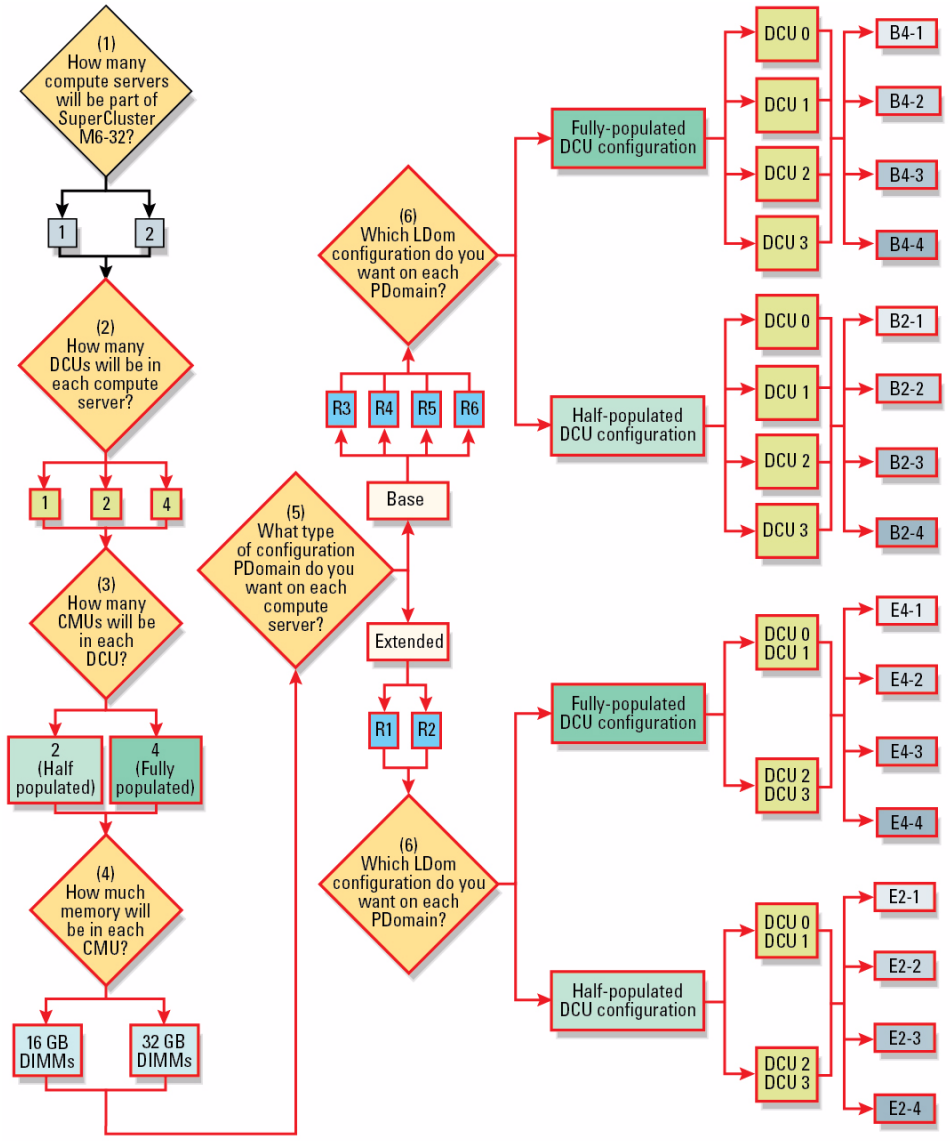

There are five areas that will affect your decisions on how to allocate a specific amount of memory resources for the LDoms in your SuperCluster M6-32:

- The number of DCUs in each compute server
- The number of CMUs in each DCU
- The amount of memory in each CMU (16 GB or 32 GB)
- The type of configuration PDomain that you want on each compute server

- The type of LDom configuration that you want on the PDomain

There are two levels that you need to consider when making the appropriate decision on how to allocate memory resources for LDoms:

- "Compute Server Level Considerations" on page 19
- "LDom Level Considerations" on page 20

Compute Server Level Considerations

At the compute server level, you should make decisions on the following factors based on where you would want future expandability:

- The number of DCUs in each compute server (1, 2 or 4)
- The number of CMUs in each DCU (2 or 4)
- The amount of memory in each CMU (16 GB or 32 GB)

For example, assume you wanted 8192 GB worth of memory resources at the compute server level. The following table explains the different ways that would allow you to get that amount of memory resources in a compute server and how each decision would affect future expandability in some way.

No. of DCUs in Compute Server	No. of CMUs in DCUs	Size of Memory in CMUs	Notes
1	4 (fully-populated)	32 GB	• Allows for future expansion of DCUs in each compute server • Cannot expand on the number of CMUs in each DCU • Cannot increase the size of memory in the CMUs

No. of DCUs in Compute Server	No. of CMUs in DCUs	Size of Memory in CMUs	Notes
2	2 (half-populated)	32 GB	• Allows for future expansion of DCUs in each compute server • Allows for future expansion of CMUs in each DCU • Cannot increase the size of memory in the CMUs
2	4 (fully-populated)	16 GB	• Allows for future expansion of DCUs in each compute server • Allows for future increase in size of memory in the CMUs • Cannot expand on the number of CMUs in each DCU
4	2 (half-populated)	16 GB	• Allows for future expansion of CMUs in each DCU • Allows for future increase in size of memory in the CMUs • Cannot expand on the number of DCUs in each compute server

Related Information

- "Memory Resources for LDoms Flowchart" on page 17
- "LDom Level Considerations" on page 20
- "Understanding the Hardware" on page 29

LDom Level Considerations

At the LDom level, you should make a decision on the type of configuration PDomain that would be on each compute server (base or extended) and how much memory resources would be allocated to a specific LDom within that configuration PDomain.

For example, assume that you want 4096 GB of memory resources to be allocated to a specific LDom. The following table explains how certain factors such as the amount of available storage and PCIe slots would differ depending on the overall configuration PDomain that would be used for this LDom.

Configuration PDomain Type	LDOM Config	# of CMUs in DCU	Size of Memory in CMUs	Calculations	Notes
Base	LDom in B4-1	4 (Fully-Populated)	16 GB	• B4-1 LDom = 4 CMUs (8 CMPs), or 256 DIMM slots • 256 DIMMs slots x 16 GB DIMMs = 4096 GB memory	• All four EMSs • 6 available HDDs • 11 available PCIe slots
Base	LDom in B2-1	2 (Half-Populated)	32 GB	• B2-1 LDom = 2 CMUs (4 CMPs), or 128 DIMM slots • 128 DIMMs slots x 32 GB DIMMs = 4096 GB memory	• All four EMSs • 6 available HDDs • 11 available PCIe slots
Base	Either LDom in B4-2	4 (Fully-Populated)	32 GB	• Either LDom in B4-2 = 2 CMUs (4 CMPs), or 128 DIMM slots • 128 DIMMs slots x 32 GB DIMMs = 4096 GB memory	• Two EMSs • 2 available HDDs • 5 or 6 available PCIe slots
Extended	LDom in E2-1	2 (Half-Populated)	16 GB	• E2-1 LDom = 4 CMUs (8 CMPs), or 256 DIMM slots • 256 DIMMs slots x 16 GB DIMMs = 4096 GB memory	• Eight EMSs (all four from each DCU) • 14 available HDDs • 22 available PCIe slots
Extended	Second LDom in E4-2	4 (Fully-Populated)	32 GB	• Second LDom in B4-2 = 2 CMU (4 CMPs), or 128 DIMM slots • 128 DIMMs slots x 32 GB DIMMs = 4096 GB memory	• Two EMSs • 2 available HDDs • 6 available PCIe slots

Related Information

- "Memory Resources for LDoms Flowchart" on page 17
- "Compute Server Level Considerations" on page 19
- "Understanding PDomains" on page 61
- "Understanding LDoms" on page 81

Understanding PCIe Cards and Slots for LDoms

- "PCIe Cards and Slots for LDoms Flowchart" on page 22
- "Compute Server Level Considerations" on page 24
- "LDom Level Considerations" on page 24

PCIe Cards and Slots for LDoms Flowchart

The following figure shows the decision points in the overall decision-making flowchart for choosing an LDom configuration based on the number of PCIe cards and slots that are available for that LDom.

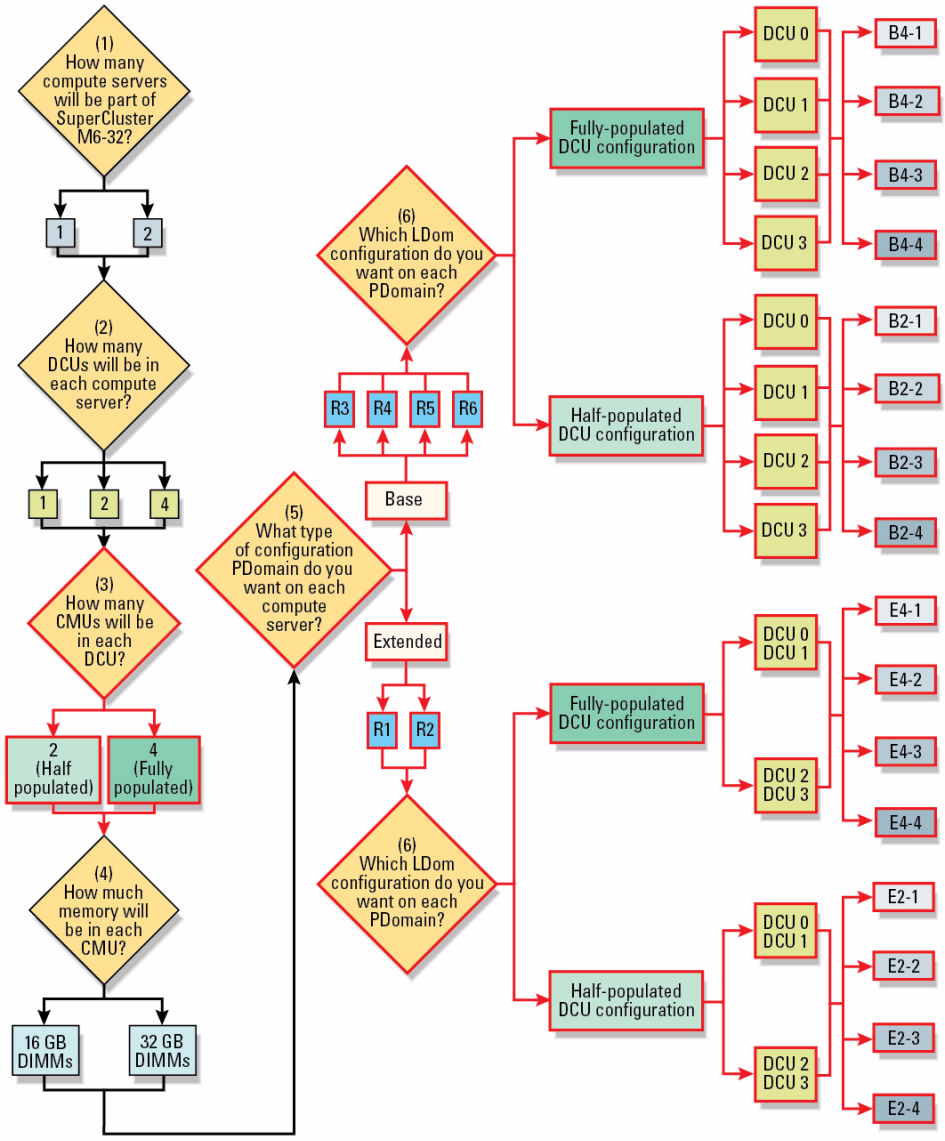

There are three areas that will affect your decision on choosing an LDom configuration based on the number of PCIe cards and slots that are available for that LDom:

- The number of DCUs in each compute server
- The type of configuration PDomain that you want on each compute server
- The type of LDom configuration that you want on the PDomain

There are two levels that you need to consider when making the appropriate decision on an LDom configuration with the number of PCIe cards and slots that are available for that LDom in mind:

- "Compute Server Level Considerations" on page 24
- "LDom Level Considerations" on page 24

Compute Server Level Considerations

Each DCU has sixteen PCIe slots. Four IB HCAs are installed in four of the PCIe slots, and one 1GbE NIC is installed in another PCIe slot. The remaining 11 PCIe slots are empty and can be used for optional Fiber Channel PCIe cards to facilitate migration of data from legacy storage subsystems to the storage servers integrated with SuperCluster M6-32.

The number of PCIe slots that are available for an LDom and the number of cards installed in those PCIe slots will not change based on the number of CMUs that are installed in the DCU. However, the number of DCUs that are installed in a compute server does have an impact, because there are certain configuration PDomains that span across two DCUs (extended configuration PDomains), thereby providing additional PCIe slots and cards for certain LDoms.

Related Information

- "PCIe Cards and Slots for LDoms Flowchart" on page 22
- "LDom Level Considerations" on page 24
- "Understanding Extended Configuration PDomains" on page 64
- "Understanding LDom Configurations for Extended Configuration PDomains" on page 99

LDom Level Considerations

The number of PCIe cards and free slots that are available for any particular LDom varies depending on the type of configuration PDomain and the size of the LDom in that configuration.

For example, the B2-3 base configuration PDomain contains three LDoms, but the number of PCIe cards and slots that are available for each LDom in that configuration varies, as described in "B2-3 LDom Configuration" on page 128.

If you want a greater number of free PCIe slots for optional Fiber Channel PCIe cards, you might chose either LDom 1 in the B2-3 configuration PDomain or a different configuration PDomain altogether that offers even larger LDoms. For

example, you might find that an extended configuration PDomain works best, where certain LDoms span across two DCUs and therefore would have PCIe slots from both DCUs available for that LDom.

As an example, consider the E2-3 extended configuration PDomain, as described in "E2-3 LDom Configuration" on page 112, which contains three LDoms, like the B2-3 base configuration PDomain. The number of PCIe slots for the second and third LDoms are similar to those of the B2-3 base configuration PDomain. However, the number of PCIe slots for the first LDom is considerably larger in the E2-3 extended configuration PDomain compared to the B2-3 base configuration PDomain.

The following table shows the number of IB HCAs and free PCIe slots that are available for the LDoms in each of those configuration PDomains.

Configuration PDomain	LDom 1	LDom 2	LDom 3
B2-3	• 2 IB HCAs • 5 free PCIe slots	• 1 IB HCA • 3 free PCIe slots	• 1 IB HCA • 3 free PCIe slots
E2-3	• 6 IB HCAs • 16 free PCIe slots	• 1 IB HCA • 3 free PCIe slots	• 1 IB HCA • 3 free PCIe slots

Related Information

- "PCIe Cards and Slots for LDoms Flowchart" on page 22
- "Compute Server Level Considerations" on page 24

Understanding Storage for LDoms

- "Storage for LDoms Flowchart" on page 25
- "Compute Server Level Considerations" on page 27
- "LDom Level Considerations" on page 27

Storage for LDoms Flowchart

The following figure shows the decision points in the overall decision-making flowchart for choosing an LDom configuration based on the storage that is available for that LDom.

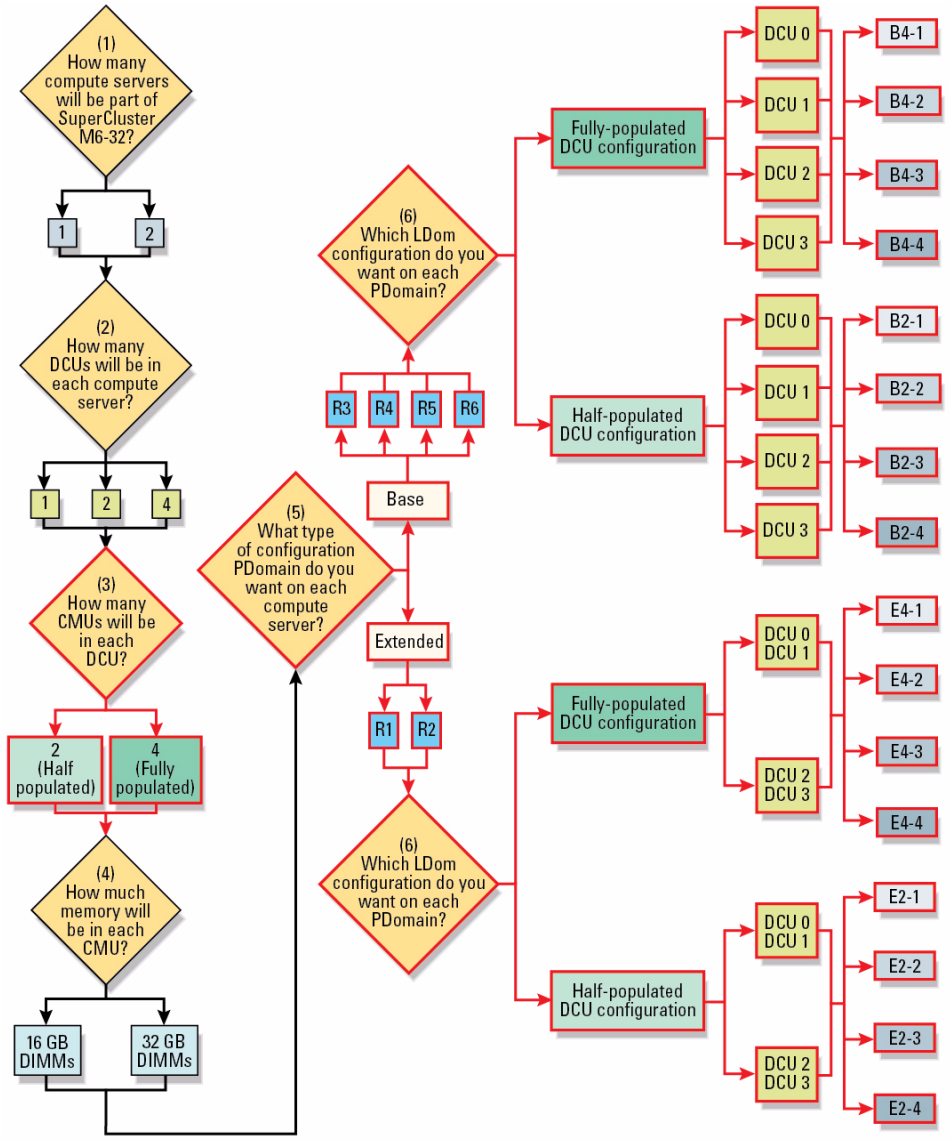

There are three areas that will affect your decision on choosing an LDom configuration based on the storage that is available for that LDom:

- The number of DCUs in each compute server
- The type of configuration PDomain that you want on each compute server
- The type of LDom configuration that you want on the PDomain

There are two levels that you need to consider when making the appropriate decision on an LDom configuration with the storage that is available for that LDom in mind:

- "Compute Server Level Considerations" on page 27
- "LDom Level Considerations" on page 27

Compute Server Level Considerations

Each DCU contains four EMS modules and eight hard drives. Each hard drive is accessed through one EMS module (each EMS module has access to two different hard drives). Every LDom has access to at least one EMS module, and therefore has access to at least two hard disk drives, which are configured as a mirrored pair holding `root`, `swap`, `dump` and `/u01` directories for user files and/or zone root file systems. Additional hard drives are left unconfigured and are available for storage.

The amount of storage that is available for an LDom will not change based on the number of CMUs that are installed in the DCU. However, the number of DCUs that are installed in a compute server does have an impact, because there are certain configuration PDomains that span across two DCUs (extended configuration PDomains), thereby providing additional storage for certain LDoms.

Related Information

- "Storage for LDoms Flowchart" on page 25
- "LDom Level Considerations" on page 27

LDom Level Considerations

The number of hard drives that are available for any particular LDom varies depending on the type of configuration PDomain and the size of the LDom in that configuration.

For example, the B2-3 base configuration PDomain contains three LDoms, but the number of EMS modules and hard disk drives that are available for each LDom in that configuration varies, as described in "B2-3 LDom Configuration" on page 128.

If you want additional storage for an LDom, you might chose either LDom 1 in the B2-3 configuration PDomain or a different configuration PDomain altogether that offers even larger LDoms. For example, you might find that an extended configuration PDomain works best, where certain LDoms span across two DCUs and therefore would have EMS modules and hard drives from both DCUs available for that LDom.

As an example, consider the E2-3 extended configuration PDomain, as described in "E2-3 LDom Configuration" on page 112, which contains three LDoms, like the B2-3 base configuration PDomain. The number of EMS modules and hard drives for the second and third LDoms are similar to those of the B2-3 base configuration PDomain. However, the number of EMS modules and hard drives for the first LDom is considerably larger in the E2-3 extended configuration PDomain compared to the B2-3 base configuration PDomain.

The following table shows the number of EMSs and the number of free hard drives that are available for the LDoms in each of those configuration PDomains (note that all LDoms will always have at least two hard drives configured as a mirrored pair holding the `root`, `swap`, `dump` and `/u01` directories).

Configuration PDomain	LDom 1	LDom 2	LDom 3
B2-3	• 2 EMSs • 2 free hard drives	• 1 EMS • No free hard drives	• 1 EMS • No free hard drives
E2-3	• 6 EMSs • 10 free hard drives	• 1 EMS • No free hard drives	• 1 EMS • No free hard drives

Related Information

- "Storage for LDoms Flowchart" on page 25
- "Compute Server Level Considerations" on page 27

Understanding the Hardware

These topics describe the features and hardware components of Oracle SuperCluster M6-32.

- "Understanding SuperCluster M6-32" on page 29
- "Identifying SuperCluster M6-32 Components" on page 33

Understanding SuperCluster M6-32

- "SuperCluster M6-32 Overview" on page 29
- "Spares Kit Components" on page 31
- "Configuration Restrictions" on page 31

SuperCluster M6-32 Overview

SuperCluster M6-32 is a complete engineered system designed to run databases and applications on a single system. Ideal for consolidation and private cloud, SuperCluster M6-32 can run database, middleware, custom and third party applications. SuperCluster M6-32 is ideal for large scale database and application consolidation and private cloud. You can run a variety of workloads including OLTP and data warehousing, complex applications, and mixed workloads for extreme performance. With big memory, SuperCluster M6-32 can run databases and applications in memory while providing the highest levels of availability and serviceability. SuperCluster M6-32 can scale vertically, allowing customers to flexibly add compute and storage resources to meet demanding data center requirements.

SuperCluster M6-32 is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and widely varied workloads. The SuperCluster M6-32 is intended for large-scale, performance-sensitive, mission-critical application deployments.

Clustering software, such as Oracle RAC and Oracle Solaris Cluster, are optional. Combined with SuperCluster M6-32, it enables a high degree of isolation between concurrently deployed applications, which have varied security, reliability, and performance requirements. SuperCluster M6-32 enables customers to develop a single environment that can support end-to-end consolidation of their entire applications portfolio.

SuperCluster M6-32 provides an optimal solution for all database workloads, ranging from scan-intensive data warehouse applications to highly concurrent OLTP applications. With its combination of smart Oracle Exadata Storage Server Software, complete and intelligent Oracle Database software, and the latest industry-standard hardware components, SuperCluster M6-32 delivers extreme performance in a highly-available, highly-secure environment. Oracle provides unique clustering and workload management capabilities so SuperCluster M6-32 is well-suited for consolidating multiple databases into a single grid. Delivered as a complete pre-optimized, and pre-configured package of software, servers, and storage, SuperCluster M6-32 is fast to implement, and it is ready to tackle your large-scale business applications.

SuperCluster M6-32 does not include any Oracle software licenses. Appropriate licensing of the following software is required when used on SuperCluster M6-32:

- Oracle Database
- Oracle Exadata Storage Server Software
- Oracle Exalogic Elastic Cloud Software

In addition, you should license the following software:

- Oracle Solaris Cluster
- Oracle RACs
- Oracle partitioning

Note – For a full list of software license products, see http://www.oracle.com/supercluster.

SuperCluster M6-32 is designed to fully leverage an internal IB fabric that connects all of the processing, storage, memory, and external network interfaces within SuperCluster M6-32 to form a single, large computing device. Each SuperCluster M6-32 is connected to data center networks through 10GbE (traffic) and 1GbE (management) interfaces.

You can integrate SuperCluster M6-32 with Exadata or Exalogic machines by using the available IB expansion ports and optional data center switches. The IB technology used by SuperCluster M6-32 offers significantly high bandwidth, low latency, hardware-level reliability, and security. If you are using applications that follow Oracle's best practices for highly scalable, fault-tolerant systems, you do not need to make any application architecture or design changes to benefit from SuperCluster

M6-32. You can connect a combination of SuperCluster M6-32 systems and Oracle Exadata Database Machines, to develop a single, large-scale environment. You can integrate SuperCluster M6-32 systems with their current data center infrastructure using the available 10GbE ports in each compute server.

Related Information

- "Spares Kit Components" on page 31
- "Configuration Restrictions" on page 31
- "Identifying SuperCluster M6-32 Components" on page 33

Spares Kit Components

SuperCluster M6-32 includes a spares kit that includes the following components:

- Either of the following sets of disks and cards:
  - 4 TB high-capacity disk and 1.6 TB Exadata Smart Flash Cache card, or
  - 1.6 TB NVMe PCI flash drive
- IB cables

Related Information

- "SuperCluster M6-32 Overview" on page 29
- "Configuration Restrictions" on page 31
- "Identifying SuperCluster M6-32 Components" on page 33

Configuration Restrictions

The following restrictions apply to hardware and software modifications to SuperCluster M6-32. Violating these restrictions can result in loss of warranty and support.

- SuperCluster M6-32 hardware cannot be modified or customized. There is one exception to this. The only allowed hardware modification to SuperCluster M6-32 is to the administrative Ethernet management switch included with SuperCluster M6-32. You may choose to do the following:
  - Replace the Ethernet management switch, at your expense, with an equivalent 1U 48-port Ethernet management switch that conforms to your internal data center network standards. You must perform this replacement, at your expense and labor, after delivery of SuperCluster M6-32. If you choose to make this change, then Oracle cannot make or assist with this change given the numerous

possible scenarios involved, and it is not included as part of the standard installation. You 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 your network through an external switch or patch panel. You must perform these changes at your expense and labor. In this case, the Ethernet management switch in the rack can be turned off and disconnected from the data center network.

