

**Oracle<sup>®</sup> SuperCluster M6-32 Zones With Oracle  
Database on Database Domains**

Configuration Guide



Part No.: E38018-01  
Nov 2014

Copyright © 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

---

Copyright © 2013, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

# Contents

---

**Using This Documentation** ix

**Planning to Set Up Zones on Database Domains** 1

Important Cautions 2

Naming Conventions 2

Understanding PDomains 2

PDomains Overview 3

PDomain Guidelines 4

Extended Configuration PDomain Overview 5

Understanding Extended Configuration PDomains 6

Understanding Four DCUs in One Compute Server (R1 Extended Configuration PDomains) 6

Four Fully-Populated DCUs (R1\_1 PDomain Configuration) 7

Two Fully-Populated DCUs and Two Half-Populated DCUs (R1\_2 PDomain Configuration) 7

Four Half-Populated DCUs (R1\_3 PDomain Configuration) 8

Understanding Four DCUs Across Two Compute Servers (R2 Extended Configuration PDomains) 8

Four Fully-Populated DCUs (R2\_1 PDomain Configuration) 9

Two Fully-Populated DCUs and Two Half-Populated DCUs (R2\_2 PDomain Configuration) 9

Four Half-Populated DCUs (R2\_3 PDomain Configuration) 10

Base Configuration Overview 11

Understanding Base Configuration PDomains 12

Understanding Four DCUs on One Compute Server (R3 Base Configuration PDomains)	13
Four Fully-Populated DCUs (R3_1 PDomain Configuration)	14
Two Fully-Populated DCUs and Two Half-Populated DCUs (R3_2 PDomain Configuration)	14
Four Half-Populated DCUs (R3_3 PDomain Configuration)	15
Understanding Four DCUs Across Two Compute Servers (R4 Base Configuration PDomains)	15
Four Fully-Populated DCUs (R4_1 PDomain Configuration)	16
Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_2 PDomain Configuration)	16
Four Half-Populated DCUs (R4_3 PDomain Configuration)	17
Two Fully-Populated DCUs and Two Half-Populated DCUs (R4_4 PDomain Configuration)	17
Understanding Two DCUs on One Compute Server (R5 Base Configuration PDomains)	18
Two Fully-Populated DCUs (R5_1 PDomain Configuration)	19
Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains)	19
Two Fully-Populated DCUs (R6_1 PDomain Configuration)	20
Understanding LDom	20
Understanding LDom Configurations for Extended Configuration PDomains	21
Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)	21
LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)	22
E4-1 LDom Configuration	23
E4-2 LDom Configuration	24
E4-3 LDom Configuration	25
E4-4 LDom Configuration	27

Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)	29
LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)	30
E2-1 LDom Configuration	31
E2-2 LDom Configuration	33
E2-3 LDom Configuration	34
E2-4 LDom Configuration	36
Understanding LDom Configurations for Base Configuration PDomains	38
Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)	39
LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)	40
B4-1 LDom Configuration	41
B4-2 LDom Configuration	42
B4-3 LDom Configuration	43
B4-4 LDom Configuration	44
Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)	46
LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)	47
B2-1 LDom Configuration	48
B2-2 LDom Configuration	49
B2-3 LDom Configuration	50
B2-4 LDom Configuration	51
Determining the Cores Available for LDoms and Zones	53
Cores Available for Domains and Zones	53
Example Configuration	56
LDom 1 (First Database Domain)	57
LDom 4 (Second Database Domain)	57

Memory Available for LDoms	58
Zones and Cluster Planning for Database Domains	59
Guidelines for Planning the Number of Zones and Clusters	60
Guidelines for Planning the Storage Server Disk Group Layout	61
Planning the Storage Server Disk Group Layout	62
Determine the Available Existing Storage Capacity	63
Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster	65
Determine the Amount of Space for Each Disk Group in Each Cluster	66
Understanding an Example Scenario	66
Determine the Available Existing Storage Capacity (Example Scenario)	67
Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster (Example Scenario)	67
Determine the Amount of Space for Each Disk Group in Each Cluster (Example Scenario)	68
Recording Your Existing and Planned Configuration Information	68
Recording Your Zone Configuration Information	68
Database Domain 1 Zone Configuration Information	69
Database Domain 2 Zone Configuration Information	72
Database Domain 3 Zone Configuration Information	74
Database Domain 4 Zone Configuration Information	77
Recording Your Cluster Configuration Information	79
Recording Your Storage Server Disk Group Layout	80
<b>Preparing to Configure Zones on Database Domains</b>	<b>83</b>
▼ Prepare to Configure Zones	83
▼ Update the Base Software	85
▼ Determine the Repository Location	85