- The storage rack, containing nine storage servers and the Oracle ZFS Storage ZS3-ES storage appliance (ZFS storage appliance), is a required component for SuperCluster M6-32. For additional storage, up to 17 optional expansion racks, either Full, Half or Quarter Racks, can be added to SuperCluster M6-32.
- The optional expansion rack can only be connected to SuperCluster M6-32 or Oracle Exadata Database Machine, and only supports databases running on the Database Domains in SuperCluster M6-32 or on the database servers in Oracle Exadata Database Machine.
- Standalone storage servers can only be connected to SuperCluster M6-32 or Oracle Exadata Database Machine, and only support databases running on the Database Domains in SuperCluster M6-32 or on the database servers in Oracle Exadata Database Machine. The standalone storage servers must be installed in a separate rack.
- Earlier Oracle DB releases can be run in Application Domains running Oracle Solaris 10. Non-Oracle DBs can be run in either Application Domains running Oracle Solaris 10 or Oracle Solaris 11, depending on the Oracle Solaris version they support.
- Oracle Exadata Storage Server Software and the operating systems cannot be modified, and you cannot install any additional software or agents on the storage servers.
- You cannot update the firmware directly on the storage servers. The firmware is updated as part of a storage server patch.
- You 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 nonstandard kernel modules to the OS of the Database Domains is allowed but discouraged. Oracle does not support questions or issues with the nonstandard modules. If a server crashes, and Oracle suspects the crash may have been caused by a nonstandard module, then Oracle support may refer the customer to the vendor of the nonstandard module or ask that the issue be reproduced without the nonstandard module. Modifying the 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 M6-32 supports separate domains dedicated to applications, with high throughput/low latency access to the database domains through IB. Since Oracle DB 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.
- You 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 servers can be used to connect to external storage servers using iSCSI or NFS. However, the FCoE protocol is not supported.
- Only switches specified for use in SuperCluster M6-32, Oracle Exadata Rack, and Oracle Exalogic Elastic Cloud may be connected to the IB network. Connecting third-party switches and other switches not used in SuperCluster M6-32, Oracle Exadata Rack, and Oracle Exalogic Elastic Cloud is not supported.

Related Information

- "SuperCluster M6-32 Overview" on page 29
- "Spares Kit Components" on page 31
- "Identifying SuperCluster M6-32 Components" on page 33

Identifying SuperCluster M6-32 Components

SuperCluster M6-32 consists of the following components:

- One or two compute servers
- A storage rack, containing nine storage servers and one ZFS storage appliance

The following topics describe SuperCluster M6-32:

- "Understanding the Compute Server" on page 33
- "Storage Rack Overview" on page 38

Understanding the Compute Server

The following topics describe the compute server:

- "Compute Server Overview" on page 34
- "Front Components" on page 35
- "Rear Components" on page 37

Compute Server Overview

The compute server is an enterprise-class server designed for mission-critical applications.

You can have one or two compute servers as part of SuperCluster M6-32. In addition to the storage rack, you can also connect up to 17 optional expansion racks on the same IB fabric without the need for any external switches. Refer to the *Oracle SuperCluster M6-32 Owner's Guide: Installation* for more information.

Feature	Description
Processor	8 to 32 SPARC M6 processors, each with 12 cores and up to 8 threads per core
Memory	32 DIMM slots per processor
I/O expansion	64 low-profile PCIe Generation 3 card slots

Feature	Description
Hard drive storage	32 drive slots supporting 2.5-inch HDDs or SSDs
Service processors	2 redundant SPs with 4 SPPs to monitor and control the server remotely

Related Information

- "Front Components" on page 35
- "Rear Components" on page 37

Front Components

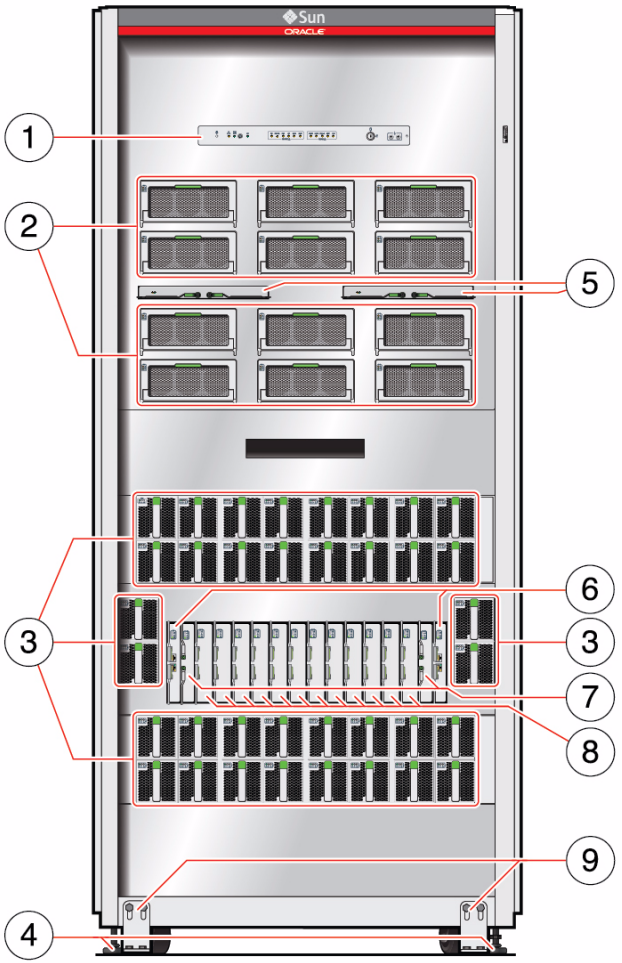

No.	Description
1	Front LED panel, containing LEDs, key switch, and ESD grounding jacks
2	PS
3	Fan module
4	Leveling feet
5	PSDB
6	SPs
7	Clock board
8	SSBs
9	Mounting brackets

Related Information

- "Rear Components" on page 37

Rear Components

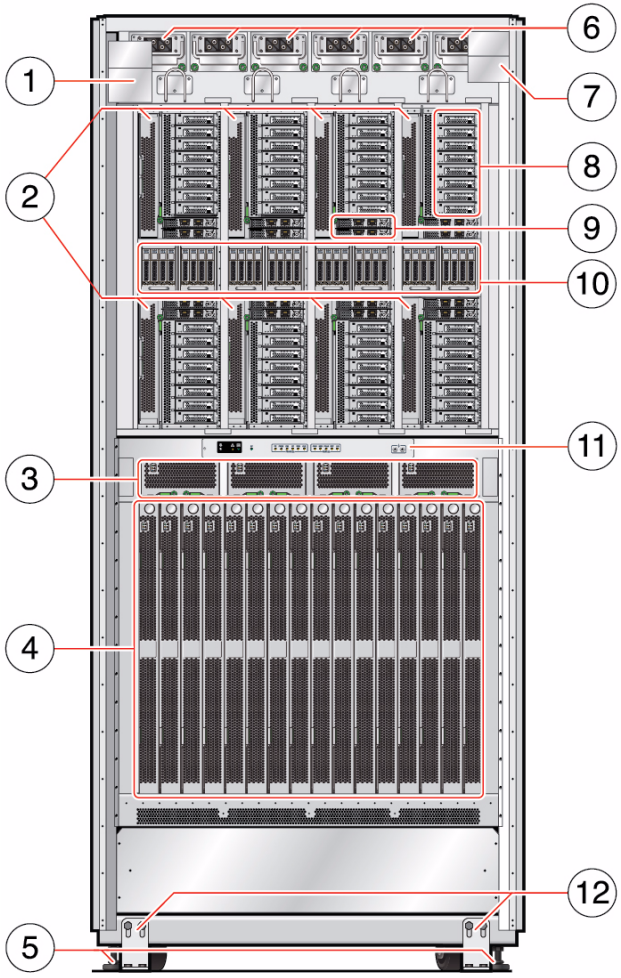

No.	Description
1	Left cable management bracket
2	I/O boards
3	SPPs
4	CMUs
5	Leveling feet
6	AC input filters (power cord connectors)

No.	Description
7	Right cable management bracket
8	PCIe hot-plug carrier for low-profile PCIe cards
9	EMSs
10	Hard drives
11	Rear LED board, containing LEDs and ESD grounding jacks
12	Mounting brackets

Related Information

- "Front Components" on page 35

Storage Rack Overview

The storage rack, containing nine storage servers and one ZFS storage appliance, is a required component to SuperCluster M6-32. In addition, up to 17 optional expansion racks, either full, half or quarter racks, can be added to SuperCluster M6-32.

For more information on the storage rack, refer to the *Oracle SuperCluster M6-32 Owner's Guide: Installation*.

Related Information

- "Identifying SuperCluster M6-32 Components" on page 33

Understanding DCUs

These topics describe DCUs and DCU configurations.

Note – Certain DCU configurations have a mixture of half-populated and fully-populated DCUs in a server. For extended configuration PDomains in those DCU configurations, both DCUs in each PDomain must have the same configuration, where both DCUs in the PDomain are either half-populated or fully-populated. See "CMU Overview" on page 5 for more information.

- "DCU Overview" on page 39
- "Understanding DCU Configurations" on page 42
- "Understanding Half-Populated DCU Root Complexes" on page 48
- "Understanding Fully-Populated DCU Root Complexes" on page 54

DCU Overview

Each server has up to four DCUs, which are the building blocks of PDomains. Each DCU has two or four CMUs that have two memory boards and two CMPs. Each DCU is managed by its own SPP, which monitors environmental sensors and manages the CMUs, memory controller, and DIMMs within the DCU. The Active-SP and SPPs communicate over a private VLAN to manage the system.

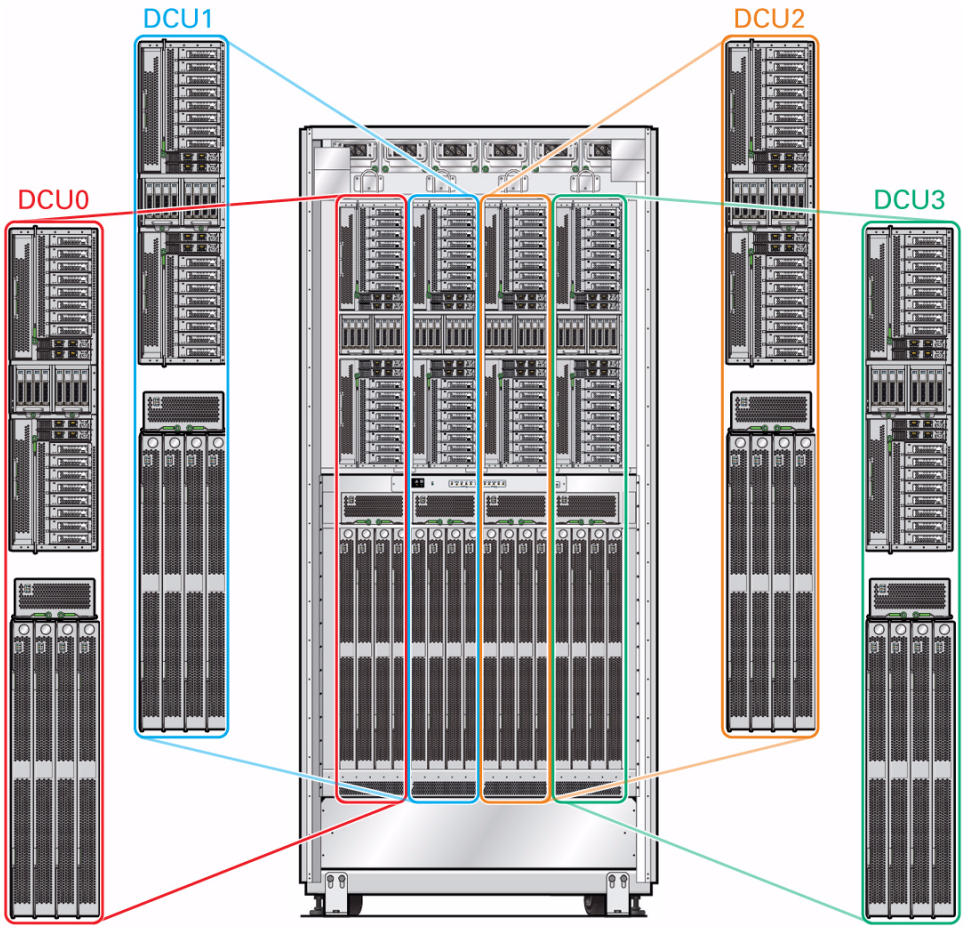

The illustration displays the individual components of a single DCU.

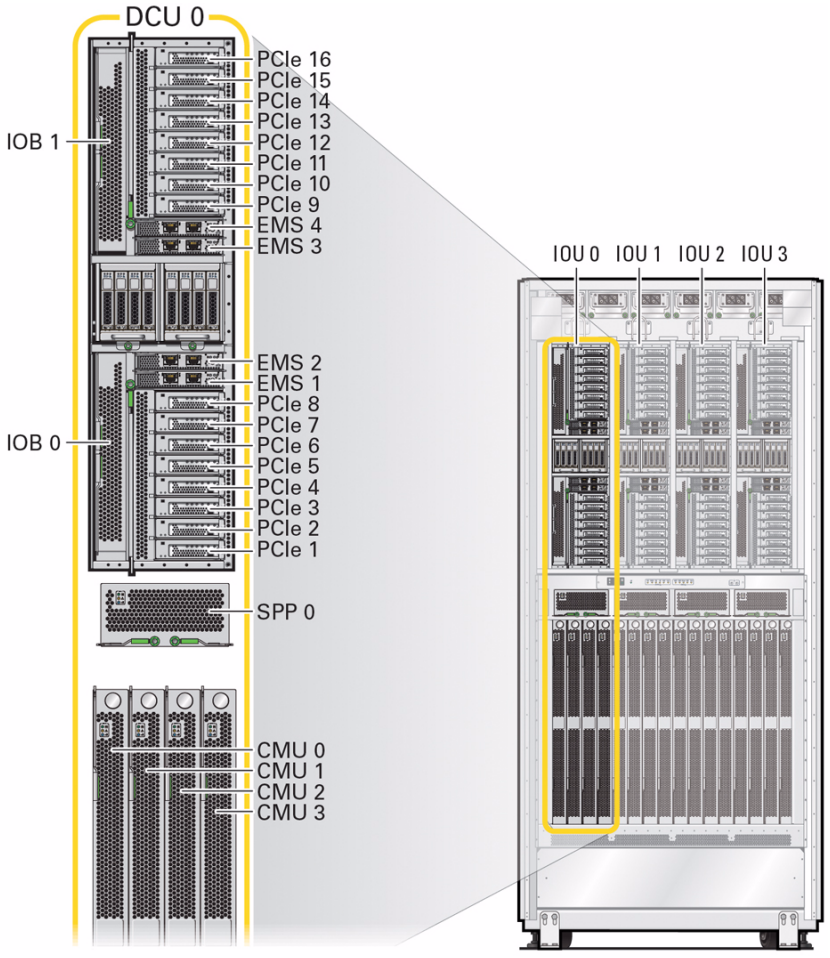

Each DCU contains the following hardware components:

- One IOU, which contains:
  - Sixteen PCIe expansion slots, populated with:
    - One quad-port GbE NIC for the management network
    - Four dual-port IB HCAs
  - Four EMS cards, each with two 10 GbE ports
  - Eight hard drives
- One SPP
- Two or four CPU memory units (see "CMU Overview" on page 5 for more information)

For SuperCluster M6-32, you can have either one DCU or two DCUs associated with each PDomain, depending on the type of configuration you choose. See "Understanding PDomains" on page 61 for more information.

Related Information

- "Rear Components" on page 37
- "Understanding PDomains" on page 61

Understanding DCU Configurations

Once you determine which of the PDomain configurations that you want, you must then choose the number of CMUs that are in the DCUs that are in your configuration:

- Half-populated DCU configuration – Two CMUs in each DCU
- Fully-populated DCU configuration – Four CMUs in each DCU

The following topics describe those configurations in more detail:

- "PCIe Device Root Complexes Overview" on page 42
- "PCIe Communication and Paths Overview" on page 43
- "Understanding DCU PCIe and EMS Slot Locations" on page 43

Related Information

- "Understanding Half-Populated DCU Root Complexes" on page 48
- "Understanding Fully-Populated DCU Root Complexes" on page 54

PCIe Device Root Complexes Overview

A root complex is the 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. Understanding the relationship of the PCIe root complexes to the PCIe I/O fabrics will help you properly assign devices when configuring Oracle VM Server for SPARC LDoms.

There are 64 root complexes in the compute server (16 per DCU).

Related Information

- "PCIe Communication and Paths Overview" on page 43
- "Understanding DCU PCIe and EMS Slot Locations" on page 43

PCIe Communication and Paths Overview

PCIe I/O paths are routed from CMP root complexes, through PCIe switches, to a destination PCIe slot and/or EMS card. Each CMP provides two root complexes, allowing it to attach to two PCIe switches. Each PCIe switch connects to two root complexes. In an ideal configuration, with all components present and functioning, two CMPs connect to each PCIe switch, and that switch is partitioned into two virtual switches with independent traffic channels. This configuration maximizes I/O bandwidth.

In the event of a failure or missing CMPs (and associated root complexes), the PCIe switch is merged into a single switch so that the remaining root complex services all downstream ports. This configuration maximizes I/O connectivity.

As long as there are no changes in the hardware components present in the configuration, the I/O paths will remain consistent following a boot or reset. If there are changes in the number of CMPs present (for example, if CMPs are added or removed), some I/O paths might change.

Related Information

- "PCIe Device Root Complexes Overview" on page 42
- "Understanding DCU PCIe and EMS Slot Locations" on page 43

Understanding DCU PCIe and EMS Slot Locations

The following topics describe the PCIe and EMS slot locations for the DCUs in a server:

- "DCU 0 PCIe and EMS Slot Locations" on page 44
- "DCU 1 PCIe and EMS Slot Locations" on page 45
- "DCU 2 PCIe and EMS Slot Locations" on page 46
- "DCU 3 PCIe and EMS Slot Locations" on page 47

DCU 0 PCIe and EMS Slot Locations

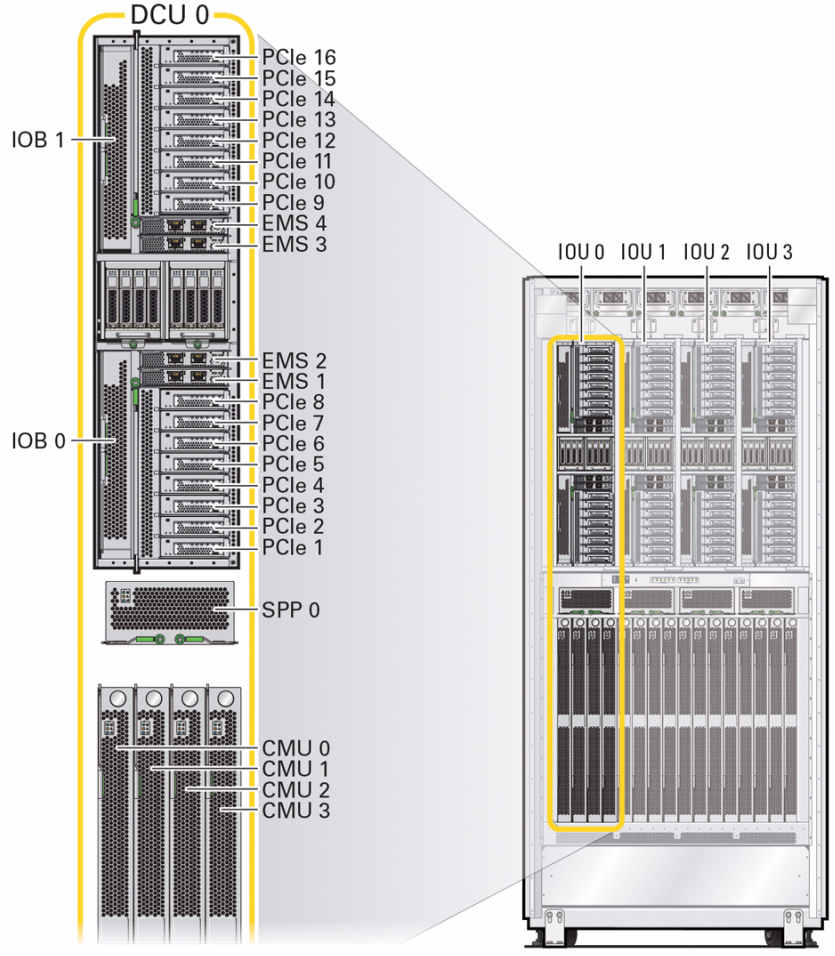

Related Information

- "DCU 1 PCIe and EMS Slot Locations" on page 45
- "DCU 2 PCIe and EMS Slot Locations" on page 46
- "DCU 3 PCIe and EMS Slot Locations" on page 47

DCU 1 PCIe and EMS Slot Locations

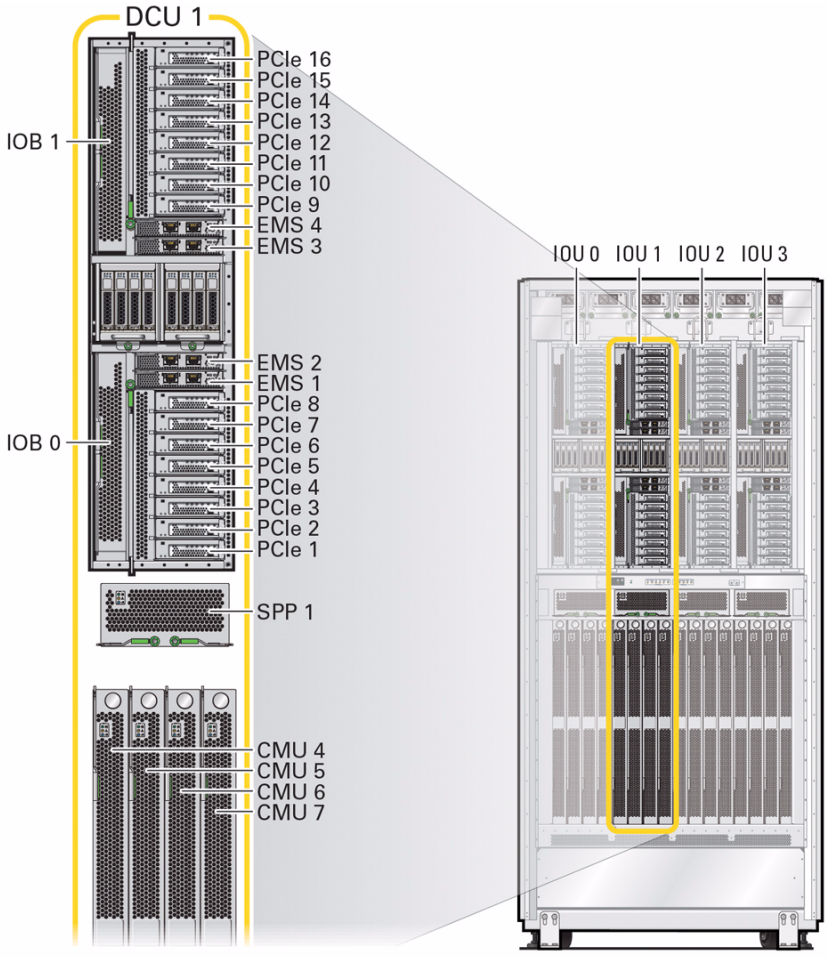

Related Information

- "DCU 0 PCIe and EMS Slot Locations" on page 44
- "DCU 2 PCIe and EMS Slot Locations" on page 46
- "DCU 3 PCIe and EMS Slot Locations" on page 47

DCU 2 PCIe and EMS Slot Locations

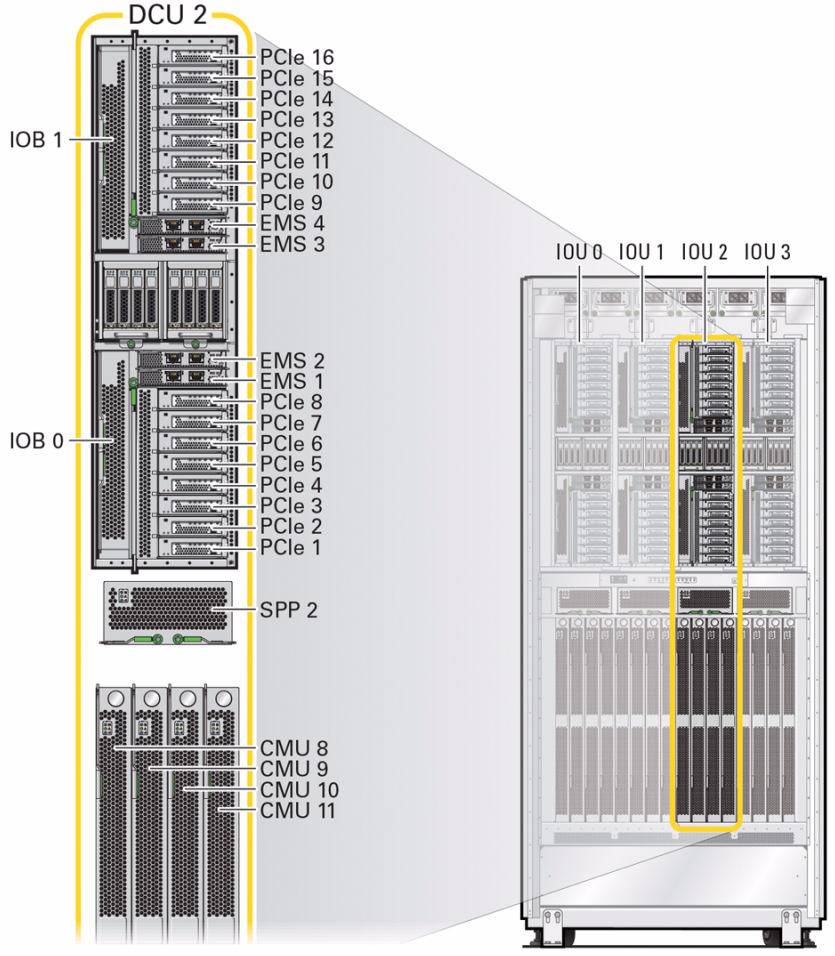

Related Information

- “DCU 0 PCIe and EMS Slot Locations” on page 44
- “DCU 1 PCIe and EMS Slot Locations” on page 45
- “DCU 3 PCIe and EMS Slot Locations” on page 47

DCU 3 PCIe and EMS Slot Locations

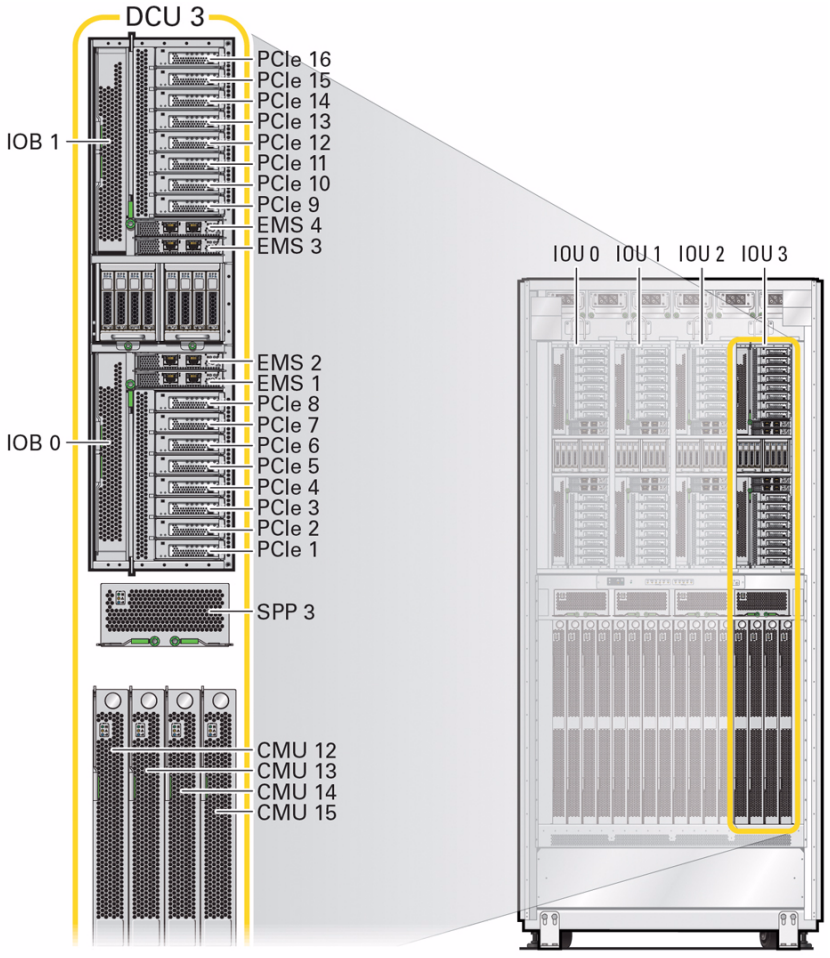

Related Information

- “DCU 0 PCIe and EMS Slot Locations” on page 44
- “DCU 1 PCIe and EMS Slot Locations” on page 45
- “DCU 2 PCIe and EMS Slot Locations” on page 46

Understanding Half-Populated DCU Root Complexes

In a half-populated DCU configuration, each DCU contains two CMUs. You can have from one to four DCUs in each compute server.

The following topics list the root complex names and Oracle Solaris OS device paths for the PCIe and EMS slots in each DCU in a half-populated DCU configuration:

- "Half-Populated DCU 0 PCIe Slot Root Complexes" on page 48
- "Half-Populated DCU 1 PCIe Slot Root Complexes" on page 49
- "Half-Populated DCU 2 PCIe Slot Root Complexes" on page 51
- "Half-Populated DCU 3 PCIe Slot Root Complexes" on page 52

Half-Populated DCU 0 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric path from each root complex to the PCIe and EMS slots in a half-populated DCU 0. A half-populated DCU 0 contains only CMU0 and CMU3.

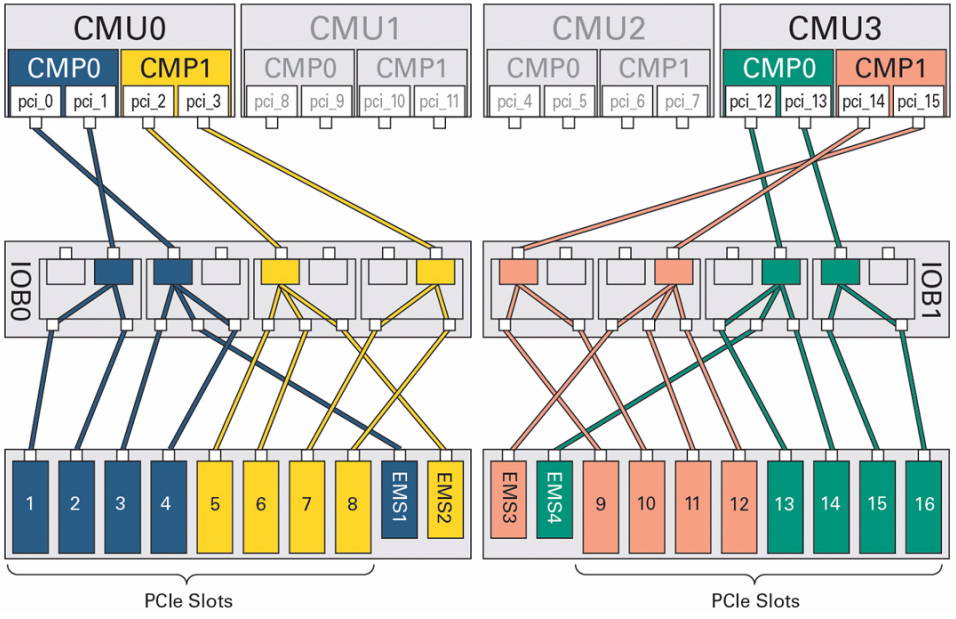

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_1	CMU0/CMP0	/pci@340/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE1
PCIe 2	pci_1	CMU0/CMP0	/pci@340/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE2
PCIe 3	pci_0	CMU0/CMP0	/pci@300/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE3
PCIe 4	pci_0	CMU0/CMP0	/pci@300/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE4
PCIe 5	pci_2	CMU0/CMP1	/pci@380/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE5
PCIe 6	pci_2	CMU0/CMP1	/pci@380/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE6
PCIe 7	pci_3	CMU0/CMP1	/pci@3c0/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE7
PCIe 8	pci_3	CMU0/CMP1	/pci@3c0/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE8
PCIe 9	pci_15	CMU3/CMP1	/pci@6c0/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE9
PCIe 10	pci_15	CMU3/CMP1	/pci@6c0/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE10
PCIe 11	pci_14	CMU3/CMP1	/pci@680/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE11
PCIe 12	pci_14	CMU3/CMP1	/pci@680/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE12
PCIe 13	pci_12	CMU3/CMP0	/pci@600/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE13
PCIe 14	pci_12	CMU3/CMP0	/pci@600/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE14
PCIe 15	pci_13	CMU3/CMP0	/pci@640/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE15
PCIe 16	pci_13	CMU3/CMP0	/pci@640/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE16
EMS1	pci_0	CMU0/CMP0	/pci@300/pci@1/pci@0/pci@c	/SYS/IOU0/EMS1
EMS2	pci_2	CMU0/CMP1	/pci@380/pci@1/pci@0/pci@0	/SYS/IOU0/EMS2
EMS3	pci_14	CMU3/CMP1	/pci@680/pci@1/pci@0/pci@c	/SYS/IOU0/EMS3
EMS4	pci_12	CMU3/CMP0	/pci@600/pci@1/pci@0/pci@0	/SYS/IOU0/EMS4

Related Information

- "Half-Populated DCU 1 PCIe Slot Root Complexes" on page 49
- "Half-Populated DCU 2 PCIe Slot Root Complexes" on page 51
- "Half-Populated DCU 3 PCIe Slot Root Complexes" on page 52

Half-Populated DCU 1 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric path from each root complex to the PCIe and EMS slots in a half-populated DCU 1. A half-populated DCU 1 contains only CMU4 and CMU7.

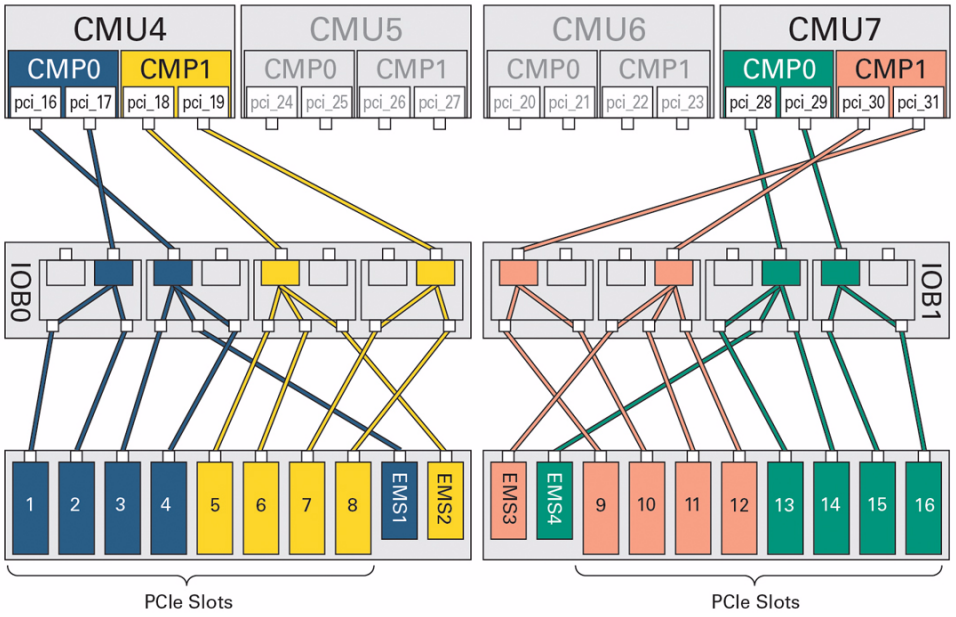

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_17	CMU4/CMP0	/pci@740/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE1
PCIe 2	pci_17	CMU4/CMP0	/pci@740/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE2
PCIe 3	pci_16	CMU4/CMP0	/pci@700/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE3
PCIe 4	pci_16	CMU4/CMP0	/pci@700/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE4
PCIe 5	pci_18	CMU4/CMP1	/pci@780/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE5
PCIe 6	pci_18	CMU4/CMP1	/pci@780/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE6
PCIe 7	pci_19	CMU4/CMP1	/pci@7c0/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE7
PCIe 8	pci_19	CMU4/CMP1	/pci@7c0/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE8
PCIe 9	pci_31	CMU7/CMP1	/pci@ac0/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE9
PCIe 10	pci_31	CMU7/CMP1	/pci@ac0/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE10
PCIe 11	pci_30	CMU7/CMP1	/pci@a80/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE11
PCIe 12	pci_30	CMU7/CMP1	/pci@a80/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE12
PCIe 13	pci_28	CMU7/CMP0	/pci@a00/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE13
PCIe 14	pci_28	CMU7/CMP0	/pci@a00/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE14
PCIe 15	pci_29	CMU7/CMP0	/pci@a40/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE15
PCIe 16	pci_29	CMU7/CMP0	/pci@a40/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE16

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
EMS1	pci_16	CMU4/CMP0	/pci@700/pci@1/pci@0/pci@c	/SYS/IOU1/EMS1
EMS2	pci_18	CMU4/CMP1	/pci@780/pci@1/pci@0/pci@0	/SYS/IOU1/EMS2
EMS3	pci_30	CMU7/CMP1	/pci@a80/pci@1/pci@0/pci@c	/SYS/IOU1/EMS3
EMS4	pci_28	CMU7/CMP0	/pci@a00/pci@1/pci@0/pci@0	/SYS/IOU1/EMS4

Related Information

- "Half-Populated DCU 0 PCIe Slot Root Complexes" on page 48
- "Half-Populated DCU 2 PCIe Slot Root Complexes" on page 51
- "Half-Populated DCU 3 PCIe Slot Root Complexes" on page 52

Half-Populated DCU 2 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a half-populated DCU 2. A half-populated DCU 2 contains only CMU8 and CMU11.

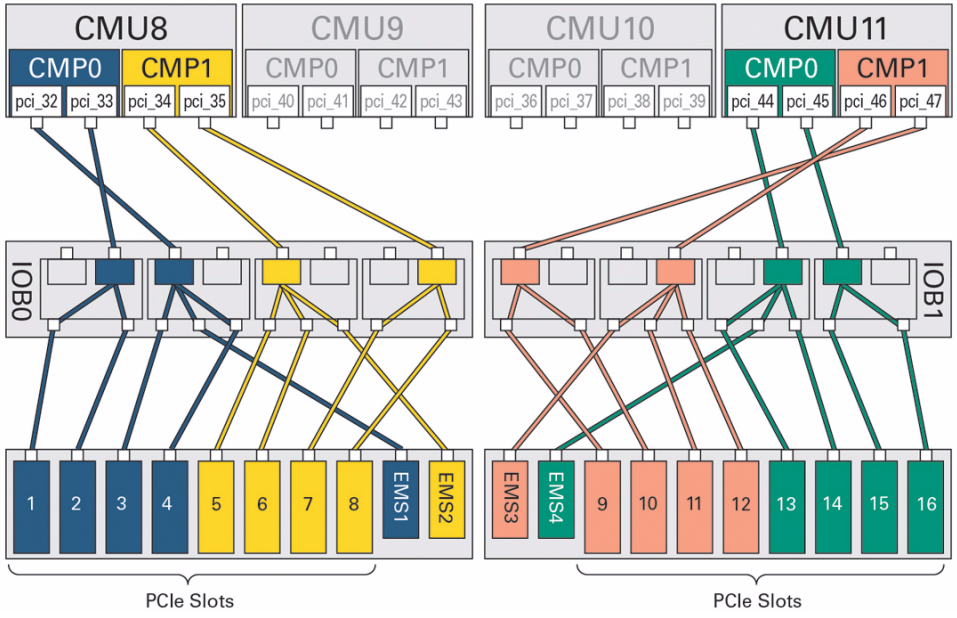

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_33	CMU8/CMP0	/pci@b40/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE1
PCIe 2	pci_33	CMU8/CMP0	/pci@b40/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE2
PCIe 3	pci_32	CMU8/CMP0	/pci@b00/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE3
PCIe 4	pci_32	CMU8/CMP0	/pci@b00/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE4
PCIe 5	pci_34	CMU8/CMP1	/pci@b80/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE5
PCIe 6	pci_34	CMU8/CMP1	/pci@b80/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE6
PCIe 7	pci_35	CMU8/CMP1	/pci@bc0/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE7
PCIe 8	pci_35	CMU8/CMP1	/pci@bc0/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE8
PCIe 9	pci_47	CMU11/CMP1	/pci@ec0/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE9
PCIe 10	pci_47	CMU11/CMP1	/pci@ec0/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE10
PCIe 11	pci_46	CMU11/CMP1	/pci@e80/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE11
PCIe 12	pci_46	CMU11/CMP1	/pci@e80/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE12
PCIe 13	pci_44	CMU11/CMP0	/pci@e00/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE13
PCIe 14	pci_44	CMU11/CMP0	/pci@e00/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE14
PCIe 15	pci_45	CMU11/CMP0	/pci@e40/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE15
PCIe 16	pci_45	CMU11/CMP0	/pci@e40/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE16
EMS1	pci_32	CMU8/CMP0	/pci@b00/pci@1/pci@0/pci@c	/SYS/IOU2/EMS1
EMS2	pci_34	CMU8/CMP1	/pci@b80/pci@1/pci@0/pci@0	/SYS/IOU2/EMS2
EMS3	pci_46	CMU11/CMP1	/pci@e80/pci@1/pci@0/pci@c	/SYS/IOU2/EMS3
EMS4	pci_44	CMU11/CMP0	/pci@e00/pci@1/pci@0/pci@0	/SYS/IOU2/EMS4

Related Information

- "Half-Populated DCU 0 PCIe Slot Root Complexes" on page 48
- "Half-Populated DCU 1 PCIe Slot Root Complexes" on page 49
- "Half-Populated DCU 3 PCIe Slot Root Complexes" on page 52

Half-Populated DCU 3 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a half-populated DCU 3. A half-populated DCU 3 contains only CMU12 and CMU15.

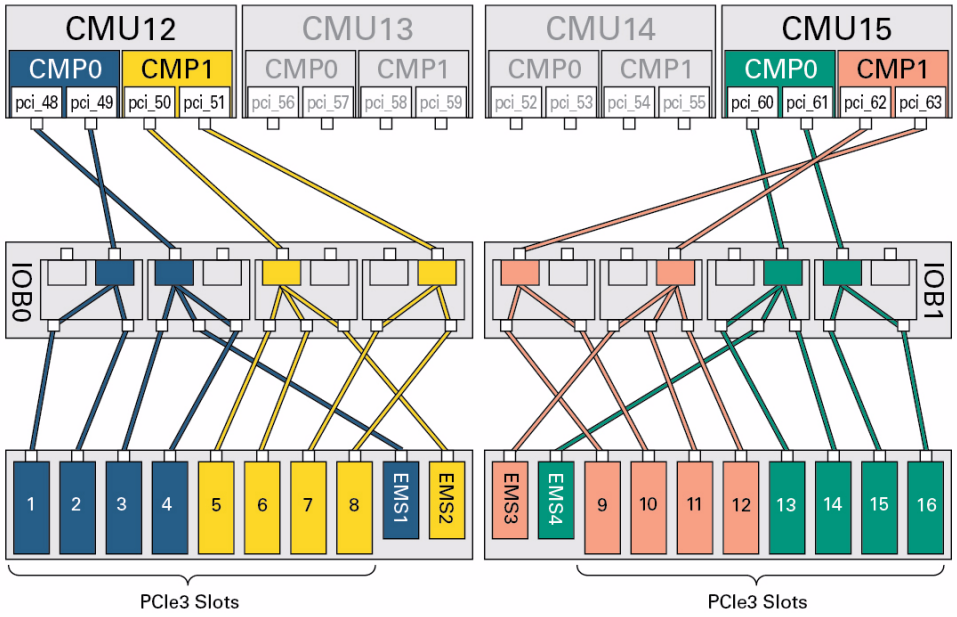

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_49	CMU12/CMP0	/pci@f40/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE1
PCIe 2	pci_49	CMU12/CMP0	/pci@f40/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE2
PCIe 3	pci_48	CMU12/CMP0	/pci@f00/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE3
PCIe 4	pci_48	CMU12/CMP0	/pci@f00/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE4
PCIe 5	pci_50	CMU12/CMP1	/pci@f80/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE5
PCIe 6	pci_50	CMU12/CMP1	/pci@f80/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE6
PCIe 7	pci_51	CMU12/CMP1	/pci@fc0/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE7
PCIe 8	pci_51	CMU12/CMP1	/pci@fc0/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE8
PCIe 9	pci_63	CMU15/CMP1	/pci@12c0/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE9
PCIe 10	pci_63	CMU15/CMP1	/pci@12c0/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE10
PCIe 11	pci_62	CMU15/CMP1	/pci@1280/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE11
PCIe 12	pci_62	CMU15/CMP1	/pci@1280/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE12
PCIe 13	pci_60	CMU15/CMP0	/pci@1200/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE13
PCIe 14	pci_60	CMU15/CMP0	/pci@1200/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE14
PCIe 15	pci_61	CMU15/CMP0	/pci@1240/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE15
PCIe 16	pci_61	CMU15/CMP0	/pci@1240/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE16

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
EMS1	pci_48	CMU12/CMP0	/pci@f00/pci@1/pci@0/pci@c	/SYS/IOU3/EMS1
EMS2	pci_50	CMU12/CMP1	/pci@f80/pci@1/pci@0/pci@0	/SYS/IOU3/EMS2
EMS3	pci_62	CMU15/CMP1	/pci@1280/pci@1/pci@0/pci@c	/SYS/IOU3/EMS3
EMS4	pci_60	CMU15/CMP0	/pci@1200/pci@1/pci@0/pci@0	/SYS/IOU3/EMS4

Related Information

- "Half-Populated DCU 0 PCIe Slot Root Complexes" on page 48
- "Half-Populated DCU 1 PCIe Slot Root Complexes" on page 49
- "Half-Populated DCU 2 PCIe Slot Root Complexes" on page 51

Understanding Fully-Populated DCU Root Complexes

In a fully-populated DCU configuration, each DCU contains four CMUs. You can have from one to four DCUs in each compute server.

The following topics list the root complex names and Oracle Solaris OS device paths for the PCIe and EMS slots in each DCU in a fully-populated DCU configuration:

- "Fully-Populated DCU 0 PCIe Slot Root Complexes" on page 54
- "Fully-Populated DCU 1 PCIe Slot Root Complexes" on page 56
- "Fully-Populated DCU 2 PCIe Slot Root Complexes" on page 57
- "Fully-Populated DCU 3 PCIe Slot Root Complexes" on page 59

Fully-Populated DCU 0 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a fully populated DCU 0.

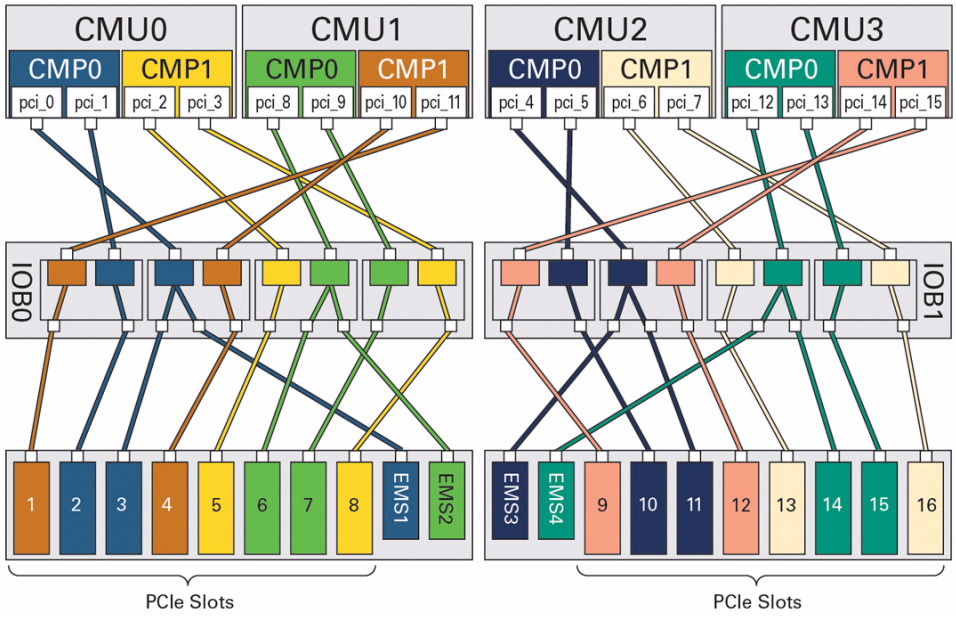

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_11	CMU1/CMP1	/pci@5c0/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE1
PCIe 2	pci_1	CMU0/CMP0	/pci@340/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE2
PCIe 3	pci_0	CMU0/CMP0	/pci@300/pci@1/pci@0/pci@6	/SYS/IOU0/PCIe3
PCIe 4	pci_10	CMU1/CMP1	/pci@580/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE4
PCIe 5	pci_2	CMU0/CMP1	/pci@380/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE5
PCIe 6	pci_8	CMU1/CMP0	/pci@500/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE6
PCIe 7	pci_9	CMU1/CMP0	/pci@540/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE7
PCIe 8	pci_3	CMU0/CMP1	/pci@3c0/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE8
PCIe 9	pci_15	CMU3/CMP1	/pci@6c0/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE9
PCIe 10	pci_5	CMU2/CMP0	/pci@440/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE10
PCIe 11	pci_4	CMU2/CMP0	/pci@400/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE11
PCIe 12	pci_14	CMU3/CMP0	/pci@680/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE12
PCIe 13	pci_6	CMU2/CMP1	/pci@480/pci@1/pci@0/pci@8	/SYS/IOU0/PCIE13
PCIe 14	pci_12	CMU3/CMP0	/pci@600/pci@1/pci@0/pci@2	/SYS/IOU0/PCIE14
PCIe 15	pci_13	CMU3/CMP0	/pci@640/pci@1/pci@0/pci@6	/SYS/IOU0/PCIE15
PCIe 16	pci_7	CMU2/CMP1	/pci@4c0/pci@1/pci@0/pci@4	/SYS/IOU0/PCIE16

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
EMS1	pci_0	CMU0/CMP0	/pci@300/pci@1/pci@0/pci@c	/SYS/IOU0/EMS1
EMS2	pci_8	CMU1/CMP0	/pci@500/pci@1/pci@0/pci@0	/SYS/IOU0/EMS2
EMS3	pci_4	CMU2/CMP0	/pci@400/pci@1/pci@0/pci@c	/SYS/IOU0/EMS3
EMS4	pci_12	CMU3/CMP0	/pci@600/pci@1/pci@0/pci@0	/SYS/IOU0/EMS4

Related Information

- "Fully-Populated DCU 1 PCIe Slot Root Complexes" on page 56
- "Fully-Populated DCU 2 PCIe Slot Root Complexes" on page 57
- "Fully-Populated DCU 3 PCIe Slot Root Complexes" on page 59

Fully-Populated DCU 1 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a fully populated DCU 1.