- ▼ Install or Update Packages From the Remote Repository 86
- ▼ Install or Update Files From the Local Repository 88
- ▼ Verify Configuration Tool Installation 91

## **Creating Configuration Files (OEDA) 95**

- Example Configuration 96
- ▼ Verify Storage Server Disk Space for Additional Zones 97
- ▼ Locate the Necessary Files to Set Up Zones 98
- ▼ Set Up the OEDA 100
- ▼ Start the OEDA 101
- ▼ Import the Most Recent OEDA Configuration File 102
- ▼ Review Existing Configuration Information 104
- ▼ Review the Information in the Identify Compute Node Operating System Page 105
- ▼ Review the Information in the Management and Private Networks Page 107
- ▼ Complete the Define Clusters Page 109
- ▼ Set Zone Default Configurations 114
- ▼ Complete the Cluster Page 117
- ▼ Complete the Cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks Page 121
- ▼ Verify Remaining Configuration Information 123
- ▼ Generate the Configuration Files 123

## **Creating the Template Zone on Each Database Domain 125**

- ▼ Create a Template Zone on a Database Domain 125
- ▼ Delete a Template Zone From a Database Domain 128

## **Creating VNETs 131**

- Determining if Additional VNETs Are Needed for a Database Domain 131

- ▼ Determine if Additional VNETs Are Needed (Using Specific LDom Information) 132
- ▼ Determine if Additional VNETs Are Needed (Using Software Commands) 132
- ▼ Locate the Control Domain 133
- ▼ Create VNETs for Zones 134
- ▼ Delete VNETs 136

### **Running the Java OneCommand 139**

- ▼ Set Up Public Key Authentication for ZFS Storage Controllers 139
- ▼ Set Up the Java OneCommand 141
- ▼ Run the Java OneCommand 142
- NTP Issue 144

### **Creating Additional Links on the IB Storage Network for Zones 147**

- ▼ Create Additional Links on the IB Storage Network for Zones 147

### **Setting Up VLAN Tagging for Zones 151**

- ▼ Set Up VLAN Tagging for Zones 151

### **Index 155**



# Using This Documentation

---

This document describes how to set up zones and Real Application Clusters (RACs) on the Database Domains in the Oracle SuperCluster M6-32.

- “Product Notes” on page ix
- “Related Documentation” on page x
- “Feedback” on page x
- “Access to Oracle Support” on page x

---

## Product Notes

For late-breaking information and known issues about this product, refer to the product notes, which can be viewed through a browser by viewing the following directory on the Oracle SuperCluster M6-32:

```
/opt/oracle/node/doc/E41531_01/index.html
```

---

## Related Documentation

Documentation	Links
All Oracle products	<a href="http://www.oracle.com/documentation">http://www.oracle.com/documentation</a>
Oracle Solaris OS and systems software library	<a href="http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw">http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw</a> <a href="http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw">http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw</a>

---

## Feedback

Provide feedback about this documentation at:

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

---

## Access to Oracle Support

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

# Planning to Set Up Zones on Database Domains

---

Once the Oracle SuperCluster M6-32 is set up, you can modify it to deploy zones in the Application Domains and Database Domains. These topics describe how to plan for configuring zones on Database Domains.