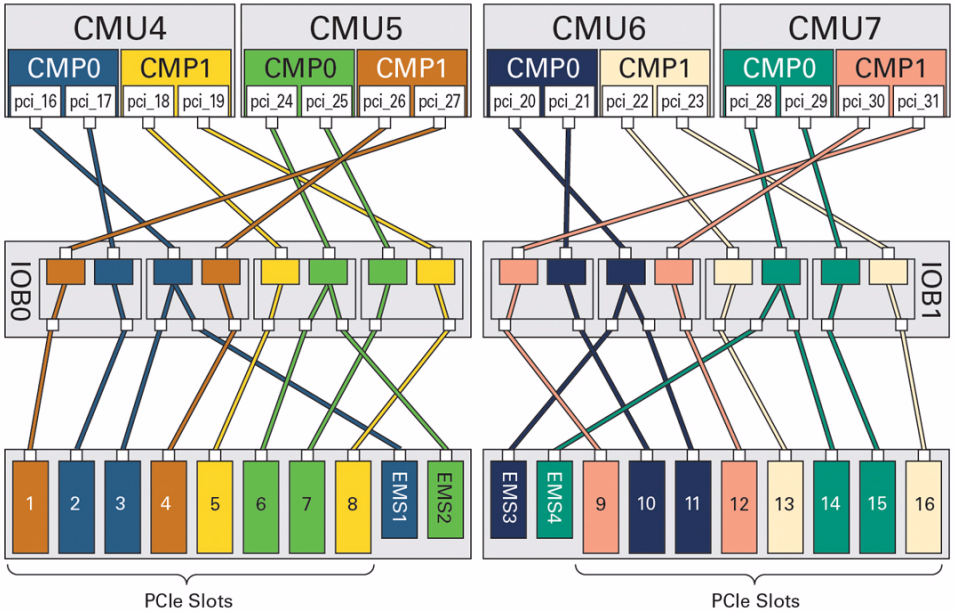

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_27	CMU5/CMP1	/pci@9c0/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE1
PCIe 2	pci_17	CMU4/CMP0	/pci@740/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE2
PCIe 3	pci_16	CMU4/CMP0	/pci@700/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE3
PCIe 4	pci_26	CMU5/CMP1	/pci@980/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE4
PCIe 5	pci_28	CMU4/CMP1	/pci@780/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE5
PCIe 6	pci_24	CMU5/CMP0	/pci@900/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE6
PCIe 7	pci_25	CMU5/CMP0	/pci@940/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE7
PCIe 8	pci_19	CMU4/CMP1	/pci@7c0/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE8
PCIe 9	pci_31	CMU7/CMP1	/pci@ac0/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE9
PCIe 10	pci_21	CMU6/CMP0	/pci@840/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE10
PCIe 11	pci_20	CMU6/CMP0	/pci@800/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE11
PCIe 12	pci_30	CMU7/CMP0	/pci@a80/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE12
PCIe 13	pci_22	CMU6/CMP1	/pci@880/pci@1/pci@0/pci@8	/SYS/IOU1/PCIE13
PCIe 14	pci_28	CMU7/CMP0	/pci@a00/pci@1/pci@0/pci@2	/SYS/IOU1/PCIE14
PCIe 15	pci_29	CMU7/CMP0	/pci@a40/pci@1/pci@0/pci@6	/SYS/IOU1/PCIE15
PCIe 16	pci_23	CMU6/CMP1	/pci@8c0/pci@1/pci@0/pci@4	/SYS/IOU1/PCIE16
EMS1	pci_16	CMU4/CMP0	/pci@700/pci@1/pci@0/pci@c	/SYS/IOU1/EMS1
EMS2	pci_24	CMU5/CMP0	/pci@900/pci@1/pci@0/pci@0	/SYS/IOU1/EMS2
EMS3	pci_20	CMU6/CMP0	/pci@800/pci@1/pci@0/pci@c	/SYS/IOU1/EMS3
EMS4	pci_28	CMU7/CMP0	/pci@a00/pci@1/pci@0/pci@0	/SYS/IOU1/EMS4

Related Information

- "Fully-Populated DCU 0 PCIe Slot Root Complexes" on page 54
- "Fully-Populated DCU 2 PCIe Slot Root Complexes" on page 57
- "Fully-Populated DCU 3 PCIe Slot Root Complexes" on page 59

Fully-Populated DCU 2 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a fully populated DCU 2.

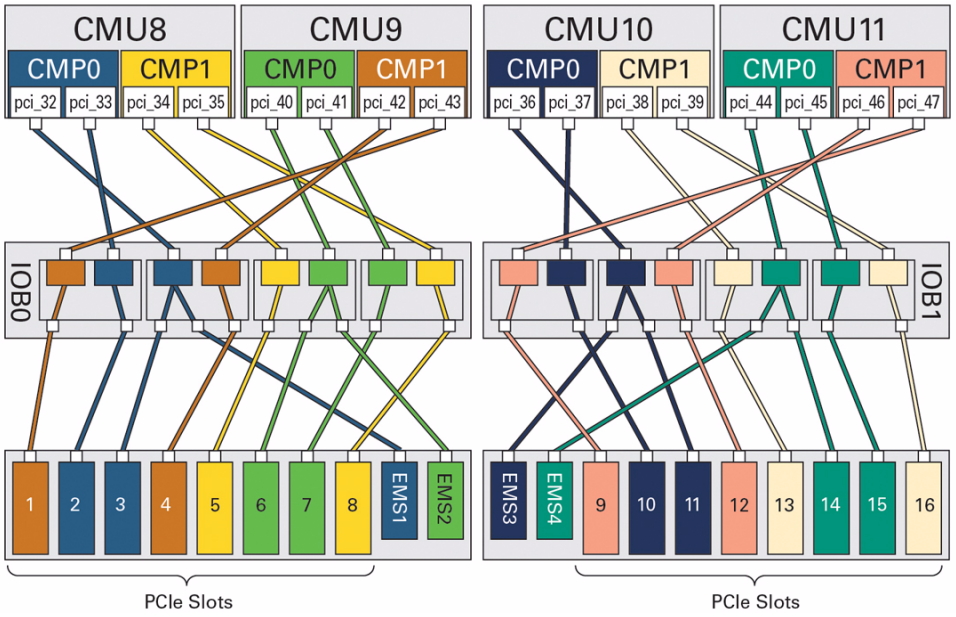

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_43	CMU9/CMP1	/pci@dc0/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE1
PCIe 2	pci_33	CMU8/CMP0	/pci@b40/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE2
PCIe 3	pci_32	CMU8/CMP0	/pci@b00/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE3
PCIe 4	pci_42	CMU9/CMP1	/pci@d80/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE4
PCIe 5	pci_34	CMU8/CMP1	/pci@b80/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE5
PCIe 6	pci_40	CMU9/CMP0	/pci@d00/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE6
PCIe 7	pci_41	CMU9/CMP0	/pci@d40/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE7
PCIe 8	pci_35	CMU8/CMP1	/pci@bc0/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE8
PCIe 9	pci_47	CMU11/CMP1	/pci@ec0/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE9
PCIe 10	pci_37	CMU10/CMP0	/pci@c40/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE10
PCIe 11	pci_36	CMU10/CMP0	/pci@c00/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE11
PCIe 12	pci_46	CMU11/CMP0	/pci@e80/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE12
PCIe 13	pci_38	CMU10/CMP1	/pci@c80/pci@1/pci@0/pci@8	/SYS/IOU2/PCIE13
PCIe 14	pci_44	CMU11/CMP0	/pci@e00/pci@1/pci@0/pci@2	/SYS/IOU2/PCIE14
PCIe 15	pci_45	CMU11/CMP0	/pci@e40/pci@1/pci@0/pci@6	/SYS/IOU2/PCIE15
PCIe 16	pci_39	CMU10/CMP1	/pci@cc0/pci@1/pci@0/pci@4	/SYS/IOU2/PCIE16

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
EMS1	pci_32	CMU8/CMP0	/pci@b00/pci@1/pci@0/pci@c	/SYS/IOU2/EMS1
EMS2	pci_40	CMU9/CMP0	/pci@d00/pci@1/pci@0/pci@0	/SYS/IOU2/EMS2
EMS3	pci_36	CMU10/CMP0	/pci@c00/pci@1/pci@0/pci@c	/SYS/IOU2/EMS3
EMS4	pci_44	CMU11/CMP0	/pci@e00/pci@1/pci@0/pci@0	/SYS/IOU2/EMS4

Related Information

- "Fully-Populated DCU 0 PCIe Slot Root Complexes" on page 54
- "Fully-Populated DCU 1 PCIe Slot Root Complexes" on page 56
- "Fully-Populated DCU 3 PCIe Slot Root Complexes" on page 59

Fully-Populated DCU 3 PCIe Slot Root Complexes

The following illustration displays the PCIe I/O fabric paths from each root complex to the PCIe and EMS slots in a fully populated DCU 3.

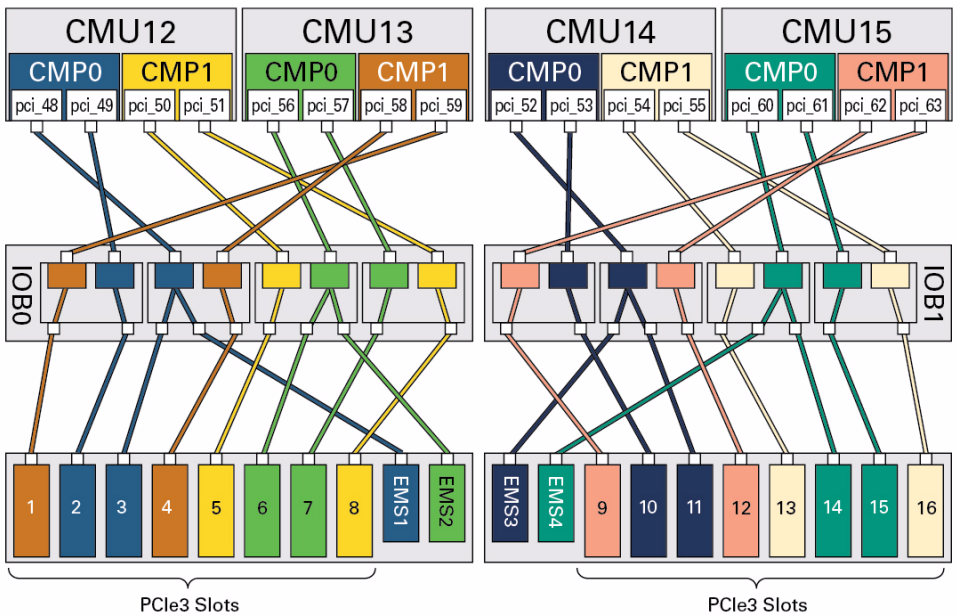

Slot	Root Complex	CMU No./CMP No.	Device Path	FRU Path
PCIe 1	pci_59	CMU13/CMP1	/pci@11c0/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE1
PCIe 2	pci_49	CMU12/CMP0	/pci@f40/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE2
PCIe 3	pci_48	CMU12/CMP0	/pci@f00/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE3
PCIe 4	pci_58	CMU13/CMP1	/pci@1180/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE4
PCIe 5	pci_50	CMU12/CMP1	/pci@f80/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE5
PCIe 6	pci_56	CMU13/CMP0	/pci@1100/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE6
PCIe 7	pci_57	CMU13/CMP0	/pci@1140/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE7
PCIe 8	pci_51	CMU12/CMP1	/pci@fc0/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE8
PCIe 9	pci_63	CMU15/CMP1	/pci@12c0/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE9
PCIe 10	pci_53	CMU14/CMP0	/pci@1040/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE10
PCIe 11	pci_52	CMU14/CMP0	/pci@1000/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE11
PCIe 12	pci_62	CMU15/CMP0	/pci@1280/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE12
PCIe 13	pci_54	CMU14/CMP1	/pci@1080/pci@1/pci@0/pci@8	/SYS/IOU3/PCIE13
PCIe 14	pci_60	CMU15/CMP0	/pci@1200/pci@1/pci@0/pci@2	/SYS/IOU3/PCIE14
PCIe 15	pci_61	CMU15/CMP0	/pci@1240/pci@1/pci@0/pci@6	/SYS/IOU3/PCIE15
PCIe 16	pci_55	CMU14/CMP1	/pci@10c0/pci@1/pci@0/pci@4	/SYS/IOU3/PCIE16
EMS1	pci_48	CMU12/CMP0	/pci@f00/pci@1/pci@0/pci@c	/SYS/IOU3/EMS1
EMS2	pci_56	CMU13/CMP0	/pci@1100/pci@1/pci@0/pci@0	/SYS/IOU3/EMS2
EMS3	pci_52	CMU14/CMP0	/pci@1000/pci@1/pci@0/pci@c	/SYS/IOU3/EMS3
EMS4	pci_60	CMU15/CMP0	/pci@1200/pci@1/pci@0/pci@0	/SYS/IOU3/EMS4

Related Information

- "Fully-Populated DCU 0 PCIe Slot Root Complexes" on page 54
- "Fully-Populated DCU 1 PCIe Slot Root Complexes" on page 56
- "Fully-Populated DCU 2 PCIe Slot Root Complexes" on page 57

Understanding PDomains

These topics describe PDomains and the PDomain configurations.

- "PDomains Overview" on page 61
- "PDomain Guidelines" on page 62
- "Extended Configuration PDomain Overview" on page 63
- "Understanding Extended Configuration PDomains" on page 64
- "Base Configuration Overview" on page 70
- "Understanding Base Configuration PDomains" on page 71

PDomains Overview

Each PDomain has its own set of fans and I/O boards, and operates as an independent server that has full hardware isolation from other PDomains in the chassis. A hardware or software failure on one PDomain does not affect the other PDomains in the chassis. There are two redundant SPs (SP0 and SP1) in the chassis. One of the SPs functions as the Active-SP and actively manages the chassis, while the other acts as the Standby-SP that assumes the Active-SP role in the event of a failure. You can use the Oracle ILOM software that is preinstalled on each SP to remotely monitor and control PDomains.

There are two types of PDomains:

PDomain Type	Description
Nonbounded PDomain	The processors of a PDomain can communicate with the rest of the processors in the server. To enable the processor communication, processor data traffic is routed through the SSBs, which might indicate lower performance compared to a bounded PDomain. A nonbounded PDomain can contain one to four DCUs.
Bounded PDomain	Each bounded PDomain contains only one DCU. The processors within a bounded PDomain can communicate only with each other and not with other processors in other DCUs. Since the processors communicate only within the DCU, bounded PDomains perform better than nonbounded PDomains.

Each server can contain either two or four PDomains, in the following configurations:

- All PDomains in a single server, with the two or four PDomains all contained within the single server.
- PDomains split across two servers, with one of these situations:
  - Two PDomains are split across the two servers (one PDomain on each server)
  - Four PDomains are split across the two servers (two PDomains on each server)

Related Information

- "PDomain Guidelines" on page 62
- "Extended Configuration PDomain Overview" on page 63
- "Understanding Extended Configuration PDomains" on page 64
- "Base Configuration Overview" on page 70
- "Understanding Base Configuration PDomains" on page 71

PDomain Guidelines

- In configurations where a DCU has only two CMUs, individual CMPs can be unconfigured. In configurations where a DCU has four CMUs, an entire CMU will be unconfigured if a CMP needs to be reconfigured.
- In a multiple-PDomain configuration, PDomain_0 should contain SAS0 and rKVMS from each DCU.

Related Information

- "PDomains Overview" on page 61
- "Extended Configuration PDomain Overview" on page 63
- "Understanding Extended Configuration PDomains" on page 64
- "Base Configuration Overview" on page 70
- "Understanding Base Configuration PDomains" on page 71

Extended Configuration PDomain Overview

The extended configuration has the following characteristics:

- Two PDomains
- Two DCUs associated with each PDomain
- PDomains on one or two servers, with one of these configurations:
  - On a single compute server, with the two PDomains within that compute server, each with two DCUs
  - Split across two compute servers, with the two PDomains split across the two compute servers (one PDomain on each server), with each PDomain containing two DCUs

Because there are two DCUs associated with each PDomain in an extended configuration, the PDomain in an extended configuration is considered a nonbounded PDomain. See "PDomains Overview" on page 61 for more information on the characteristics of this type of PDomain.

Note – Certain DCU configurations have a mixture of half-populated and fully-populated DCUs in a compute server. For extended configuration PDomains in those DCU configurations, both DCUs in each PDomain must have the same configuration, where both DCUs in the PDomain are either half-populated or fully-populated. See "CMU Overview" on page 5 for more information.

The following figure illustrates the available extended configurations.

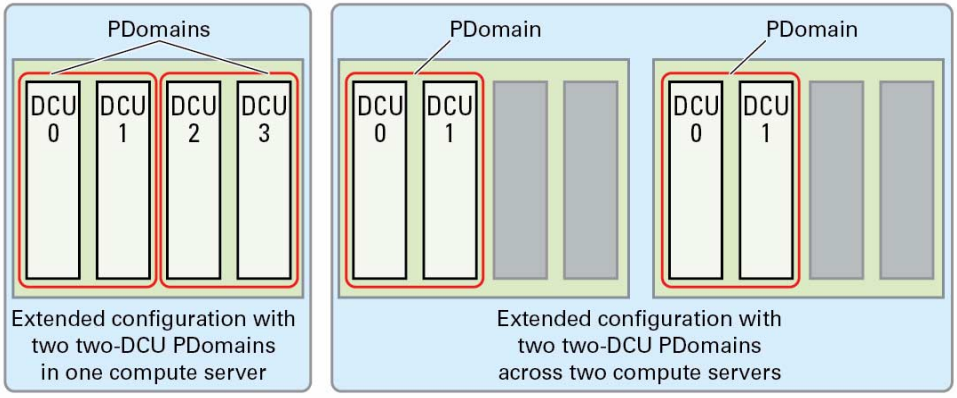

Related Information

- "PDomains Overview" on page 61
- "PDomain Guidelines" on page 62
- "Understanding Extended Configuration PDomains" on page 64

Understanding Extended Configuration PDomains

The following extended configuration PDomains are available:

- "Understanding Four DCUs in One Compute Server (R1 Extended Configuration PDomains)" on page 64
- "Understanding Four DCUs Across Two Compute Servers (R2 Extended Configuration PDomains)" on page 67

Understanding Four DCUs in One Compute Server (R1 Extended Configuration PDomains)

This configuration has the following characteristics:

- Two PDomains, with two DCUs associated with each PDomain
- Both PDomains on a single compute server

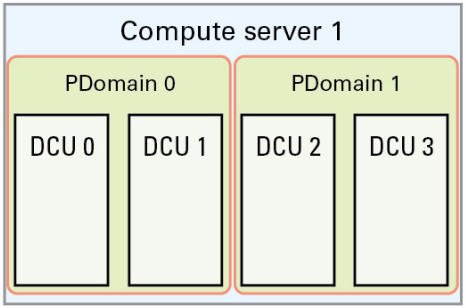

Three DCU configuration are options available for this configuration:

- "Four Fully-Populated DCUs (R1_1 PDomain Configuration)" on page 65
- "Two Fully-Populated DCUs and Two Half-Populated DCUs (R1_2 PDomain Configuration)" on page 66
- "Four Half-Populated DCUs (R1_3 PDomain Configuration)" on page 66

Four Fully-Populated DCUs (R1_1 PDomain Configuration)

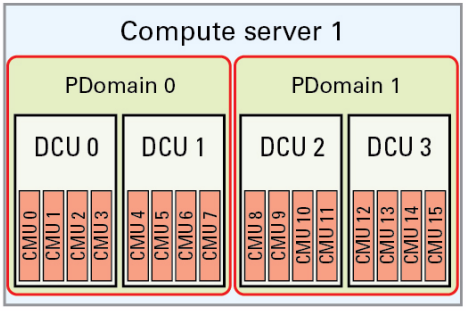

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99

Two Fully-Populated DCUs and Two Half-Populated DCUs (R1_2 PDomain Configuration)

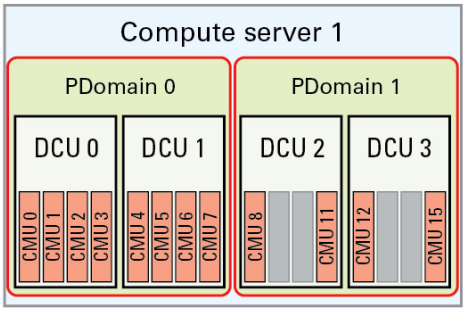

For this configuration, the following LDom configurations are available for each DCU:

- DCU 0 and DCU 1: "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99
- DCU 2 and DCU 3: "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Four Half-Populated DCUs (R1_3 PDomain Configuration)

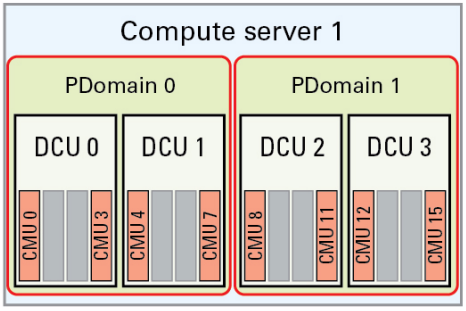

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Understanding Four DCUs Across Two Compute Servers (R2 Extended Configuration PDomains)

This configuration has the following characteristics:

- Two PDomains, with two DCUs associated with each PDomain
- The two PDomains split across two compute servers (one PDomain on each compute server)

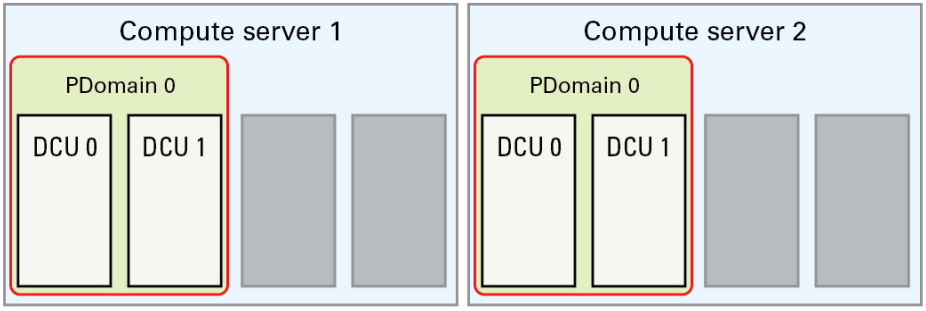

Three DCU configuration are options available for this configuration:

- "Four Fully-Populated DCUs (R2_1 PDomain Configuration)" on page 68
- "Two Fully-Populated DCUs and Two Half-Populated DCUs (R2_2 PDomain Configuration)" on page 69
- "Four Half-Populated DCUs (R2_3 PDomain Configuration)" on page 69

Four Fully-Populated DCUs (R2_1 PDomain Configuration)

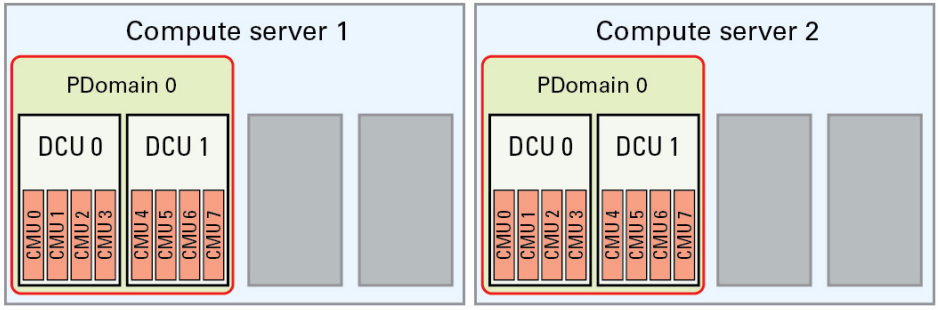

For this configuration, the following LDom configurations are available for each DCU:

- Extended PDomain on first compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99
- Extended PDomain on second compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99

Two Fully-Populated DCUs and Two Half-Populated DCUs (R2_2 PDomain Configuration)

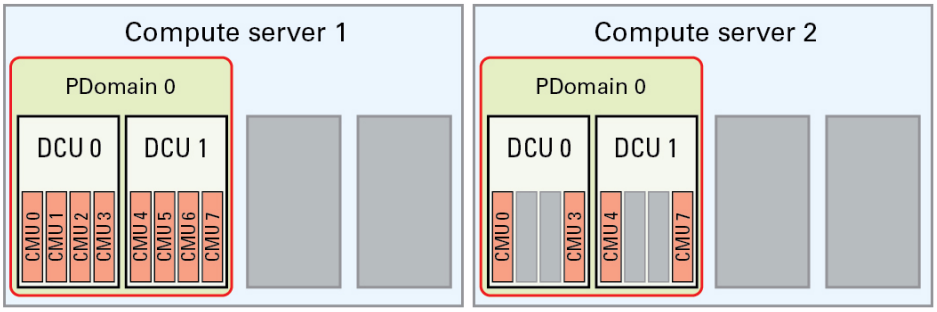

For this configuration, the following LDom configurations are available for each DCU:

- Extended PDomain on first compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99
- Extended PDomain on second compute server: "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Four Half-Populated DCUs (R2_3 PDomain Configuration)

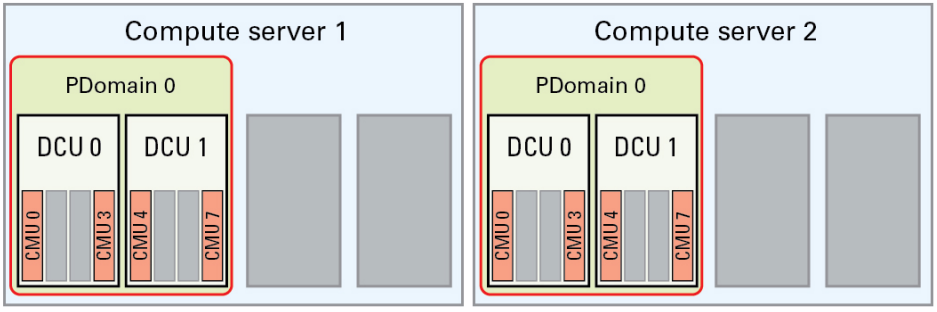

For this configuration, the following LDom configurations are available for each DCU:

- Extended PDomain on first compute server: "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108
- Extended PDomain on second compute server: "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Base Configuration Overview

The base configurations have the following characteristics:

- Two or four PDomains
- One DCU associated with each PDomain
- PDomains on one or two servers, with one of these configurations:
  - On a single server, with the two or four PDomains within that server, each with one DCU
  - Split across two servers, with either two PDomains split across the two servers (one PDomain on each server) or four PDomains split across the two servers (two PDomains on each server), with each PDomain containing a single DCU

Because there is one DCU associated with each PDomain in a base configuration, the PDomain in a base configuration is considered a bounded PDomain. See "PDomains Overview" on page 61 for more information on the characteristics of this type of PDomain.

The following figure illustrates the available base configurations.

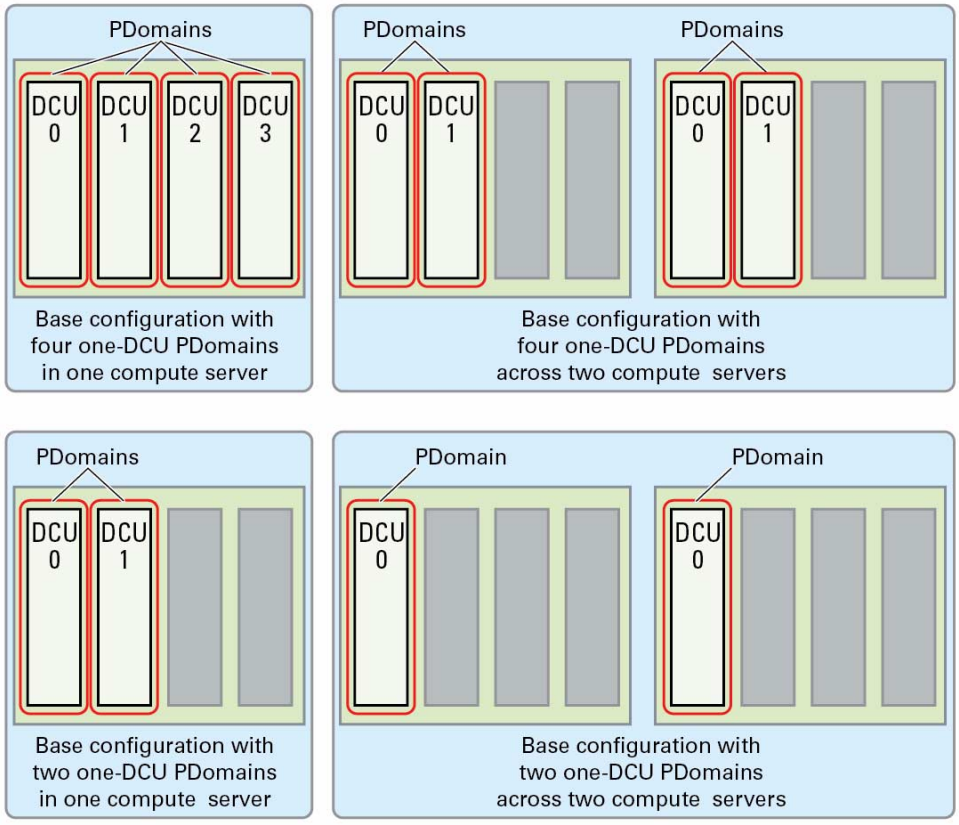

Related Information

- "PDomains Overview" on page 61
- "PDomain Guidelines" on page 62
- "Understanding Base Configuration PDomains" on page 71

Understanding Base Configuration PDomains

The following base configuration PDomains are available:

- "Understanding Four DCUs on One Compute Server (R3 Base Configuration PDomains)" on page 72

- "Understanding Four DCUs Across Two Compute Servers (R4 Base Configuration PDomains)" on page 74
- "Understanding Two DCUs on One Compute Server (R5 Base Configuration PDomains)" on page 77
- "Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains)" on page 78

Understanding Four DCUs on One Compute Server (R3 Base Configuration PDomains)

This configuration has the following characteristics:

- Four PDomains, with one DCU associated with each PDomain
- All four PDomains on a single compute server

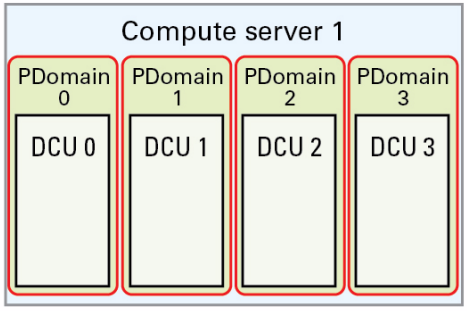

There are three DCU configuration options available for four DCUs on one compute server:

- "Four Fully-Populated DCUs (R3_1 PDomain Configuration)" on page 73
- "Two Fully-Populated DCUs and Two Half-Populated DCUs (R3_2 PDomain Configuration)" on page 73
- "Four Half-Populated DCUs (R3_3 PDomain Configuration)" on page 74

Four Fully-Populated DCUs (R3_1 PDomain Configuration)

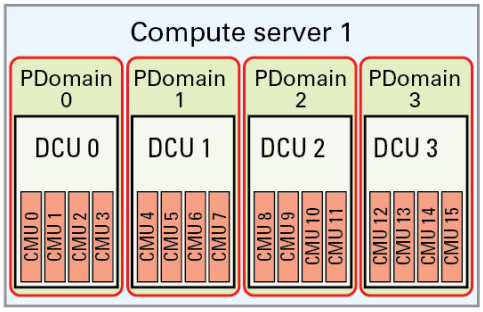

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117

Two Fully-Populated DCUs and Two Half-Populated DCUs (R3_2 PDomain Configuration)

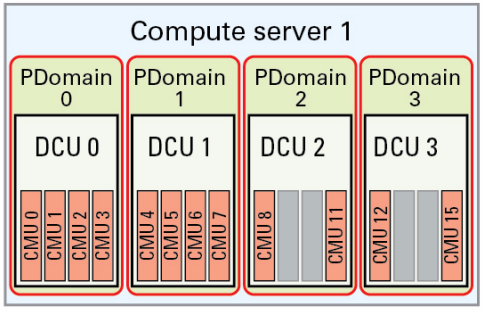

For this configuration, the following LDom configurations are available for each DCU:

- DCU 0 and DCU 1: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
- DCU 2 and DCU 3: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Four Half-Populated DCUs (R3_3 PDomain Configuration)

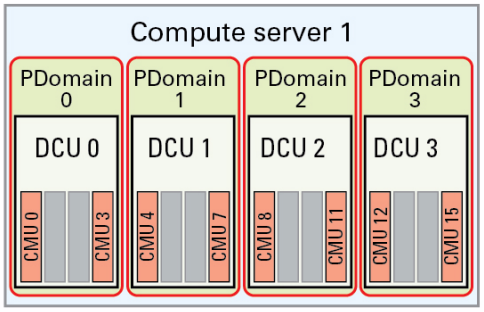

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Understanding Four DCUs Across Two Compute Servers (R4 Base Configuration PDomains)

This configuration has the following characteristics:

- Four PDomains, with one DCU associated with each PDomain
- The four PDomains split across two compute servers (two PDomains on each compute server)

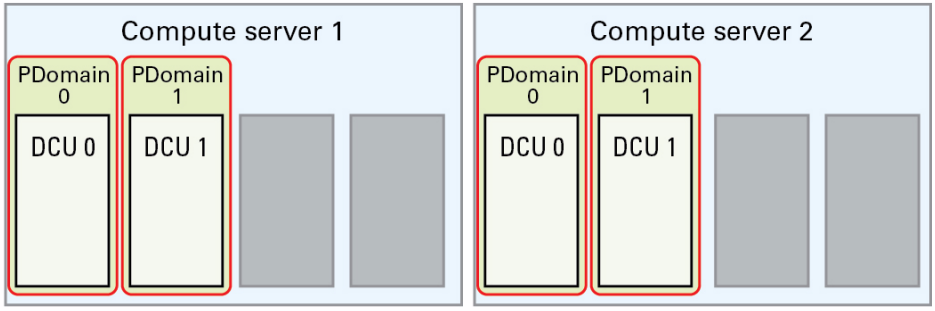

There are four DCU configuration options available for this configuration:

- "Four Fully-Populated DCUs (R4_1 PDomain Configuration)" on page 75

- "Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_2 PDomain Configuration)" on page 76
- "Four Half-Populated DCUs (R4_3 PDomain Configuration)" on page 76
- "Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_4 PDomain Configuration)" on page 77

Four Fully-Populated DCUs (R4_1 PDomain Configuration)

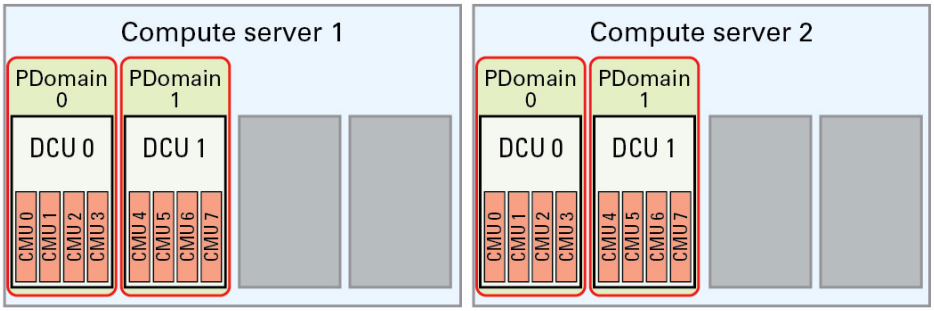

For this configuration, the following LDom configurations are available for each DCU:

- Two base PDomains on first compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
- Two base PDomains on second compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117

Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_2 PDomain Configuration)

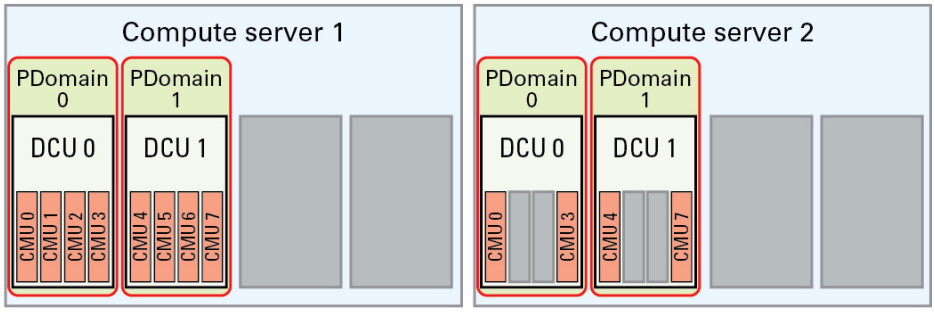

For this configuration, the following LDom configurations are available for each DCU:

- Two base PDomains on first compute server: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
- Two base PDomains on second compute server: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Four Half-Populated DCUs (R4_3 PDomain Configuration)

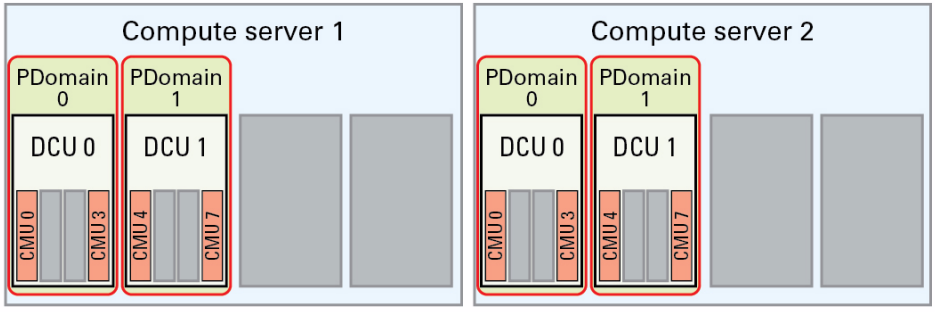

For this configuration, the following LDom configurations are available for each DCU:

- Two base PDomains on first compute server: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

- Two base PDomains on second compute server: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_4 PDomain Configuration)

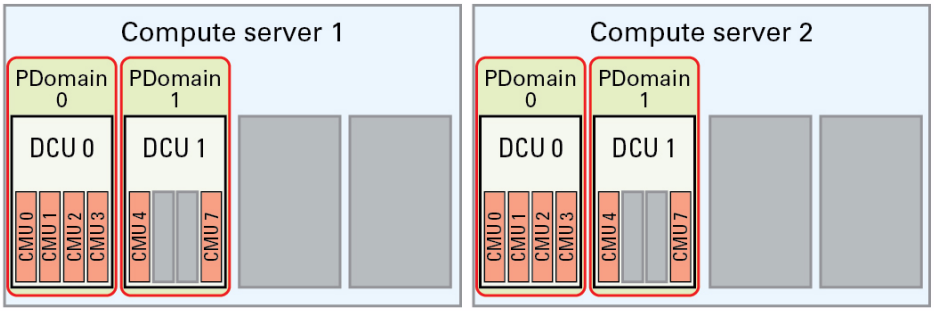

For this configuration, the following LDom configurations are available for each DCU:

- First compute server:
  - DCU 0: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
  - DCU 1: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124
- Second compute server:
  - DCU 0: "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
  - DCU 1: "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Understanding Two DCUs on One Compute Server (R5 Base Configuration PDomains)

This configuration has the following characteristics:

- Two PDomains, with one DCU associated with each PDomain
- Both PDomains on a single compute server

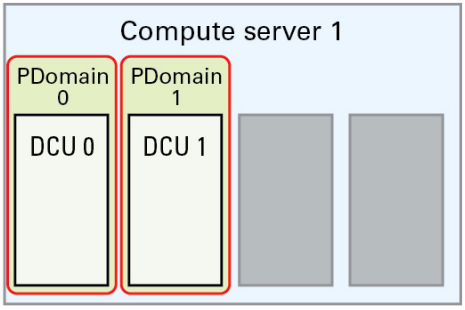

There is one DCU configuration option available for two DCUs on one compute server:

- "Two Fully-Populated DCUs (R5_1 PDomain Configuration)" on page 78

Two Fully-Populated DCUs (R5_1 PDomain Configuration)

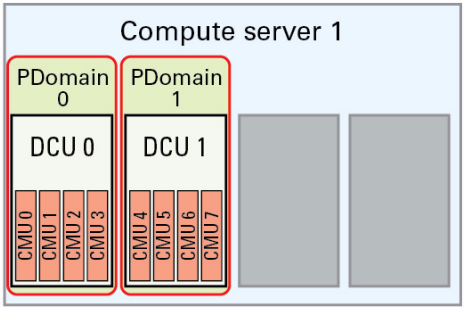

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117

Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains)

This configuration has the following characteristics:

- Two PDomains, with one DCU associated with each PDomain

- The two PDomains split across two compute servers (one PDomain on each compute server)

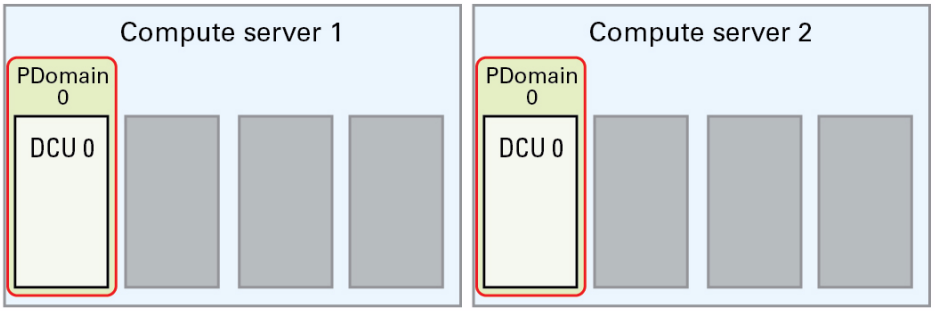

There is one DCU configuration option available for two DCUs across two compute servers:

- "Two Fully-Populated DCUs (R6_1 PDomain Configuration)" on page 79

Two Fully-Populated DCUs (R6_1 PDomain Configuration)

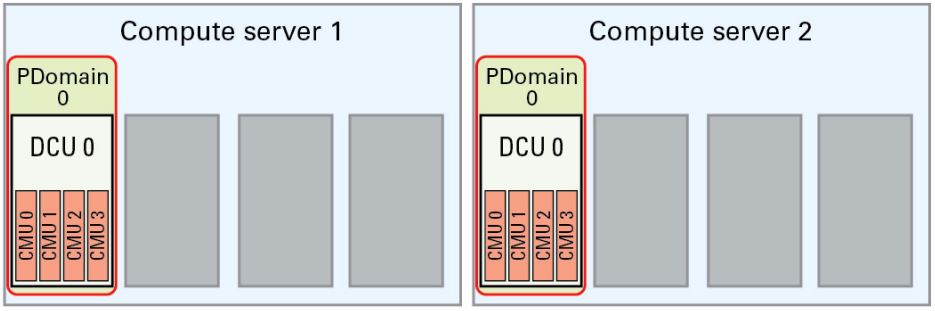

For this configuration, the following LDom configurations are available for each DCU:

- "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117

Understanding LDoms

LDoms provide the flexibility to create different specialized virtual systems within a single hardware platform. Each DCU that is installed in each server is configured with at least one LDom.

The LDom configurations that are available varies, depending on the following factors:

- Type of PDomain configuration:
  - Extended configuration
  - Base configuration
- Type of DCU configuration:
  - Fully-populated DCU configuration
  - Half-populated DCU configuration

The following sections provide more information on these configuration options:

- "Understanding Compute Server Hardware and Networks" on page 82
- "Understanding Domains" on page 87
- "Understanding LDom Configurations for Extended Configuration PDomains" on page 99
- "Understanding LDom Configurations for Base Configuration PDomains" on page 116

Understanding Compute Server Hardware and Networks

These topics describe the compute server hardware and networks in relation to LDoms.

- "CPU and Memory Resources Overview" on page 82
- "LDoms and the PCIe Slots Overview" on page 83
- "Management Network Overview" on page 84
- "10GbE Client Access Network Overview" on page 84
- "Understanding the IB Network" on page 85

CPU and Memory Resources Overview

When LDoms are created on SuperCluster M6-32, the amount of CPU and memory resources allocated to the LDom is set between 25% and 100% of the CPU and memory resources for one or two DCUs, depending on the type of PDomain configuration (base or extended) and the type of the DCU configuration (half-populated or fully-populated).

You can use the CPU/Memory tool (`osc-setcoremem`) to change this default allocation after the initial installation of your system. However, you must reboot the resized domain after you have reallocated the CPU and memory resources for the changes to take affect. Refer to the *Oracle SuperCluster M6-32 Owner's Guide: Administration* for more information.

Related Information

- "LDoms and the PCIe Slots Overview" on page 83
- "Management Network Overview" on page 84
- "10GbE Client Access Network Overview" on page 84
- "Understanding the IB Network" on page 85

LDoms and the PCIe Slots Overview

LDom configurations supported on SuperCluster M6-32 have the following characteristics:

- One to four LDoms on one or two DCUs
- Each LDom can be one of the following types:
  - Database Domain (dedicated domain)
  - Application Domain running Oracle Solaris 10 (dedicated domain)
  - Application Domain running Oracle Solaris 11 (dedicated domain)
  - Root Domain

Note – For more information on the different LDom types, see "Understanding Domains" on page 87.

Each DCU in the compute server has sixteen PCIe slots. The following cards are installed in certain PCIe slots and are used to connect to these networks:

- One quad-port 1GbE NIC – Used to connect to the management network
- Four IB HCAs – Used to connect to the private IB network

Each DCU also has four EMS cards, each with two 10GbE ports, which are used to connect to the 10GbE client access network.

Note – Optional Fibre Channel PCIe cards are also available to facilitate migration of data from legacy storage subsystems to the storage servers integrated with SuperCluster M6-32. The PCIe slots that are available for those optional Fibre Channel PCIe cards varies, depending on your configuration. Refer to the *Oracle SuperCluster M6-32 Owner's Guide: Installation* for more information.

The PCIe slots and EMSs used for each configuration varies, depending on the type and number of LDoms that are used for that configuration.

Related Information

- "Understanding Domains" on page 87
- "CPU and Memory Resources Overview" on page 82
- "Management Network Overview" on page 84
- "10GbE Client Access Network Overview" on page 84
- "Understanding the IB Network" on page 85

Management Network Overview

The management network connects to your existing management network, and is used for administrative work. Each DCU in the compute server provides access to the following management networks:

- Oracle ILOM management network – Connected through the Oracle ILOM Ethernet interface on the SPs in each compute server. Connections to this network are the same, regardless of the type of PDomain or LDom configuration.
- 1GbE host management network – Connected through the four 1GbE host management interfaces (NET0 - NET3) on the quad-port 1GbE NIC installed in each DCU. Connections to this network varies, depending on the type of configuration that is set up on the system. In most cases, the four 1GbE host management ports use IPMP to provide redundancy for the management network interfaces to the LDoms. However, the ports that are grouped together varies depending on the type of PDomain configuration.

Related Information

- "CPU and Memory Resources Overview" on page 82
- "LDoms and the PCIe Slots Overview" on page 83
- "10GbE Client Access Network Overview" on page 84
- "Understanding the IB Network" on page 85

10GbE Client Access Network Overview

This required 10GbE network connects the compute servers to your existing client network and is used for client access to the LDoms. EMSs in each DCU are used for connection to this network. The number of EMSs used for each LDom varies depending on the type of configuration.

Related Information

- "CPU and Memory Resources Overview" on page 82
- "LDoms and the PCIe Slots Overview" on page 83
- "Management Network Overview" on page 84
- "Understanding the IB Network" on page 85

Understanding the IB Network

These topics describe the IB network.

- "IB Network Overview" on page 85
- "IB Network Data Paths for a Database Domain" on page 85
- "IB Network Data Paths for an Application Domain" on page 86

IB Network Overview

The IB network connects the compute servers, storage appliance, and storage servers using the IB HCAs installed in the PCIe slots in each DCU and the IB switches installed in the expansion rack. This nonroutable network is fully contained in SuperCluster M6-32, and does not connect to your existing network.

When SuperCluster M6-32 is configured with the appropriate types of LDoms, the IB network is partitioned to define the data paths between the compute servers, and between the compute servers and the storage appliances.

The defined IB data path coming out of the compute servers varies, depending on the type of LDom created:

- "IB Network Data Paths for a Database Domain" on page 85
- "IB Network Data Paths for an Application Domain" on page 86

IB Network Data Paths for a Database Domain

Note – The information in this section applies to a Database Domain that is either a dedicated domain or a Database I/O Domain.

When a Database Domain is created, the Database Domain has the following IB paths:

- Compute server to both IB leaf switches in the expansion rack
- Compute server to each storage server, through the IB leaf switches in the expansion rack
- Compute server to the storage appliance, through the IB leaf switches in the expansion rack

The number of IB HCAs that are assigned to the Database Domain varies, depending on the type of configuration.

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

- **Storage private network:** One IB private network for the Database Domains to communicate with each other and with the Application Domains running Oracle Solaris 10, and with the storage appliance
- **Exadata private network**: One IB private network for the Oracle Database 11g RAC interconnects, and for communication between the Database Domains and the storage servers

The two ports on each IB HCA connect to different IB leaf switches to provide redundancy between the compute servers and the leaf switches.

Related Information

- "IB Network Data Paths for an Application Domain" on page 86

IB Network Data Paths for an Application Domain

Note – The information in this section applies to an Application Domain that is either a dedicated domain or an Application I/O Domain.

When an Application Domain is created (an Application Domain running either Oracle Solaris 10 or Oracle Solaris 11), it has the following IB paths:

- Compute server to both IB leaf switches in the expansion rack
- Compute server to the storage appliance, through the IB leaf switches in the expansion rack

The Application Domain does not access the storage servers, which are used only for the Database Domain.

The number of IB HCAs that are assigned to the Application Domain varies, depending on the type of configuration.

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

- **Storage private network**: One IB private network for Application Domains to communicate with each other and with the Database Domains, and with the storage appliance
- **Oracle Solaris Cluster private network**: Two IB private networks for the optional Oracle Solaris Cluster interconnects

The two ports on each IB HCA connects to different IB leaf switches to provide redundancy between the compute servers and the leaf switches.

Related Information

- "IB Network Data Paths for a Database Domain" on page 85

Understanding Domains

These topics describe the domain types:

- "Dedicated Domains" on page 87
- "Understanding SR-IOV Domain Types" on page 88

Dedicated Domains

The following SuperCluster-specific domain types have always been available:

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

These SuperCluster-specific domain types have been available in software version 1.x and are now known as **dedicated domains**.

Note – The Database Domains can also be in two states: with zones or without zones.

When SuperCluster M6-32 is set up as part of the initial installation, each domain is assigned one of these three SuperCluster-specific dedicated domain types. With these dedicated domains, every domain in SuperCluster M6-32 has direct access to the EMSs and IB HCAs (and Fibre Channel cards, if those are installed in the card slots). The following graphic shows this concept on SuperCluster M6-32 with four domains.

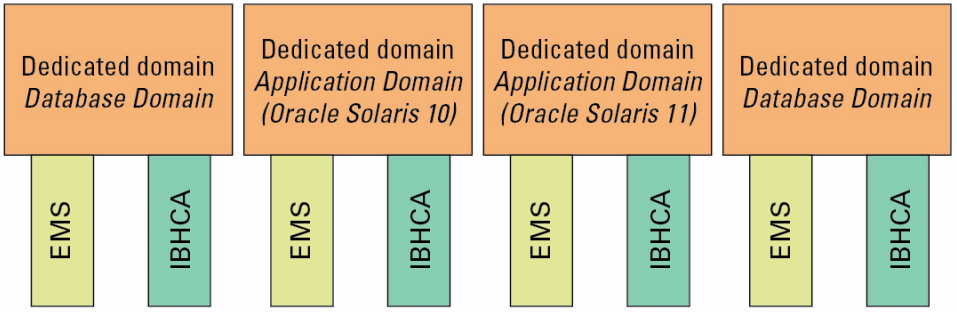

With dedicated domains, connections to the 10GbE client access network go through the physical ports on each EMS, and connections to the IB network go through the physical ports on each IB HCA, as shown in the following graphic.