For instructions on setting up zones in the Application Domains, refer to the following documentation:

- For Oracle Solaris 10, refer to *Oracle Solaris Administration: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones* at:  
[http://download.oracle.com/docs/cd/E23823\\_01](http://download.oracle.com/docs/cd/E23823_01)
- For Oracle Solaris 11, refer to *Oracle Solaris Administration: Network Interfaces and Network Virtualization* and *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* at:  
[http://download.oracle.com/docs/cd/E23824\\_01](http://download.oracle.com/docs/cd/E23824_01)

The following sections provide information for planning on setting up zones on Database Domains:

- “Important Cautions” on page 2
- “Naming Conventions” on page 2
- “Understanding PDomains” on page 2
- “Understanding LDOMs” on page 20
- “Determining the Cores Available for LDOMs and Zones” on page 53
- “Memory Available for LDOMs” on page 58
- “Zones and Cluster Planning for Database Domains” on page 59
- “Guidelines for Planning the Number of Zones and Clusters” on page 60
- “Guidelines for Planning the Storage Server Disk Group Layout” on page 61
- “Recording Your Existing and Planned Configuration Information” on page 68

---

## Important Cautions



---

**Caution** – Ensure that you back up all existing databases before running the Java OneCommand.

---



---

**Caution** – You must use the latest version of OEDA and the Java OneCommand patch (patch 19766190 or later). Refer to the Oracle Exadata Deployment Assistant (OEDA) section in MOS Note 888828.1 for details.

---



---

**Caution** – Any version of Java OneCommand prior to patch 19766190 can destroy the storage cell disks and griddisk if an undo option is performed on certain steps (Create Cell disk, for example). This can cause complete destruction of all the griddisks on the storage cells. In addition, re-running the griddisk creation step or mistakenly specifying a non-unique diskgroup in OEDA will result in the destruction of existing griddisks. Be aware that older versions of Java OneCommand also destroy cell disks and griddisks with the Create Cell Disks step.

---

---

## Naming Conventions

This document uses these naming conventions:

- *oscnnumber* for the Oracle SuperCluster M6-32
- *rnumber* for the cluster number
- *dbnumber* for the Database Domain number
- *znumber* for the zone number

---

## Understanding PDomains

These topics describe PDomains and the PDomain configurations.

- [“PDomains Overview” on page 3](#)

- “PDomain Guidelines” on page 4
- “Extended Configuration PDomain Overview” on page 5
- “Understanding Extended Configuration PDomains” on page 6
- “Base Configuration Overview” on page 11
- “Understanding Base Configuration PDomains” on page 12

## 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	<p>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.</p> <p>A nonbounded PDomain can contain one to four DCUs.</p>
Bounded PDomain	<p>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.</p>

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 4](#)
- [“Extended Configuration PDomain Overview” on page 5](#)
- [“Understanding Extended Configuration PDomains” on page 6](#)
- [“Base Configuration Overview” on page 11](#)
- [“Understanding Base Configuration PDomains” on page 12](#)

## 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

- [“PDomain Guidelines” on page 4](#)

- “Extended Configuration PDomain Overview” on page 5
- “Understanding Extended Configuration PDomains” on page 6
- “Base Configuration Overview” on page 11
- “Understanding Base Configuration PDomains” on page 12

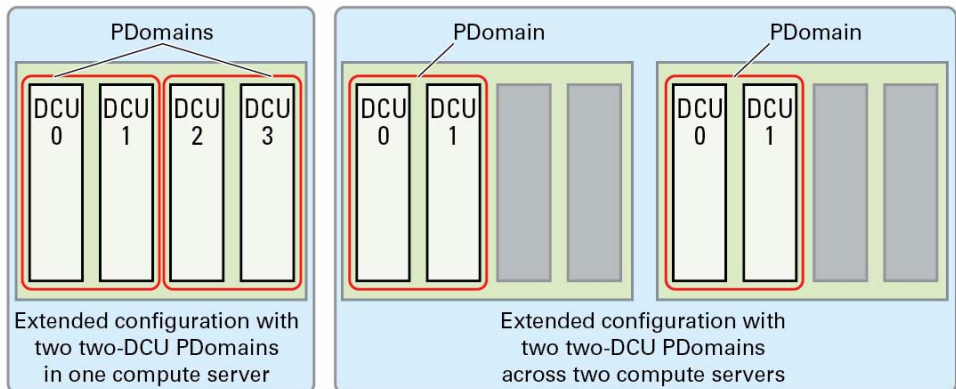
## Extended Configuration PDomain Overview

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 3 for more information on the characteristics of this type of PDomain.