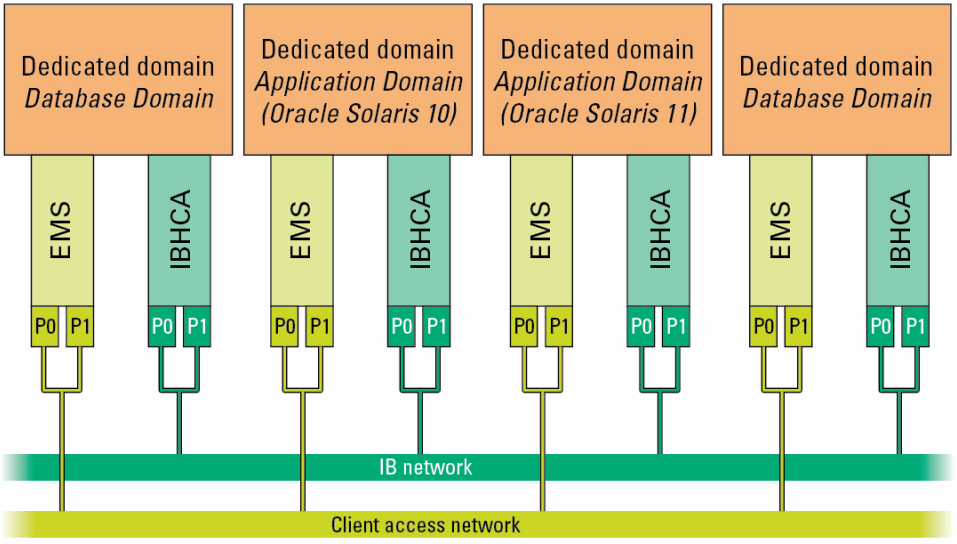

With dedicated domains, the domain configuration for SuperCluster M6-32 (the number of domains and the SuperCluster-specific types assigned to each) are set at the time of the initial installation, and can only be changed by an Oracle representative.

Understanding SR-IOV Domain Types

In addition to the dedicated domain types (Database Domains and Application Domains running either Oracle Solaris 10 or Oracle Solaris 11), the following version 2.x SR-IOV (Single-Root I/O Virtualization) domain types are now also available:

- "Root Domains" on page 89
- "I/O Domains" on page 93

Root Domains

A Root Domain is an SR-IOV domain that hosts the physical I/O devices, or physical functions (PFs), such as the IB HCAs, EMSs, and any Fibre Channel cards installed in the PCIe slots. Almost all of its CPU and memory resources are parked for later use by I/O Domains. Logical devices, or virtual functions (VFs), are created from each PF, with each PF hosting 32 VFs.

Because Root Domains host the physical I/O devices, just as dedicated domains currently do, Root Domains essentially exist at the same level as dedicated domains.

With the introduction of Root Domains, the following parts of the domain configuration for SuperCluster M6-32 are set at the time of the initial installation and can only be changed by an Oracle representative:

- Type of domain:
  - Root Domain
  - Application Domain running Oracle Solaris 10 (dedicated domain)
  - Application Domain running Oracle Solaris 11 (dedicated domain)
  - Database Domain (dedicated domain)
- Number of Root Domains and dedicated domains on the server

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 you have a domain that has more than two IB HCAs associated with it (for example, the B4-1 domain), then that domain must be a dedicated domain.

When deciding which domains will be a Root Domain, the last domain must always be the first Root Domain, and you would start from the last domain in your configuration and go in for every additional Root Domain. For example, assume you have four domains in your configuration, and you want two Root Domains and two dedicated domains. In this case, the first two domains would be dedicated domains and the last two domains would be Root Domains.

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.

The following domains have only one or two IB HCAs associated with them and can therefore be used as a Root Domain:

- Fully-populated DCUs (either base or extended configurations):
  - Domains associated with two CMPs (one IB HCA)
  - Domains associated with four CMPs (two IB HCAs)
- Half-populated DCUs (either base or extended configurations):
  - Domains associated with one CMP (one IB HCA)
  - Domains associated with two CMPs (two IB HCAs)

In addition, the first domain in the system (the Control Domain) will always be a dedicated domain. The Control Domain cannotbe a Root Domain. Therefore, you cannot have all of the domains on your server as Root Domains, but you can have a mixture of Root Domains and dedicated domains on your server or all of the domains as dedicated domains.

A certain amount of CPU core and memory is 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 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
  - Four cores and 64 GB of memory reserved for a Root Domain with two IB HCAs
- Any other domain in a domain configuration:
  - One core and 16 GB of memory reserved for a Root Domain with one IB HCA
  - Two cores and 32 GB of memory reserved for a Root Domain with two IB HCAs

---

Note – The amount of CPU core and memory reserved for Root Domains is sufficient to support only the PFs in each Root Domain. There is insufficient CPU core 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 core and memory resources associated with each Root Domain are parked in CPU and memory repositories, as shown in the following graphic.

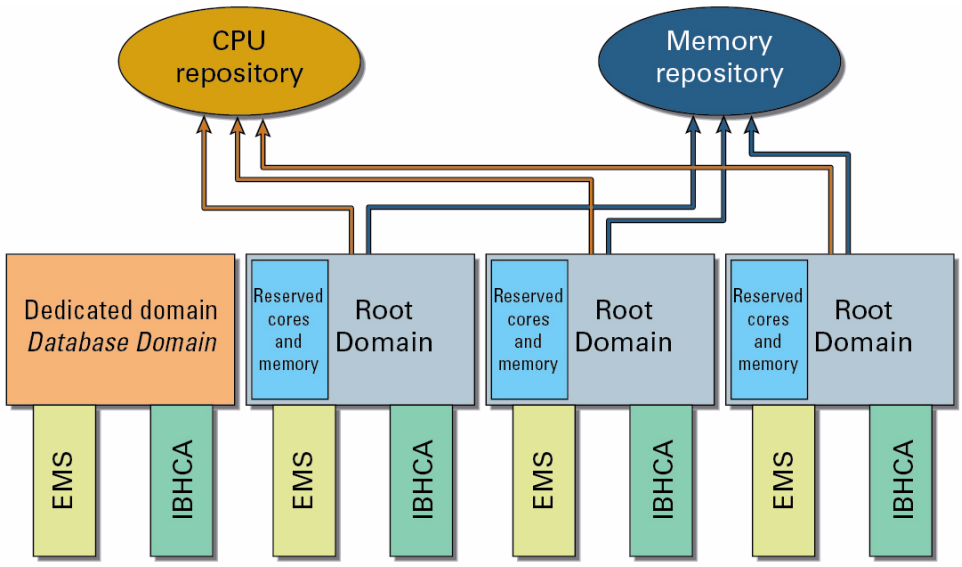

CPU and memory repositories contain resources not only from the Root Domains, but also any parked resources from the dedicated domains. Whether CPU core and memory resources originated from dedicated domains or from Root Domains, once those resources have been parked in the CPU and memory repositories,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

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

For example, assume you have the B2-4 LDom configuration, with four domains on your PDomain, with three of the four domains as Root Domains, as shown in the previous graphic. Assume each domain has the following IB HCAs and EMSs, and the following CPU core and memory resources:

- One IB HCA and one EMS
- 12 cores
- 512 GB of memory (16GB DIMMs)

In this situation, the following CPU core 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. 10 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.
  - 11 cores and 496 GB of memory available from each of these Root Domains for the CPU and memory repositories.
  - A total of 22 cores (11 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 32 cores (10 + 22 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 10GbE client access network go through the physical ports on each EMS, and connections to the IB network go through the physical ports on each IB HCA, just as they did with dedicated domains. However, cards 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.

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.

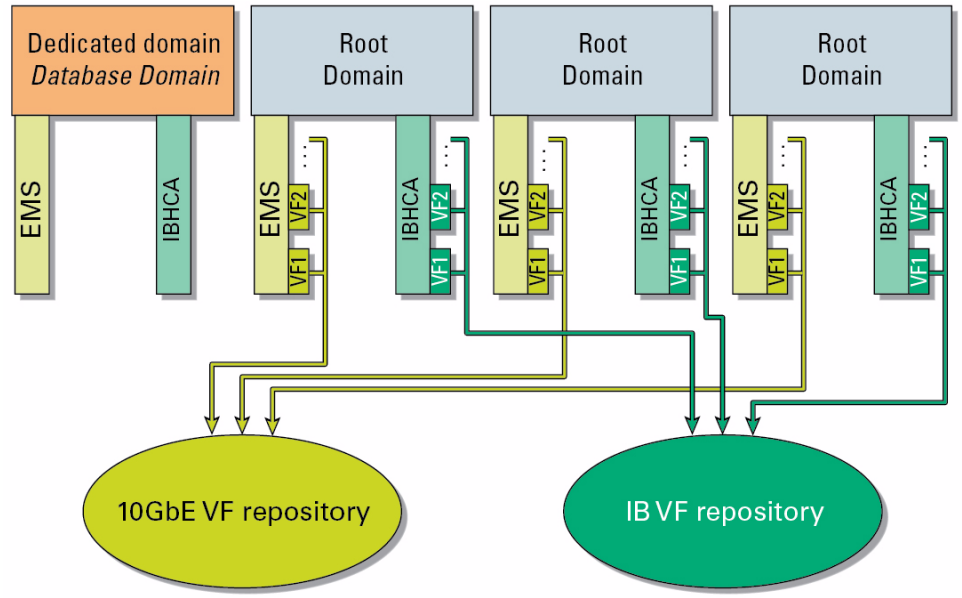

Even though the VFs from each Root Domain are parked in the VF repositories, the VFs are created on each EMS and IB HCA, so those VFs are associated with the Root Domain that contains those specific EMSs and IB HCA cards. For example, looking at the example configuration in the previous graphic, the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

I/O Domains

An I/O Domain is an SR-IOV domain that owns its own VFs, each of which is a virtual device based on a PF in one of the Root Domains. Root domains function solely as a provider of VFs to the I/O Domains, based on the physical I/O devices associated with each Root Domain. Applications and zones are supported only in I/O Domains, not in Root Domains.

You can create multiple I/O Domains using the I/O Domain Creation tool. As part of the domain creation process, you also associate one of the following SuperCluster-specific domain types to each I/O Domain:

- Application Domain running Oracle Solaris 11, or Application I/O Domain
- Database I/O Domain

Note that only Database Domains that are dedicated domains can host database zones. Database I/O Domains cannot host database zones.

The CPU cores 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 the following graphic.

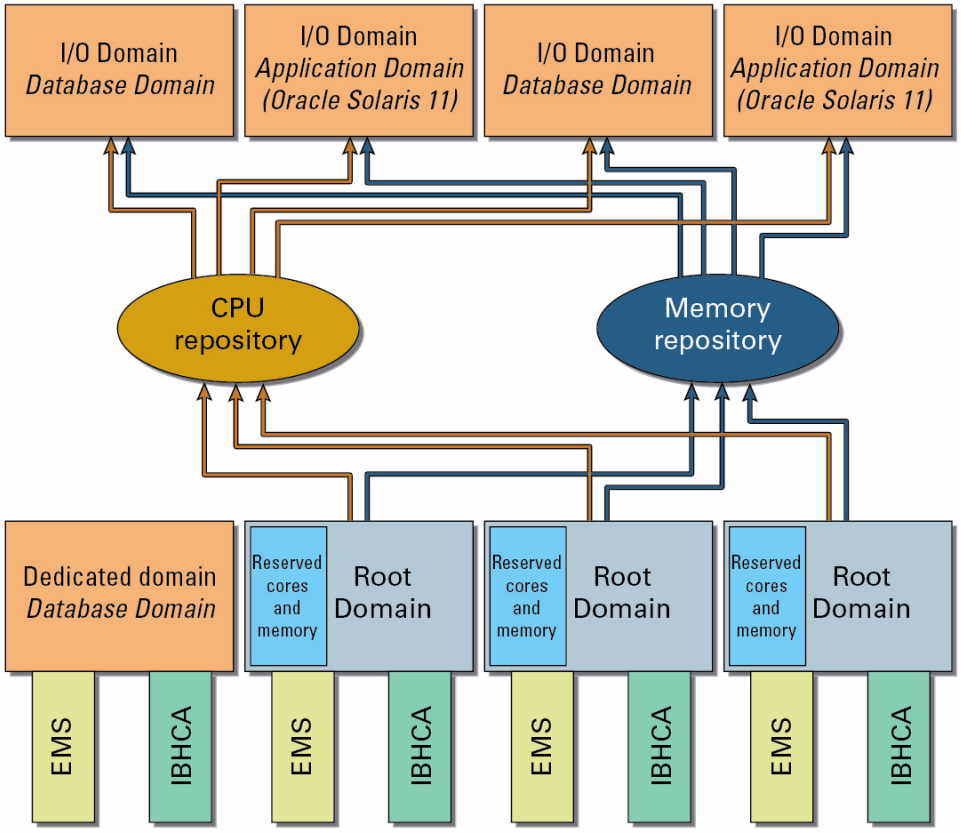

You use the I/O Domain Creation tool to assign the CPU core and memory resources to the I/O Domains, based on the amount of CPU core and memory resources that you want to assign to each I/O Domain and the total amount of CPU core and memory resources available in the CPU and memory repositories.

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.

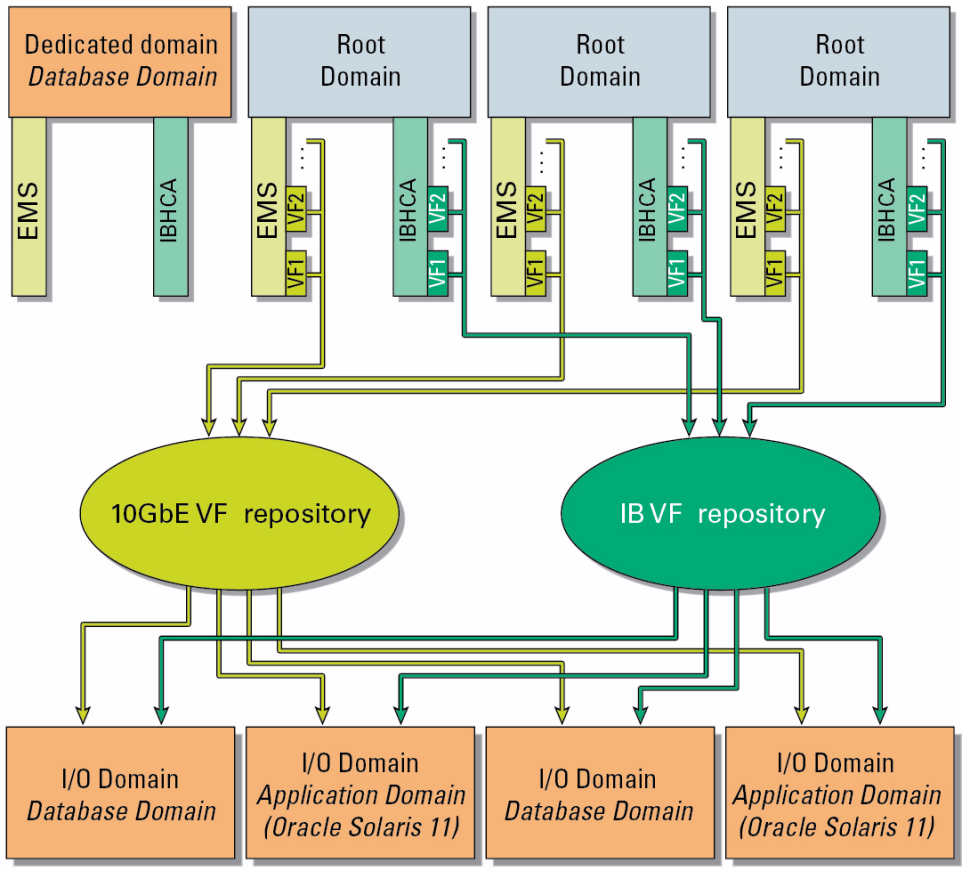

Again, you use the I/O Domain Creation tool to assign the IB VFs and 10GbE VFs to the I/O Domains using the resources available in the IB VF and 10GbE VF repositories. However, because VFs are created on each EMS and IB HCA, the VFs assigned to an I/O Domain will always come from the specific Root Domain that is associated with the EMSs and IB HCAs that contain those VFs.

The number and size of the I/O Domains that you can create depends on several factors, including the amount of CPU core and memory resources that are available in the CPU and memory repositories and the amount of CPU core and memory resources that you want to assign to each I/O Domain. However, while it is useful to know the total amount of resources are 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.

For example, assume that you have 32 cores parked in the CPU repository and 1472 GB of memory parked in the memory repository. You could therefore create I/O Domains in any of the following ways:

- One or more large I/O Domains, with each large I/O Domain using one socket's worth of resources (for example, 12 cores and 512 GB of memory)
- One or more medium I/O Domains, with each medium I/O Domain using four cores and 64 GB of memory
- One or more small I/O Domains, with each small I/O Domain using one core and 16 GB of memory

When you go through the process of creating I/O Domains, at some point, the I/O Domain Creation tool will inform you that you cannot create additional I/O Domains. This could be due to several factors, such as reaching the limit of total CPU core and memory resources in the CPU and memory repositories, reaching the limit of resources available specifically to you as a user, or reaching the limit on the number of I/O Domains allowable for this system.

Note – The following examples describe how resources might be divided up between domains using percentages to make the conceptual information easier to understand. However, you actually divide CPU core and memory resources between domains at a socket granularity or core granularity level. Refer to the *Oracle SuperCluster M6-32 Owner's Guide: Administration* 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.

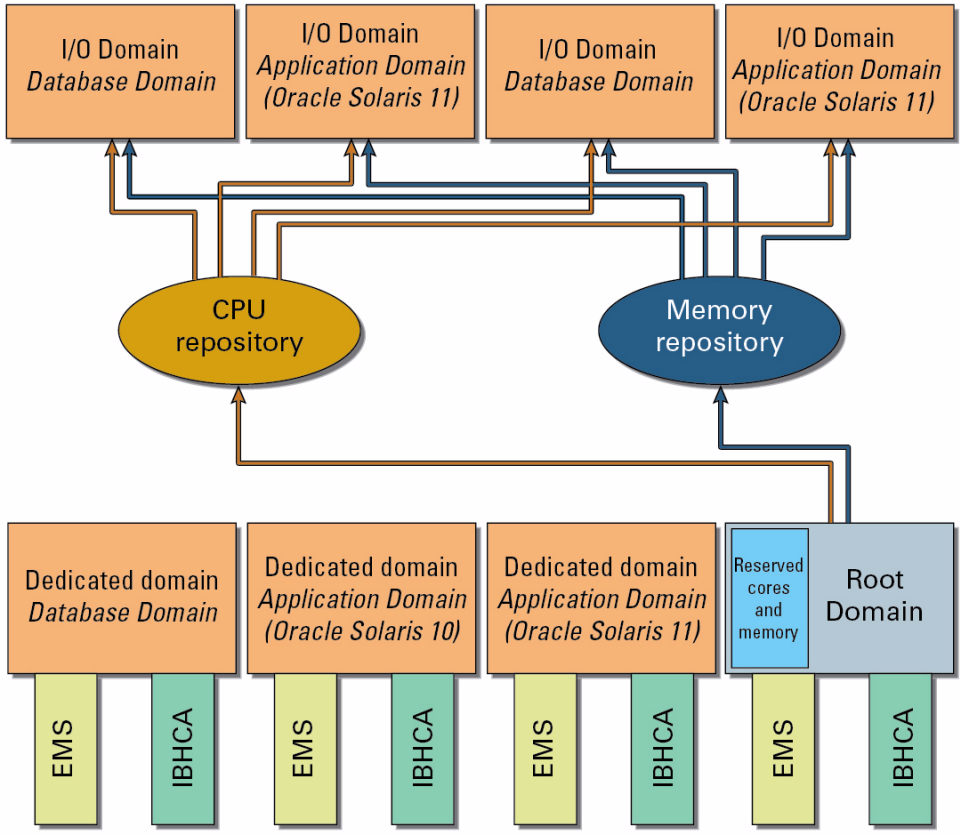

Even though dedicated domains and Root Domains are all shown as equal-sized domains in the preceding figure, that does not mean that CPU core and memory resources must be split evenly across all four domains (where each domain would get 25% of the CPU core and memory resources). Using information that you provide in the configuration worksheets, you can request different sizes of CPU core and memory resources for each domain when your SuperCluster M6-32 is initially installed.

For example, you could request that each dedicated domain have 30% of the CPU core 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 would mean 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 could also request that some of the resources from the dedicated domains be parked at the time of the initial installation of your system, which would further increase the amount of CPU core and memory resources available for I/O Domains to pull from the repositories.

You could also use the CPU/Memory tool after the initial installation to resize the amount of CPU core 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 core and memory resources used by those domains. However, you must reboot those resized dedicated domains if you change the amount of resources using the CPU/Memory tool.
- 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 core and memory resources used by those dedicated domains. You can also use the tool to park some of the CPU core and memory resources from the dedicated domains, which would park those resources in the CPU and Memory repositories, making them available for the I/O Domains. However, you must reboot those resized dedicated domains if you change the amount of resources using the CPU/Memory tool.
  - For the Root Domains, you cannot resize the amount of CPU core 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 M6-32 Owner's Guide: Administration* 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 core and memory resources (total of 90% resources allocated to dedicated domains), and the remaining 10% allocated to the Root Domain. You could then make the following changes to the resource allocation, depending on your situation:

- If you are satisfied with the amount of CPU core and memory resources allocated to the Root Domain, but you find that one dedicated domain needs more resources while another needs less, you could 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 core and memory resources allocated to the Root Domain is insufficient, you could park resources from the dedicated domains, which would park those resources in the CPU and Memory repositories, making them available for the 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 could park 10% of the resources from one or more of the dedicated domains, which would increase the amount of resources in the CPU and Memory repositories by that amount for the I/O Domains.

Understanding LDom Configurations for Extended Configuration PDomains

An extended configuration PDomain has the following characteristics:

- Two PDomains
- Two DCUs associated with each PDomain
- PDomains on one or two compute servers:
  - On a single compute server, with the two PDomains within that compute server, each with two DCUs, or
  - Split across two compute servers, with the two PDomains split across the two compute servers (one PDomain on each server), with each PDomain containing two DCUs

---

Note – Certain DCU configurations have a mixture of half-populated and fully-populated DCUs in a compute server. For extended configuration PDomains in those DCU configurations, both DCUs in each PDomain must have the same configuration, where both DCUs in the PDomain are either half-populated or fully-populated. See "CMU Overview" on page 5 for more information.

Extended configuration PDomains are available for fully-populated DCUs and half-populated DCUs:

Description	Links
Choose the LDom configuration for extended configuration PDomains on fully-populated DCUs.	"Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99
Choose the LDom configuration for extended configuration PDomains on half-populated DCUs.	"Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)

From an overall PDomain level, the extended configuration PDomains with fully-populated DCUs have the following resources:

- Two DCUs, with four CMUs in each DCU, for a total of eight CMUs
- Sixteen processors (two processors per CMU)
- 512 DIMM slots (64 DIMM slots per CMU), for a total of 8 TB (16 GB DIMMs) or 16 TB (32 GB DIMMs) of total available memory
- Eight IB HCAs

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

These topics describe the LDom configurations available for extended configuration PDomains on fully-populated DCUs.

- "LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 100
- "E4-1 LDom Configuration" on page 101
- "E4-2 LDom Configuration" on page 103
- "E4-3 LDom Configuration" on page 104
- "E4-4 LDom Configuration" on page 106

Related Information

- "Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117

LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)

This figure provides information on the available LDom configurations for the fully-populated DCU configuration with an extended configuration PDomain. The CMU no. information in the figure varies, depending on which DCU is being used in this configuration.