The following figure illustrates the available extended configurations.



### Related Information

- “PDomains Overview” on page 3
- “PDomain Guidelines” on page 4

- [“Understanding Extended Configuration PDomains” on page 6](#)

## Understanding Extended Configuration PDomains


The following extended configuration PDomains are available:

- [“Understanding Four DCUs in One Compute Server \(R1 Extended Configuration PDomains\)” on page 6](#)
- [“Understanding Four DCUs Across Two Compute Servers \(R2 Extended Configuration PDomains\)” on page 8](#)

### 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 7](#)
- [“Two Fully-Populated DCUs and Two Half-Populated DCUs \(R1\\_2 PDomain Configuration\)” on page 7](#)
- [“Four Half-Populated DCUs \(R1\\_3 PDomain Configuration\)” on page 8](#)




### *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 21


### *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 21
- DCU 2 and DCU 3: [“Understanding LDom Configurations for Half-Populated DCUs \(Extended Configuration PDomains\)”](#) on page 29

## Four Half-Populated DCUs (R1\_3 PDomain Configuration)



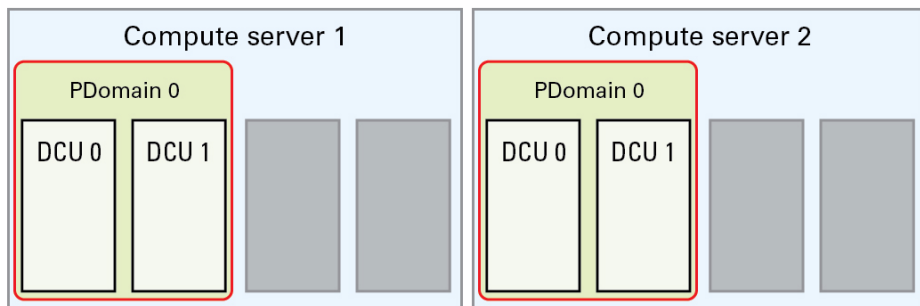
For this configuration, the following LDom configurations are available for each DCU:

- [“Understanding LDom Configurations for Half-Populated DCUs \(Extended Configuration PDomains\)”](#) on page 29

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

This configuration has the following characteristics:

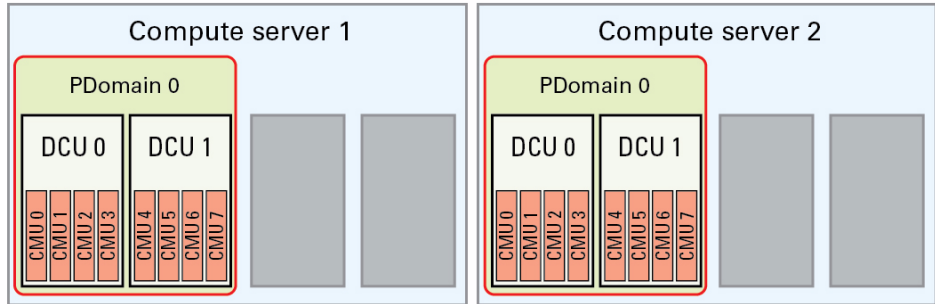
- Two PDomains, with two DCUs associated with each PDomain
- The two PDomains split across two compute servers (one PDomain on each compute server)



Three DCU configuration are options available for this configuration:

- “Four Fully-Populated DCUs (R2\_1 PDomain Configuration)” on page 9
- “Two Fully-Populated DCUs and Two Half-Populated DCUs (R2\_2 PDomain Configuration)” on page 9
- “Four Half-Populated DCUs (R2\_3 PDomain Configuration)” on page 10

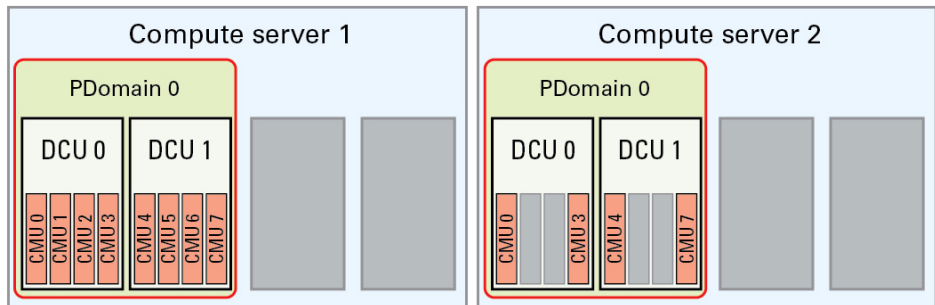
### *Four Fully-Populated DCUs (R2\_1 PDomain Configuration)*



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 21
- Extended PDomain on second compute server: “Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)” on page 21

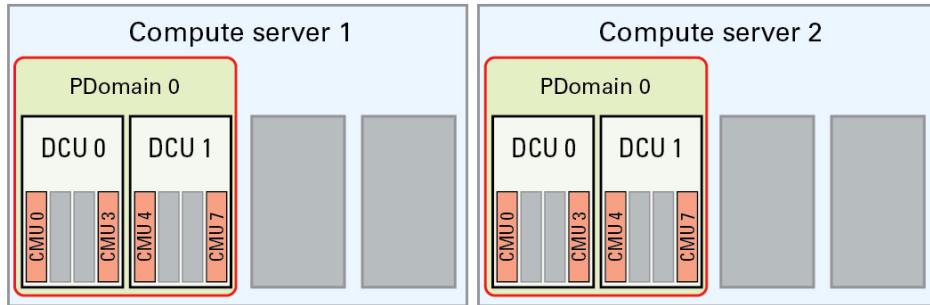
### *Two Fully-Populated DCUs and Two Half-Populated DCUs (R2\_2 PDomain Configuration)*



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 21
- Extended PDomain on second compute server: “Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)” on page 29

### *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 29
- Extended PDomain on second compute server: [“Understanding LDom Configurations for Half-Populated DCUs \(Extended Configuration PDomains\)”](#) on page 29