FIGURE: LDom Configurations for Extended Configuration PDomains (Fully-Populated DCU Configuration)

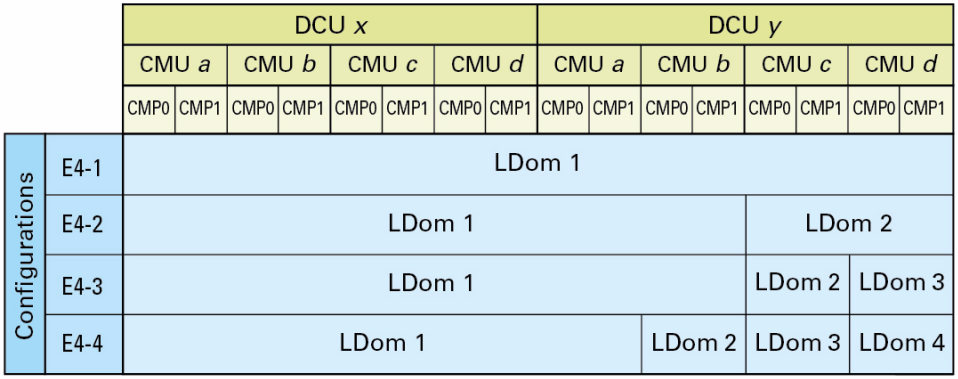

CMU No.	DCU *x* and DCU *y* DCU 0	DCU 1	DCU 2	DCU 3
CMU *a*	CMU 0	CMU 4	CMU 8	CMU 12
CMU *b*	CMU 1	CMU 5	CMU 9	CMU 13
CMU *c*	CMU 2	CMU 6	CMU 10	CMU 14
CMU *d*	CMU 3	CMU 7	CMU 11	CMU 15

Related Information

- "E4-1 LDom Configuration" on page 101
- "E4-2 LDom Configuration" on page 103
- "E4-3 LDom Configuration" on page 104
- "E4-4 LDom Configuration" on page 106

E4-1 LDom Configuration

The following tables provide information on the E4-1 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E4-1 Configuration)

Item	LDom 1
EMSs	EMS 1 through EMS 4 in both DCUs
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DCU and HDD0 through HDD7 in second DCU available for additional storage
IB HCAs	PCIe slots 2, 6, 10, and 14 in both DCUs
Empty (free) PCIe slots	1, 4, 5, 7, 8, 9, 11, 12, 13, 15, 16 in both DCUs
Default CPU/Memory Resources	100% of resources in both DCUs

TABLE: Networks (E4-1 Configuration)

		LDom 1
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU
	Standby	NET7, using P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU
	Standby	EMS 4, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU
	Standby	P0 in IB HCA in slot 14 in second DCU
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2, 6, 10, and 14 in both DCUs
	Standby	P1 in IB HCAs in slots 2, 6, 10, and 14 in both DCUs
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 6 in first DCU
	Standby	P1 in IB HCA in slot 10 in second DCU

Related Information

- "LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 100

E4-2 LDom Configuration

The following tables provide information on the E4-2 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E4-2 Configuration)

Item	LDom 1	LDom 2
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 and 2 in second DCU	EMS 3 and EMS 4 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 through HDD3 in second DCU available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • HDD6 and HDD7 in second DCU available for additional storage
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slots 2 and 6 in second DCU	PCIe slots 10 and 14 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 1, 4, 5, 7, and 8 in second DCU	PCI slots 9, 11, 12, 13, 15 and 16 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 50% of resources in second DCU	50% of resources in second DCU

TABLE: Networks (E4-2 Configuration)

		LDom 1	LDom 2
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU
	Standby	NET3, using P3 in 1GbE NIC in first DCU	NET3, using VNET through P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 3, P0 in second DCU
	Standby	EMS 2, P1 in second DCU	EMS 4, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 10 in second DCU
	Standby	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 14 in second DCU

TABLE: Networks (E4-2 Configuration)

		LDom 1	LDom 2
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCAs in slots 2 and 6 in second DCU	P0 in IB HCAs in slots 10 and 14 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCAs in slots 2 and 6 in second DCU	P1 in IB HCAs in slots 10 and 14 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	P1 in IB HCA in slot 6 second DCU	P1 in IB HCA in slot 14 in second DCU

Related Information

- "LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 100

E4-3 LDom Configuration

The following tables provide information on the E4-3 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E4-3 Configuration)

Item	LDom 1	LDom 2	LDom 3
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 and 2 in second DCU	EMS 3 in second DCU	EMS 4 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 through HDD3 in second DCU available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD6 and HDD7 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E4-3 Configuration) *(Continued)*

Item	LDom 1	LDom 2	LDom 3
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slots 2 and 6 in second DCU	PCIe slot 10 in second DCU	PCIe slot 14 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 1, 4, 5, 7, and 8 in second DCU	11, 13, and 16 in second DCU	PCIe slots 9, 12, and 15 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 50% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

TABLE: Networks (E4-3 Configuration)

		LDom 1	LDom 2	LDom 3
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU	NET0, using VNET through P2 in 1GbE NIC in second DCU
	Standby	NET3, using P3 in 1GbE NIC in first DCU	NET1, using VNET through P1 in 1GbE NIC in second DCU	NET1, using VNET through P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 3, P0 in second DCU	EMS 4, P0 in second DCU
	Standby	EMS 2, P1 in second DCU	EMS 3, P1 in second DCU	EMS 4, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU
	Standby	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCAs in slots 2 and 6 in second DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCAs in slots 2 and 6 in second DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	P1 in IB HCA in slot 6 second DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU

Related Information

- "LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 100

E4-4 LDom Configuration

The following tables provide information on the E4-4 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E4-4 Configuration)

	LDom 1	LDom 2	LDom 3	LDom 4
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 in second DCU	EMS 2 in second DCU	EMS 3 in second DCU	EMS 4 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 and HDD1 in second DCU available for additional storage	• HDD2 and HDD3 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD6 and HDD7 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slot 2 in second DCU	PCIe slot 6 in second DCU	PCIe slot 10 in second DCU	PCIe slot 14 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 5 and 8 in second DCU	PCIe slots 1, 4, and 7 in second DCU	PCIe slots 11, 13, and 16 in second DCU	PCIe slots 9, 12, and 15 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

TABLE: Networks (E4-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 4
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P2 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU	NET0, using VNET through P2 in 1GbE NIC in second DCU
	Standby	NET1, using P1 in 1GbE NIC in first DCU	NET1, using VNET through P3 in 1GbE NIC in first DCU	NET1, using VNET through P1 in 1GbE NIC in second DCU	NET1, using VNET through P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 2, P0 in second DCU	EMS 3, P0 in second DCU	EMS 4, P0 in second DCU
	Standby	EMS 1, P1 in second DCU	EMS 2, P1 in second DCU	EMS 3, P1 in second DCU	EMS 4, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU
	Standby	P0 in IB HCA in slot 2 in second DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCA in slot 2 in second DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCA in slots 2 in second DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	P1 in IB HCA in slot 2 second DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 10 in second DCU	P1 in IB HCA in slot 14 in second DCU

Related Information

- "LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 100

Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)

From an overall PDomain level, the extended configuration PDomains with half-populated DCUs have the following resources:

- Two DCUs, with two CMUs in each DCU, for a total of four CMUs
- Eight processors (two processors per CMU)
- 256 DIMM slots (64 DIMM slots per CMU), for a total of 4 TB (16 GB DIMMs) or 8 TB (32 GB DIMMs) of total available memory
- Eight IB HCAs

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

These topics describe the LDom configurations available for extended configuration PDomains on half-populated DCUs.

- "LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108
- "E2-1 LDom Configuration" on page 109
- "E2-2 LDom Configuration" on page 111
- "E2-3 LDom Configuration" on page 112
- "E2-4 LDom Configuration" on page 114

Related Information

- "Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)

This figure provides information on the available LDom configurations for the half-populated DCU configuration with an extended configuration PDomain. The CMU no. information in the figure varies, depending on which DCU is being used in this configuration.

FIGURE: LDom Configurations for Extended Configuration PDomains (Half-Populated DCU Configuration)

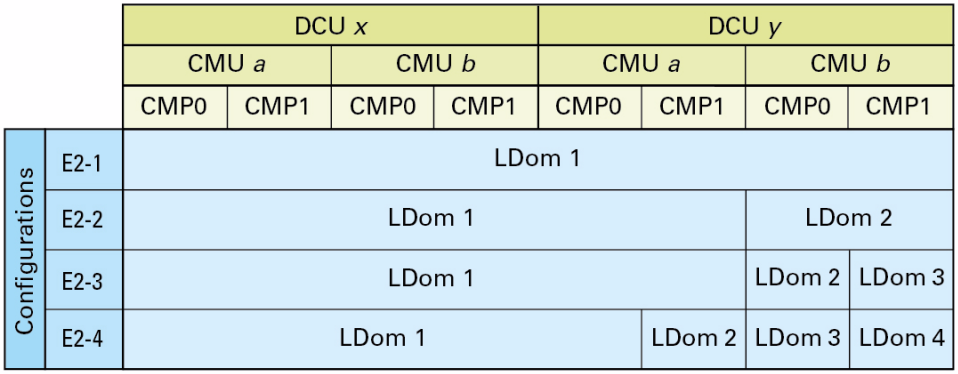

CMU No.	DCU *x* and DCU *y* DCU 0	DCU 1	DCU 2	DCU 3
CMU *a*	CMU 0	CMU 4	CMU 8	CMU 12
CMU *b*	CMU 3	CMU 7	CMU 11	CMU 15

Related Information

- "E2-1 LDom Configuration" on page 109
- "E2-2 LDom Configuration" on page 111
- "E2-3 LDom Configuration" on page 112
- "E2-4 LDom Configuration" on page 114

E2-1 LDom Configuration

The following tables provide information on the E2-1 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E2-1 Configuration)

Item	LDom 1
EMSs	EMS 1 through EMS 4 in both DCUs
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DCU and HDD0 through HDD7 in second DCU available for additional storage
IB HCAs	PCIe slots 2, 6, 10, and 14 in both DCUs
Empty (free) PCIe slots	1, 4, 5, 7, 8, 9, 11, 12, 13, 15, 16 in both DCUs
Default CPU/Memory Resources	100% of resources in both DCUs

TABLE: Networks (E2-1 Configuration)

		LDom 1
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU
	Standby	NET7, using P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU
	Standby	EMS 3, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU
	Standby	P0 in IB HCA in slot 10 in second DCU
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2, 6, 10, and 14 in both DCUs
	Standby	P1 in IB HCAs in slots 2, 6, 10, and 14 in both DCUs
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 6 in first DCU
	Standby	P1 in IB HCA in slot 14 in second DCU

Related Information

- "LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

E2-2 LDom Configuration

The following tables provide information on the E2-2 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E2-2 Configuration)

Item	LDom 1	LDom 2
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 and 2 in second DCU	EMS 3 and EMS 4 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 through HDD3 in second DCU available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • HDD6 and HDD7 in second DCU available for additional storage
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slots 2 and 6 in second DCU	PCIe slots 10 and 14 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 1, 4, 5, 7, and 8 in second DCU	PCI slots 9, 11, 12, 13, 15 and 16 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 50% of resources in second DCU	50% of resources in second DCU

TABLE: Networks (E2-2 Configuration)

		LDom 1	LDom 2
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU
	Standby	NET3, using P3 in 1GbE NIC in first DCU	NET3, using VNET through P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 4, P0 in second DCU
	Standby	EMS 2, P1 in second DCU	EMS 3, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 14 in second DCU
	Standby	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 10 in second DCU

TABLE: Networks (E2-2 Configuration)

		LDom 1	LDom 2
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCAs in slots 2 and 6 in second DCU	P0 in IB HCAs in slots 10 and 14 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCAs in slots 2 and 6 in second DCU	P1 in IB HCAs in slots 10 and 14 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	P1 in IB HCA in slot 6 second DCU	P1 in IB HCA in slot 10 in second DCU

Related Information

- "LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

E2-3 LDom Configuration

The following tables provide information on the E2-3 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E2-3 Configuration)

Item	LDom 1	LDom 2	LDom 3
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 and 2 in second DCU	EMS 4 in second DCU	EMS 3 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 through HDD3 in second DCU available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD6 and HDD7 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E2-3 Configuration) *(Continued)*

Item	LDom 1	LDom 2	LDom 3
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slots 2 and 6 in second DCU	PCIe slot 14 in second DCU	PCIe slot 10 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 1, 4, 5, 7, and 8 in second DCU	13, 15 and 16 in second DCU	9, 11 and 12 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 50% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

TABLE: Networks (E2-3 Configuration)

		LDom 1	LDom 2	LDom 3
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU	NET0, using VNET through P2 in 1GbE NIC in second DCU
	Standby	NET3, using P3 in 1GbE NIC in first DCU	NET1, using VNET through P1 in 1GbE NIC in second DCU	NET1, using VNET through P3 in 1GbE NIC in second DCU
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 4, P0 in second DCU	EMS 3, P0 in second DCU
	Standby	EMS 2, P1 in second DCU	EMS 4, P1 in second DCU	EMS 3, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU
	Standby	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU

TABLE: Networks (E2-3 Configuration)

		LDom 1	LDom 2	LDom 3
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCAs in slots 2 and 6 in second DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCAs in slots 2 and 6 in second DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	P1 in IB HCA in slot 6 second DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU

Related Information

- "LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

E2-4 LDom Configuration

The following tables provide information on the E2-4 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (E2-4 Configuration)

	LDom 1	LDom 2	LDom 3	LDom 4
EMSs	• EMS 1 through EMS 4 in first DCU • EMS 1 in second DCU	EMS 2 in second DCU	EMS 4 in second DCU	EMS 3 in second DCU
Hard Drives	• HDD0 and HDD1 in first DCU as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 in first DDCU and HDD0 and HDD1 in second DCU available for additional storage	• HDD2 and HDD3 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD4 and HDD5 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage	• HDD6 and HDD7 in second DCU as boot and `/u01` primary and mirrored drives • No drives available for additional storage
IB HCAs	• PCIe slots 2, 6, 10, and 14 in first DCU • PCIe slot 2 in second DCU	PCIe slot 6 in second DCU	PCIe slot 14 in second DCU	PCIe slot 10 in second DCU
Empty (free) PCIe slots	• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU • PCIe slots 1 and 4 in second DCU	PCIe slots 5, 7 and 8 in second DCU	PCIe slots 13, 15 and 16 in second DCU	PCIe slots 9, 11 and 12 in second DCU
Default CPU/Memory Resources	• 100% of resources in first DCU • 25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

TABLE: Networks (E2-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 4
Management Network	Active	NET0, using P0 in 1GbE NIC in first DCU	NET0, using VNET through P2 in 1GbE NIC in first DCU	NET0, using VNET through P0 in 1GbE NIC in second DCU	NET0, using VNET through P2 in 1GbE NIC in second DCU
	Standby	NET1, using P1 in 1GbE NIC in first DCU	NET1, using VNET through P3 in 1GbE NIC in first DCU	NET1, using VNET through P1 in 1GbE NIC in second DCU	NET1, using VNET through P3 in 1GbE NIC in second DCU

TABLE: Networks (E2-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 4
10GbE Client Access Network	Active	EMS 1, P0 in first DCU	EMS 2, P0 in second DCU	EMS 4, P0 in second DCU	EMS 3, P0 in second DCU
	Standby	EMS 1, P1 in second DCU	EMS 2, P1 in second DCU	EMS 4, P1 in second DCU	EMS 3, P1 in second DCU
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2 in first DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU
	Standby	P0 in IB HCA in slot 2 in second DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
IB Network: Exadata Private Network (DB Domains)	Active	• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P0 in IB HCA in slot 2 in second DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU • P1 in IB HCA in slots 2 in second DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2 in first DCU	P0 in IB HCA in slot 6 in second DCU	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	P1 in IB HCA in slot 2 second DCU	P1 in IB HCA in slot 6 in second DCU	P1 in IB HCA in slot 14 in second DCU	P1 in IB HCA in slot 10 in second DCU

Related Information

- "LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

Understanding LDom Configurations for Base Configuration PDomains

A base configuration has the following characteristics:

- Two or four PDomains
- One DCU associated with each PDomain
- PDomains on one or two compute servers, with one of these situations:
  - On a single compute server, with the two or four PDomains within that compute server, each with one DCU
  - Split across two compute servers, with either two PDomains split across the two compute servers (one PDomain on each compute server) or four PDomains split across the two compute servers (two PDomains on each compute server), with each PDomain containing a single DCU

Base configuration PDomains are available for fully-populated DCUs and half-populated DCUs:

Description	Links
Choose the LDom configuration for base configuration PDomains on fully-populated DCUs.	"Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 117
Choose the LDom configuration for base configuration PDomains on half-populated DCUs.	"Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 124

Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)

From an overall PDomain level, the base configuration PDomains with fully-populated DCUs have the following resources:

- One DCU, with four CMUs in the DCU
- Eight processors (two processors per CMU)
- 256 DIMM slots (64 DIMM slots per CMU), for a total of 4 TB (16 GB DIMMs) or 8 TB (32 GB DIMMs) of total available memory
- Four IB HCAs

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

These topics describe the LDom configurations available for base configuration PDomains of fully-populated DCUs:

- "LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 118
- "B4-1 LDom Configuration" on page 119

- "B4-2 LDom Configuration" on page 120
- "B4-3 LDom Configuration" on page 121
- "B4-4 LDom Configuration" on page 123

Related Information

- "Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)" on page 99

LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)

This figure provides information on the available LDom configurations for the fully-populated DCU configuration with a base configuration PDomain. The CMU no. information in the figure varies, depending on which DCU is being used in this configuration.

FIGURE: LDom Configurations for Base Configuration PDomains (Fully-Populated DCU Configuration)

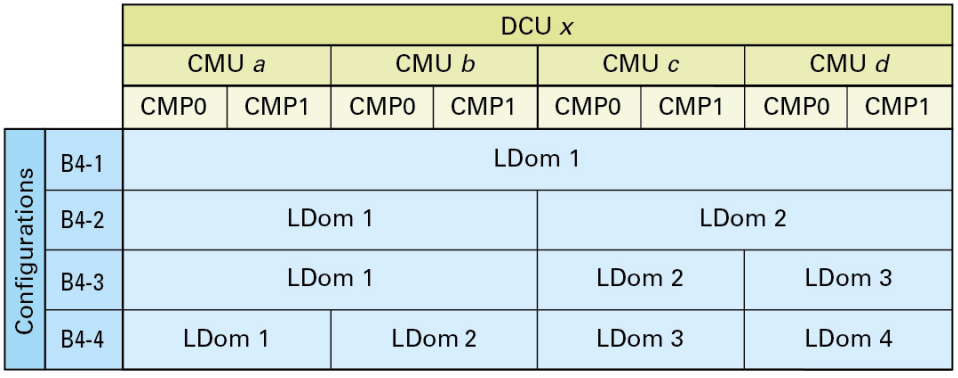

CMU No.	DCU *x* DCU 0	DCU 1	DCU 2	DCU 3
CMU *a*	CMU 0	CMU 4	CMU 8	CMU 12
CMU *b*	CMU 1	CMU 5	CMU 9	CMU 13
CMU *c*	CMU 2	CMU 6	CMU 10	CMU 14
CMU *d*	CMU 3	CMU 7	CMU 11	CMU 15

Related Information

- "B4-1 LDom Configuration" on page 119
- "B4-2 LDom Configuration" on page 120
- "B4-3 LDom Configuration" on page 121
- "B4-4 LDom Configuration" on page 123

B4-1 LDom Configuration

The following tables provide information on the B4-1 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B4-1 Configuration)

Item	LDom 1
EMSs	EMS 1, EMS 2, EMS 3 and EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 available for additional storage
IB HCAs	PCIe slots 2, 6, 10, and 14
Empty (free) PCIe slots	1, 4, 5, 7, 8, 9, 11, 12, 13, 15, 16
Default CPU/Memory Resources	100%

TABLE: Networks (B4-1 Configuration)

		LDom 1
Management Network	Active	NET0, using P0 in 1GbE NIC
	Standby	NET3, using P3 in 1GbE NIC
10GbE Client Access Network	Active	EMS 1, P0
	Standby	EMS 4, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2
	Standby	P0 in IB HCA in slot 14

TABLE: Networks (B4-1 Configuration)

		LDom 1
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2, 6, 10, and 14
	Standby	P1 in IB HCAs in slots 2, 6, 10, and 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 6
	Standby	P1 in IB HCA in slot 10

Related Information

- "LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 118

B4-2 LDom Configuration

The following tables provide information on the B4-2 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B4-2 Configuration)

Item	LDom 1	LDom 2
EMSs	EMS 1 and EMS 2	EMS 3 and EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 and HDD3 available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • HDD6 and HDD7 available for additional storage
IB HCAs	PCIe slot 2 and 6	PCIe slot 10 and 14
Empty (free) PCIe slots	1, 4, 5, 7 and 8	9, 11, 12, 13, 15 and 16
Default CPU/Memory Resources	50%	50%

TABLE: Networks (B4-2 Configuration)

		LDom 1	LDom 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	EMS 1, P0	EMS 3, P0
	Standby	EMS 2, P1	EMS 4, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 10
	Standby	P0 in IB HCA in slot 6	P0 in IB HCA in slot 14
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2 and 6	P0 in IB HCA in slot 10 and 14
	Standby	P1 in IB HCAs in slots 2 and 6	P1 in IB HCA in slot 10 and 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 10
	Standby	P1 in IB HCA in slot 6	P1 in IB HCA in slot 14

Related Information

- "LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 118

B4-3 LDom Configuration

The following tables provide information on the B4-3 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B4-3 Configuration)

Item	LDom 1	LDom 2	LDom 3
EMSs	EMS 1 and EMS 2	EMS 3	EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 and HDD3 available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD6 and HDD7 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage
IB HCAs	PCIe slots 2 and 6	PCIe slot 10	PCIe slot 14
Empty (free) PCIe slots	1, 4, 5, 7 and 8	11, 13, and 16	9, 12, and 15
Default CPU/Memory Resources	50%	25%	25%

TABLE: Networks (B4-3 Configuration)

		LDom 1	LDom 2	LDom 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	EMS 1, P0	EMS 3, P0	EMS 4, P0
	Standby	EMS 2, P1	EMS 3, P1	EMS 4, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14
	Standby	P0 in IB HCA in slot 6	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2 and 6	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
	Standby	P1 in IB HCAs in slots 2 and 6	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
	Standby	P1 in IB HCA in slot 6	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14

Related Information

- "LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 118

B4-4 LDom Configuration

The following tables provide information on the B4-4 PDomain configuration for the fully-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B4-4 Configuration)

	LDom 1	LDom 2	LDom 3	LDom 4
EMSs	EMS 1	EMS 2	EMS 3	EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD2 and HDD3 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD6 and HDD7 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage
IB HCAs	PCIe slot 2	PCIe slot 6	PCIe slot 10	PCIe slot 14
Empty (free) PCIe slots	5 and 8	1, 4, and 7	11, 13, and 16	9, 12, and 15
Default CPU/Memory Resources	25%	25%	25%	25%

TABLE: Networks (B4-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 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	EMS 1, P0	EMS 2, P0	EMS 3, P0	EMS 4, P0
	Standby	EMS 1, P1	EMS 2, P1	EMS 3, P1	EMS 4, P1

TABLE: Networks (B4-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 4
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14
	Standby	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
	Standby	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 10	P0 in IB HCA in slot 14
	Standby	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 10	P1 in IB HCA in slot 14

Related Information

- "LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)" on page 118

Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)

From an overall PDomain level, the base configuration PDomains with half-populated DCUs have the following resources:

- One DCU, with two CMUs in the DCU
- Four processors (two processors per CMU)
- 128 DIMM slots (64 DIMM slots per CMU), for a total of 2 TB (16 GB DIMMs) or 4 TB (32 GB DIMMs) of total available memory
- Four IB HCAs

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

These topics describe the LDom configurations available for base configuration PDomains on half-populated DCUs.

- "LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 125

- "B2-1 LDom Configuration" on page 126
- "B2-2 LDom Configuration" on page 127
- "B2-3 LDom Configuration" on page 128
- "B2-4 LDom Configuration" on page 130

Related Information

- "Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)" on page 108

LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)

This figure provides information on the available LDom configurations for the half-populated DCU configuration with a base configuration PDomain. The CMU no. information in the figure varies, depending on which DCU is being used in this configuration.

FIGURE: LDom Configurations for Base Configuration PDomains (Half-Populated DCU Configuration)

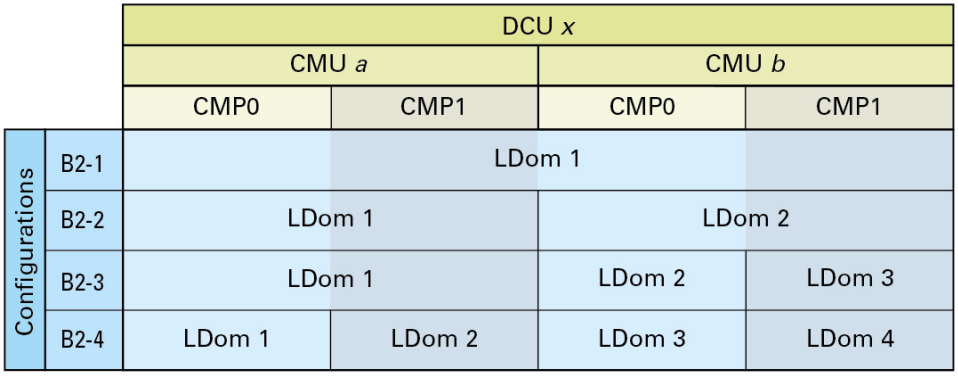

	DCU *x*			
CMU No.	**DCU 0**	**DCU 1**	**DCU 2**	**DCU 3**
CMU *a*	CMU 0	CMU 4	CMU 8	CMU 12
CMU *b*	CMU 3	CMU 7	CMU 11	CMU 15

Related Information

- "B2-1 LDom Configuration" on page 126
- "B2-2 LDom Configuration" on page 127
- "B2-3 LDom Configuration" on page 128
- "B2-4 LDom Configuration" on page 130

B2-1 LDom Configuration

The following tables provide information on the B2-1 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B2-1 Configuration)

Item	LDom 1
EMSs	EMS 1, EMS 2, EMS 3 and EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 through HDD7 available for additional storage
IB HCAs	PCIe slots 2, 6, 10, and 14
Empty (free) PCIe Slots	1, 4, 5, 7, 8, 9, 11, 12, 13, 15, 16
Default CPU/Memory Resources	100%

TABLE: Networks (B2-1 Configuration)

		LDom 1
Management Network	Active	NET0, using P0 in 1GbE NIC
	Standby	NET3, using P3 in 1GbE NIC
10GbE Client Access Network	Active	EMS 1, P0
	Standby	EMS 3, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2
	Standby	P0 in IB HCA in slot 10

TABLE: Networks (B2-1 Configuration)

		LDom 1
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2, 6, 10, and 14
	Standby	P1 in IB HCAs in slots 2, 6, 10, and 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 6
	Standby	P1 in IB HCA in slot 14

Related Information

- "LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 125

B2-2 LDom Configuration

The following tables provide information on the B2-2 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B2-2 Configuration)

Item	LDom 1	LDom 2
EMSs	EMS 1 and EMS 2	EMS 3 and EMS 4
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 and HDD3 available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • HDD6 and HDD7 available for additional storage
IB HCAs	PCIe slot 2 and 6	PCIe slot 10 and 14
Empty (free) PCIe slots	1, 4, 5, 7 and 8	9, 11, 12, 13, 15 and 16
Default CPU/Memory Resources	50%	50%

TABLE: Networks (B2-2 Configuration)

		LDom 1	LDom 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	EMS 1, P0	EMS 4, P0
	Standby	EMS 2, P1	EMS 3, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 14
	Standby	P0 in IB HCA in slot 6	P0 in IB HCA in slot 10
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2 and 6	P0 in IB HCA in slot 10 and 14
	Standby	P1 in IB HCAs in slots 2 and 6	P1 in IB HCA in slot 10 and 14
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 14
	Standby	P1 in IB HCA in slot 6	P1 in IB HCA in slot 10

Related Information

- "LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 125

B2-3 LDom Configuration

The following tables provide information on the B2-3 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B2-3 Configuration)

Item	LDom 1	LDom 2	LDom 3
EMSs	EMS 1 and EMS 2	EMS 4	EMS 3
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • HDD2 and HDD3 available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD6 and HDD7 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage
IB HCAs	PCIe slots 2 and 6	PCIe slot 14	PCIe slot 10
Empty (free) PCIe slots	1, 4, 5, 7 and 8	13, 15 and 16	9, 11 and 12
Default CPU/Memory Resources	50%	25%	25%

TABLE: Networks (B2-3 Configuration)

		LDom 1	LDom 2	LDom 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	EMS 1, P0	EMS 4, P0	EMS 3, P0
	Standby	EMS 2, P1	EMS 4, P1	EMS 3, P1
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 14	P1 in IB HCA in slot 10
	Standby	P0 in IB HCA in slot 6	P0 in IB HCA in slot 14	P0 in IB HCA in slot 10
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCAs in slots 2 and 6	P0 in IB HCA in slot 14	P0 in IB HCA in slot 10
	Standby	P1 in IB HCAs in slots 2 and 6	P1 in IB HCA in slot 14	P1 in IB HCA in slot 10
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 14	P0 in IB HCA in slot 10
	Standby	P1 in IB HCA in slot 6	P1 in IB HCA in slot 14	P1 in IB HCA in slot 10

Related Information

- "LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 125

B2-4 LDom Configuration

The following tables provide information on the B2-4 PDomain configuration for the half-populated DCU configuration.

TABLE: EMSs, PCIe Slots and Cards, and CPU/Memory Resources (B2-4 Configuration)

	LDom 1	LDom 2	LDom 3	LDom 4
EMSs	EMS 1	EMS 2	EMS 4	EMS 3
Hard Drives	• HDD0 and HDD1 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD2 and HDD3 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD4 and HDD5 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage	• HDD6 and HDD7 as boot and `/u01` primary and mirrored drives • No hard drives available for additional storage
IB HCAs	PCIe slot 2	PCIe slot 6	PCIe slot 14	PCIe slot 10
Empty (free) PCIe slots	1 and 4	5, 7 and 8	13, 15 and 16	9, 11 and 12
Default CPU/Memory Resources	25%	25%	25%	25%

TABLE: Networks (B2-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 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	EMS 1, P0	EMS 2, P0	EMS 4, P0	EMS 3, P0
	Standby	EMS 1, P1	EMS 2, P1	EMS 4, P1	EMS 3, P1

TABLE: Networks (B2-4 Configuration)

		LDom 1	LDom 2	LDom 3	LDom 4
IB Network: Storage Private Network (DB or App Domains)	Active	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 14	P1 in IB HCA in slot 10
	Standby	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 14	P0 in IB HCA in slot 102
IB Network: Exadata Private Network (DB Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 14	P0 in IB HCA in slot 10
	Standby	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 14	P1 in IB HCA in slot 10
IB Network: Oracle Solaris Cluster Private Network (App Domains)	Active	P0 in IB HCA in slot 2	P0 in IB HCA in slot 6	P0 in IB HCA in slot 14	P0 in IB HCA in slot 10
	Standby	P1 in IB HCA in slot 2	P1 in IB HCA in slot 6	P1 in IB HCA in slot 142	P1 in IB HCA in slot 10

Related Information

- "LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)" on page 125

Understanding Clustering Software

These topics describe clustering software.

- "Cluster Software Overview" on page 133
- "Cluster Software for the Database Domain" on page 134
- "Cluster Software for the Oracle Solaris Application Domains" on page 135

Cluster Software Overview

Clustering software is typically used on multiple interconnected servers so that they appear as if they are one server to end users and applications. For SuperCluster M6-32, clustering software is used to cluster certain LDoms on the compute servers together with the same type of domain. The benefits of clustering software include the following:

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

Related Information

- "Cluster Software for the Database Domain" on page 134
- "Cluster Software for the Oracle Solaris Application Domains" on page 135

Cluster Software for the Database Domain

Oracle RAC enables the clustering of the Oracle Database on the Database Domain. Oracle RAC uses Oracle Clusterware for the infrastructure to cluster the Database Domains on the compute servers together.

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

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

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

Related Information

- "Cluster Software Overview" on page 133
- "Cluster Software for the Oracle Solaris Application Domains" on page 135

Cluster Software for the Oracle Solaris Application Domains

The Oracle Solaris Cluster software is an optional clustering tool used for the Oracle Solaris Application Domains. On Oracle SuperCluster M6-32, the Oracle Solaris Cluster software is used to cluster the Oracle Solaris Application Domains on the compute servers together.

Related Information

- "Cluster Software Overview" on page 133
- "Cluster Software for the Database Domain" on page 134

Understanding System Administration Resources

These topics describe the architecture and administration resources of SuperCluster M6-32:

- "Oracle ILOM Overview" on page 137
- "Understanding Platform-Specific Oracle ILOM Features" on page 138
- "Oracle Solaris OS Overview" on page 139
- "OpenBoot Overview" on page 140
- "Oracle ILOM Remote Console Plus Overview" on page 140
- "Oracle Hardware Management Pack Overview" on page 141
- "Time Synchronization and NTP Service" on page 141
- "SNMP Service" on page 142
- "Multidomain Extensions to Oracle ILOM MIBs" on page 142
- "LDAP/SSL" on page 144
- "Active Directory" on page 144

Oracle ILOM Overview

Oracle ILOM is system management firmware that is preinstalled on some Oracle servers. Oracle ILOM enables you to actively manage and monitor components installed in your compute server. Oracle ILOM provides a browser-based interface and a CLI, as well as SNMP and IPMI interfaces.

The Oracle ILOM SP runs independently of the compute server and regardless of the compute server power state as long as AC power is connected to the server. When you connect the compute server to AC power, the ILOM service processor immediately starts up and begins monitoring the compute server. All environmental monitoring and control are handled by Oracle ILOM.

The -> prompt indicates that you are interacting with the Oracle ILOM SP directly. This prompt is the first prompt you see when you log in to the compute server through the SER MGT port or NET MGT port, regardless of the host's power state.

You can also access the Oracle ILOM SP prompt from the OpenBoot `ok` prompt, or from the Oracle Solaris shell prompt, provided the system console is configured to be accessible through the SER MGT and NET MGT ports.

For more information about how to work with Oracle ILOM features that are common to all platforms managed by Oracle ILOM, refer to the Oracle ILOM documentation.

For more information about Oracle ILOM features that are specific to the compute server, see "Understanding Platform-Specific Oracle ILOM Features" on page 138.

Related Information

- Oracle ILOM Documentation Library at: `http://www.oracle.com/goto/ILOM/docs`

Understanding Platform-Specific Oracle ILOM Features

Oracle ILOM operates on many platforms, supporting features that are common to all platforms. Some Oracle ILOM features belong to only a subset of platforms. These topics describe the difference between Oracle ILOM features supported on this server and the common set of features, which are described in the Oracle ILOM base documentation.

- "Server-Specific and New Oracle ILOM Features and Requirements" on page 138
- "Unsupported Oracle ILOM Features" on page 139

Server-Specific and New Oracle ILOM Features and Requirements

Oracle ILOM has the following requirements and supports the following features on this server:

- You can create up to 60 user accounts in Oracle ILOM. As many as 25 concurrent user sessions (SSH or web) are supported per SP.

- Certain Oracle ILOM tasks can be performed for the platform and for any available PDomain. This means that user roles must be properly assigned at the platform or domain level, and that specific commands must be provided for either the platform or PDomain. For information about the commands that must be executed at the domain level, refer to the *Oracle SuperCluster M6-32 Owner's Guide: Administration*.
- The Oracle ILOM MIB file `SUN-ILOM-CONTROL-MIB`, which provides objects for configuring and managing all Oracle ILOM functions has been modified to include a hostgroups tables to support per-domain user roles. Extensions to the Oracle ILOM MIBs are also provided. Refer to the *Oracle SuperCluster M6-32 Owner's Guide: Administration* for more information.
- A new POST diagnostics hardware change property (`trigger`) is the *default setting* for the server, and causes POST to run each time the server is AC power-cycled. If you want to ensure that POST runs on each power cycle, the trigger property must be set to `power-on-reset` or `all-resets`. For more information on enabling SPARC diagnostics to run at boot, refer to the Oracle ILOM documentation.

Related Information

- Oracle ILOM Documentation Library at: `http://www.oracle.com/goto/ILOM/docs`
- "Oracle ILOM Overview" on page 137

Unsupported Oracle ILOM Features

Among the Oracle ILOM features commonly supported on other Oracle servers, Oracle ILOM does *not* support the following features on this server:

- `/SP/policy` is not available on this compute server.
- The POST diagnostics `user-reset` trigger is not supported.
- The Storage Redirection CLI is not supported for use with Oracle ILOM 3.2.

Oracle Solaris OS Overview

The Oracle Solaris OS includes commands and other software resources to use for server administration. For an introduction to management tools in the Oracle Solaris release, refer to *System Administration Guide: Basic Administration* in the Oracle Solaris documentation collection.

The Oracle Solaris software includes the Oracle VTS software. Oracle VTS tests and validates Oracle hardware by verifying the connectivity and functionality of hardware devices, controllers, and peripherals.

In addition to the Oracle VTS information in the Oracle Solaris documentation, Oracle VTS documentation collections are available at:

http://www.oracle.com/goto/VTS/docs

Related Information

- "OpenBoot Overview" on page 140

OpenBoot Overview

The OpenBoot firmware starts the OS, validates installed hardware, and can be used for other server administration tasks below the OS level. For more information about OpenBoot commands, refer to the *OpenBoot 4.x Command Reference Manual* in the Oracle Solaris documentation collection under "Important Information From Previous Releases" at:

http://docs.oracle.com/cd/E23824_01/

Related Information

- "Oracle Solaris OS Overview" on page 139

Oracle ILOM Remote Console Plus Overview

Oracle ILOM Remote Console Plus is a Java application that enables you to remotely redirect and control the following devices on the host server. This group of devices is commonly abbreviated as KVMS.

- Keyboard
- Video display
- Mouse

- Serial connection
- Storage devices or images (CD/DVD)

Related Information

- *Oracle SuperCluster M6-32 Owner's Guide: Administration*
- *Oracle ILOM Administrator's Guide for Configuration and Maintenance*

Oracle Hardware Management Pack Overview

Oracle Hardware Management Pack provides tools for managing and configuring Oracle servers from the host operating system. To use these tools, you must install the software on your compute server. After installing the software, you are able to perform the following server management tasks:

- Monitor Oracle hardware with the host IP address.
- Monitor storage devices.
- Query, update, and validate firmware versions on supported SAS storage devices.
- Restore, set, and view Oracle ILOM configuration settings.
- Use the IPMI tool to access and manage Oracle servers.

You can download the Oracle Hardware Management Pack software from:

`http://support.oracle.com`

You can download HMP documentation from:

`http://www.oracle.com/goto/OHMP/docs`

Time Synchronization and NTP Service

When PDomains are powered on, their clocks synchronize to the NTP server when the system is configured to listen to NTP multicast (the default for the current Oracle Solaris OS). If the PDomains and SPs use the same NTP server, events logged in the Oracle Solaris OS and on the SP can be correlated based on their time stamps. If the PDomains and SPs use different NTP servers, their times might drift, and correlating

log files could become difficult. If you connect a domain to an NTP server other than the one used by the SP, ensure that both are low-stratum NTP servers that provide the same degree of accuracy.

Related Information

- *Oracle ILOM Administrator's Guide for Configuration and Maintenance*

SNMP Service

The SNMP agent is preinstalled on this server and runs on Oracle ILOM, so all SNMP management occurs through Oracle ILOM. To manage the server using SNMP, you must install an SNMP client application (for example, HMP, Openview, or Tivoli).

The SNMP agent is active only on the Active SP. In the event of failover, the SNMP agent is restarted on the newly assigned Active SP.

Related Information

- *Oracle SuperCluster M6-32 Owner's Guide: Administration*
- *Oracle ILOM Protocol Management Reference for SNMP, IPMI, CIM, WS-MAN*

Multidomain Extensions to Oracle ILOM MIBs

The Oracle ILOM MIB files `SUN-ILOM-CONTROL-MIB` and `SUN-HW-CTRL-MIB` have been extended to provide a multidomain version of existing Oracle ILOM MIB objects. The MIB files containing the extensions are available under `/SP/services/snmp/mibs` in the Oracle ILOM CLI.

Refer to the *Oracle ILOM Protocol Management Reference for SNMP, IPMI, CIM, MS-MAN* for examples the single domain versions of these objects.

In the case of `SUN-ILOM-CONTROL-MIB`, the extensions correspond to the following multidomain tables.

- ilomCtrlSPARCDiagsTable provides a listing of properties for configuring SPARC-specific diagnostics for each domain.
- ilomCtrlSPARCHostControlTable provides a listing of properties for configuring SPARC-specific host software for each domain.
- ilomCtrlSPARCBootModeTable provides a listing of properties for configuring SPARC-specific boot mode features for each domain.
- ilomCtrlSPARCKeySwitchTable provides a listing of properties for controlling the SPARC-specific virtual key switch for each domain.
- ilomCtrlSPARCDomainDCUTable provides a listing of properties for configuring SPARC-specific assigned DCUs for each domain.
- ilomCtrlNetInterconnectGlobalCfgTable provides a listing of interconnect configuration entries that can be modified in a multi-domain system.
- ilomCtrlNetInterconnectGlobalOperTable provides a listing of interconnect configuration entries that are read-only in a multidomain system.

In addition, the following SPARC-specific chassis-level diagnostics are available:

- ilomCtrlSPARCChassisDiagsMode
- ilomCtrlSPARCChassisDiagsTrigger
- ilomCtrlSPARCChassisDiagsHWChangeLevel
- ilomCtrlSPARCChassisDiagsPowerOnLevel
- ilomCtrlSPARCChassisDiagsErrorRestLevel

Refer to the SUN-ILOM-CONTROL-MIB file for detailed information about these objects.

In the case of SUN-HW-CTRL-MIB, the extensions correspond to the following multidomain tables:

- sunHwCtrlDomainPowerMgmtConsumptionTable provides a listing of power consumption entries that can be displayed and modified in a multi-domain system.
- sunHwCtrlDomainPowerMgmtBudgetSettingsTable provides a listing of power management budget settings in a multidomain system.
- sunHwCtrlDomainPowerMgmtConsumptionThresholdsTable provides a listing of domain power management consumption entries.
- sunHwCtrlTpmTable provides a listing of properties for controlling the access modes of the TPM device for each domain.

Refer to the SUN-HW-CTRL-MIB file for detailed information about these objects.

LDAP/SSL

LDAP/SSL offers enhanced security to LDAP users by way of SSL technology. To configure LDAP/SSL in a SP, you enter basic data (such as primary server, port number, and certificate mode) and optional data (such as alternate server or event or severity levels). You can enter this data using the LDAP/SSL configuration page of the Oracle ILOM web interface, the CLI, or SNMP.

Related Information

- *Oracle SuperCluster M6-32 Owner's Guide: Administration*
- *Oracle ILOM Administrator's Guide for Configuration and Maintenance*

Active Directory

Oracle ILOM supports Active Directory, the distributed directory service included with Microsoft Windows Server operating systems. Like an LDAP directory service implementation, Active Directory is used to authenticate user credentials.

Related Information

- *Oracle SuperCluster M6-32 Owner's Guide: Administration*
- *Oracle ILOM Protocol Management Reference for SNMP, IPMI, CIM, WS-MAN*

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*.

B

base configuration PDomain A SuperCluster M6-32 configuration comprised of two or four PDomains, where one DCU is associated with each PDomain. A base configuration PDomain can reside on a single compute server or can be split across two compute servers. See also *compute server*, *DCU*, *extended configuration PDomain*, and *PDomain*.

C

CFM Cubic feet per minute.

Cisco Catalyst Ethernet switch Provides the SuperCluster M6-32 management network. Referred to in this documentation using the shortened name "Ethernet management switch." See also *Ethernet management switch*.

CMP Chip multiprocessing. Each CMU contains 2 CMP processors. The compute server can contain a maximum of 32 CMPs.

CMU CPU memory unit. Each CMU in the compute server contains two CMPs and two sets of DIMM slots.

COD Capacity on Demand.

compute server Shortened name for the SPARC M6-32 server, a major component of SuperCluster M6-32. See also *SPARC M6-32 server*.

D

Database Domain The domain that contains the SuperCluster M6-32 database.

DB Oracle Database.

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

DCU Domain configurable unit. The smallest building block for PDomains. Each DCU in the compute server contains two or four CMUs and one IOU. 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 10 or 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.

EMS Express module SAS. Each EMS contains two 10GBASE-T network connections and provides access to four hard drives on the compute server.

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*.

expansion rack Shortened name for optional Oracle Exadata Storage Expansion Racks (up to 17) that can be added to SuperCluster M6-32. See also *Oracle Exadata Storage Expansion Rack*.

extended configuration PDomain A SuperCluster M6-32 configuration comprised of two PDomains, where two DCUs are associated with each PDomain. An extended configuration PDomain can reside on a single compute server or can be split across two compute servers. See also, *base configuration PDomain*, *compute server*, *DCU*, and *PDomain*.

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.

fully-populated DCU configuration A configuration where each DCU in the compute servers contains four CMUs. See also *DCU* and *half-populated DCU configuration*.

G

GB Gigabyte. 1 gigabyte = 1024 megabytes.

GbE Gigabit Ethernet.

GNS Grid Naming Service.

H

half-populated DCU configuration A configuration where each DCU in the compute servers contains two CMUs. See also *DCU* and *fully-populated DCU configuration*.

HCA Host channel adapter.

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

I

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*.

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 I/O Domain Creation tool 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*.

IOU I/O unit. The compute server contains up to 4 IOUs, one for each DCU. Each IOU supports up to 16 PCIe slots, 8 10GBASE-T ports on 4 EMS modules, and 8 drives.

IPMI Intelligent Platform Management Interface.

IPMP IP network multipathing.

iSCSI Internet Small Computer System Interface.

K

KVMS Keyboard video mouse storage.

L

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

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*.

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.

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

OCM Oracle Configuration Manager.

ONS Oracle Notification Service.

Oracle Exadata Storage Expansion Rack — Optional expansion racks (in full, half, or quarter configurations) that can be added to SuperCluster M6-32 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 M6-32 — Full name of this SuperCluster model. Referred to in this documentation using the shortened name "SuperCluster M6-32." See also *SuperCluster M6-32*.

Oracle SuperCluster M6-32 storage rack — Full name of first storage rack that contains the storage servers, ZFS storage appliance, IB switches, and Ethernet management switch. Referred to in this documentation using the shortened name "storage rack." See also *storage rack*.

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 — Located in the storage rack, it provides SuperCluster M6-32 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 I/O Domain Creation tool.

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*, *DCU*, 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 32 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 (see also *DCU*) into one or more logical groups (see also *PDomain*).

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

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.

SFP and 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 M6-32, SPs are located on the compute servers, storage servers, ZFS storage appliance controllers, and IB switches. See also *Oracle ILOM*.

SPARC M6-32 server A major component of SuperCluster M6-32 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 M6-32 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 CMUs, memory controllers, and DIMMs within the DCU. 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 rack Shortened name for the Oracle SuperCluster M6-32 storage rack that contains the storage servers. See also *Oracle SuperCluster M6-32 storage rack*.

storage server Storage servers in SuperCluster M6-32.

Sun Datacenter InfiniBand Switch 36 Interconnects SuperCluster M6-32 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 M6-32 Shortened name for Oracle SuperCluster M6-32. See also *Oracle SuperCluster M6-32*.

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 32 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 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

A

Active Directory, overview, 144
Application Domain
  clustering software, 135
  IB network, 86

B

B2-1 LDom configuration, 126
B2-2 LDom configuration, 127
B2-3 LDom configuration, 128
B2-4 LDom configuration, 130
B4-1 LDom configuration, 119
B4-2 LDom configuration, 120
B4-3 LDom configuration, 121
B4-4 LDom configuration, 123
bounded PDomains, description, 62

C

client access network
  overview, 84
clustering software, 133
CMUs
  determining correct number of, 4
  overview, 5
compute servers
  allocating CPU resources, 16
  allocating memory resources, 19
  determining correct number of, 3
  front, 36
  overview, 34
  PCIe slots, 24
  rear, 37
  storage, 27
configuration restrictions, 31
configurations
  DCUs, 42
  LDoms, 81
  PDomains, 61
CPU resources
  allocating
    compute server level, 16
    LDom level, 16
  flowchart, 14
  overview, 82

D

Database Domain
  clustering software, 134
  IB network, 85
DCU 0
  overview, 44
  root complexes
    fully-populated, 54
    half-populated, 48
DCU 1
  overview, 45
  root complexes
    fully-populated, 56
    half-populated, 49
DCU 2
  overview, 46
  root complexes
    fully-populated, 57
    half-populated, 51
DCU 3
  overview, 47
  root complexes
    fully-populated, 59
    half-populated, 52
DCUs
  components, 41
  configurations, 42
  determining correct number of, 4

fully-populated, 5
fully-populated configuration, 42
half-populated, 5
half-populated configuration, 42
locating, 40
overview, 39
PCIe device root complexes, 42
root complexes
fully-populated DCUs, 54
half-populated DCUs, 48
decision points
determining best LDom configuration, 7
determining best PDomain configuration, 7
determining correct amount of memory, 6
determining correct number of CMUs, 4
determining correct number of compute servers, 3
determining correct number of DCUs, 4

E

E2-1 LDom configuration, 109
E2-2 LDom configuration, 111
E2-3 LDom configuration, 112
E2-4 LDom configuration, 114
E4-1 LDom configuration, 101
E4-2 LDom configuration, 103
E4-3 LDom configuration, 104
E4-4 LDom configuration, 106

F

failover, 142
flowchart
allocating CPU resources, 14
allocating memory resources, 17
allocating storage for LDoms, 25
high redundancy, 11
overall, 1
PCIe cards and slots for LDoms, 22
fully-populated DCUs
LDom configurations for, 100
root complexes, 54
understanding, 42

H

half-populated DCUs
root complexes, 48
understanding, 42
high redundancy
compute server level, 13
flowchart, 11
PDomain configuration level, 13

I

IB network
Application Domain, 86
Database Domain, 85
overview, 85

L

LDAP/SSL, overview, 144
LDom configurations
B2-1, 126
B2-2, 127
B2-3, 128
B2-4, 130
B4-1, 119
B4-2, 120
B4-3, 121
B4-4, 123
base PDomains
fully-populated DCUs, 118
half-populated DCUs, 125
E2-1, 109
E2-2, 111
E2-3, 112
E2-4, 114
E4-1, 101
E4-2, 103
E4-3, 104
E4-4, 106
extended PDomains
fully-populated DCUs, 100
half-populated DCUs, 108
LDoms
allocating
CPU resources, 16
memory resources, 20
storage, 27
determining best configuration, 7
flowchart for allocating storage, 25
PCIe slots
allocation, 24
overview, 83

M

management network
  overview, 84
memory resources
  allocating
    compute server level, 19
    LDom level, 20
  flowchart, 17
  overview, 82
memory, determining correct amount of, 6

N

nonbounded PDomains, description, 62

O

Oracle Hardware Management Pack overview, 141
Oracle ILOM
  overview, 137
  platform-specific features, 138
Oracle VTS, functionality, 140
overview
  base configuration PDomains, 116
  client access network, 84
  CMUs, 5
  compute servers, 34
  CPU resources, 82
  DCU 0, 44
  DCU 1, 45
  DCU 2, 46
  DCU 3, 47
  DCUs, 39
  extended configuration PDomains, 99
  IB network, 85
  management network, 84
  memory resources, 82
  PCIe slots, 83
  PDomains, 61
    base configurations, 70
    extended configurations, 63
  SPs, 61
  storage rack, 38
  SuperCluster M6-32, 29

P

PCIe root complexes, 42
PCIe slots, overview, 83
PDomains
  base configurations
    overview, 70, 116
    R3, 72
    R3_1, 73
    R3_2, 73
    R3_3, 74
    R4, 74
    R4_1, 75
    R4_2, 76
    R4_3, 76
    R4_4, 77
    R5, 77
    R5_1, 78
    R6, 78
    R6_1, 79
  bounded PDomains, description, 62
  configuration guidelines, 62
  determining best configuration, 7
  extended configurations
    overview, 63, 99
    R1, 64
    R1_1, 65
    R1_2, 66
    R1_3, 66
    R2, 67
    R2_1, 68
    R2_2, 69
    R2_3, 70
  nonbounded PDomains, description, 62
  overview, 61
physical domains
  See PDomain

R

R1 extended configuration PDomains, 64
R1_1 extended configuration PDomain, 65
R1_2 extended configuration PDomain, 66
R1_3 extended configuration PDomain, 66
R2 extended configuration PDomain, 67
R2_1 extended configuration PDomain, 68
R2_2 extended configuration PDomain, 69
R2_3 extended configuration PDomain, 70
R3 base configuration PDomain, 72
R3_1 base configuration PDomain, 73
R3_2 base configuration PDomain, 73
R3_3 base configuration PDomain, 74
R4 base configuration PDomain, 74

R4_1 base configuration PDomain, 75
R4_2 base configuration PDomain, 76
R4_3 base configuration PDomain, 76
R4_4 base configuration PDomain, 77
R5 base configuration PDomain, 77
R5_1 base configuration PDomain, 78
R6 base configuration PDomain, 78
R6_1 base configuration PDomain, 79
root complexes
  DCU 0
    fully-populated, 54
    half-populated, 48
  DCU 1
    fully-populated, 56
    half-populated, 49
  DCU 2
    fully-populated, 57
    half-populated, 51
  DCU 3
    fully-populated, 59
    half-populated, 52

S

SNMP overview, 142
SP, clock properties, 141
spares kit components, 31
SPs
  overview, 61
storage rack
  overview, 38
SuperCluster M6-32
  configuration flowchart, 1
  configuration restrictions, 31
  hardware components, 33
  overview, 29

synchronizing clock times, 141