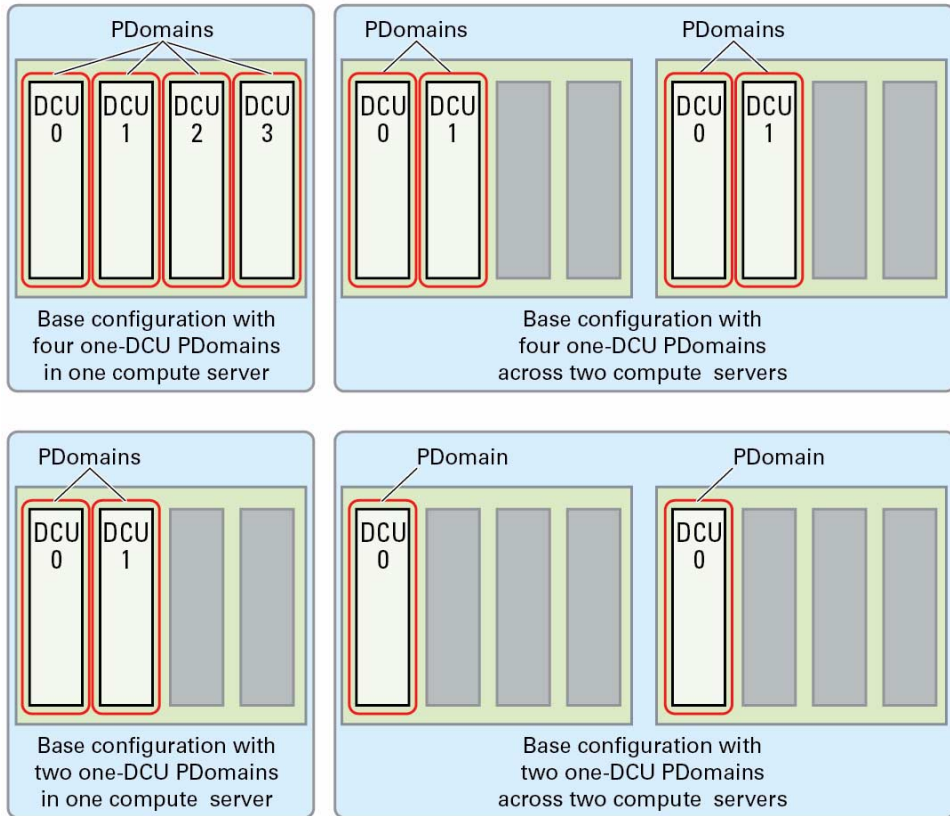
## Base Configuration Overview

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 3 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 3
- [“PDomain Guidelines”](#) on page 4
- [“Understanding Base Configuration PDomains”](#) on page 12

## Understanding Base Configuration PDomains

The following base configuration PDomains are available:


- [“Understanding Four DCUs on One Compute Server \(R3 Base Configuration PDomains\)”](#) on page 13
- [“Understanding Four DCUs Across Two Compute Servers \(R4 Base Configuration PDomains\)”](#) on page 15

- “Understanding Two DCUs on One Compute Server (R5 Base Configuration PDomains)” on page 18
- “Understanding Two DCUs Across Two Compute Servers (R6 Base Configuration PDomains)” on page 19

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

This configuration has the following characteristics:

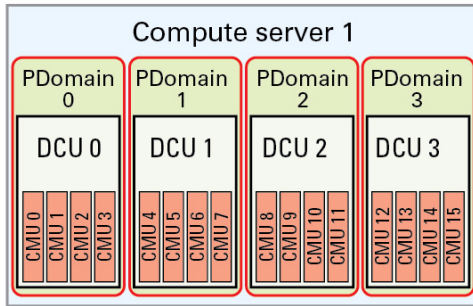
- Four PDomains, with one DCU associated with each PDomain
- All four PDomains on a single compute server



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

- “Four Fully-Populated DCUs (R3\_1 PDomain Configuration)” on page 14
- “Two Fully-Populated DCUs and Two Half-Populated DCUs (R3\_2 PDomain Configuration)” on page 14
- “Four Half-Populated DCUs (R3\_3 PDomain Configuration)” on page 15

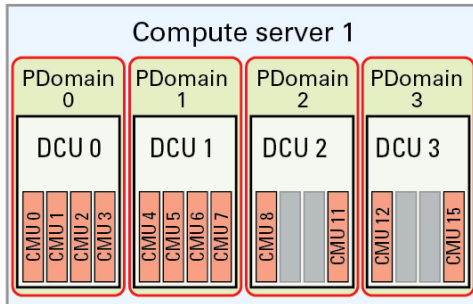
### *Four Fully-Populated DCUs (R3\_1 PDomain Configuration)*



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

- [“Understanding LDom Configurations for Fully-Populated DCUs \(Base Configuration PDomains\)”](#) on page 39

### *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 39
- DCU 2 and DCU 3: [“Understanding LDom Configurations for Half-Populated DCUs \(Base Configuration PDomains\)”](#) on page 46



### Four Half-Populated DCUs (R3\_3 PDomain Configuration)



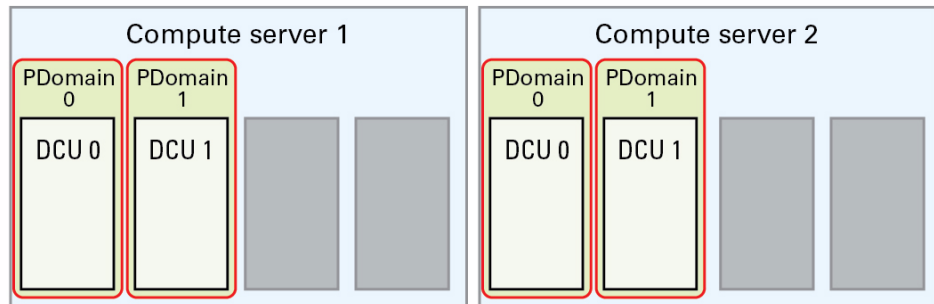
For this configuration, the following LDom configurations are available for each DCU:

- [“Understanding LDom Configurations for Half-Populated DCUs \(Base Configuration PDomains\)”](#) on page 46

### 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)

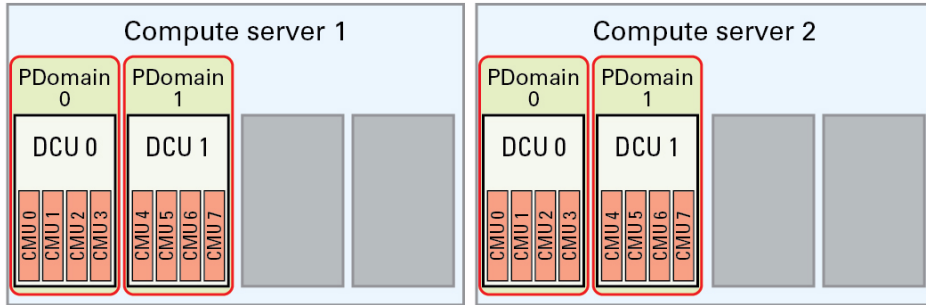


There are four DCU configuration options available for this configuration:

- [“Four Fully-Populated DCUs \(R4\\_1 PDomain Configuration\)”](#) on page 16
- [“Two Fully-Populated DCUs and Two Half-Populated DCUs \(R4\\_2 PDomain Configuration\)”](#) on page 16
- [“Four Half-Populated DCUs \(R4\\_3 PDomain Configuration\)”](#) on page 17

- “Two Fully-Populated DCUs and Two Half-Populated DCUs (R4\_4 PDomain Configuration)” on page 17

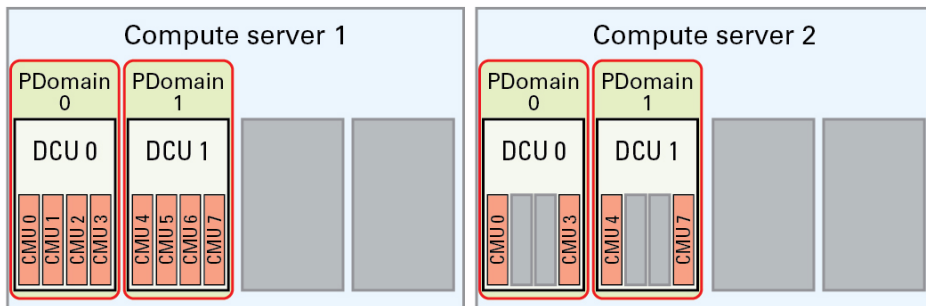
*Four Fully-Populated DCUs (R4\_1 PDomain Configuration)*



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 39
- Two base PDomains on second compute server: “[Understanding LDom Configurations for Fully-Populated DCUs \(Base Configuration PDomains\)](#)” on page 39

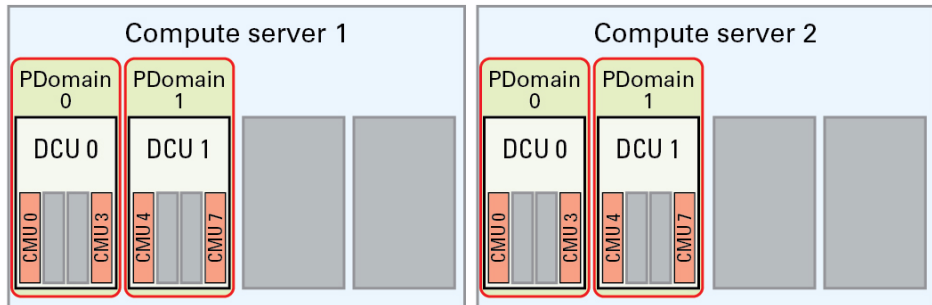
*Two Fully-Populated DCUs and Two Half-Populated DCUs (R4\_2 PDomain Configuration)*



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 39
- Two base PDomains on second compute server: “Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)” on page 46

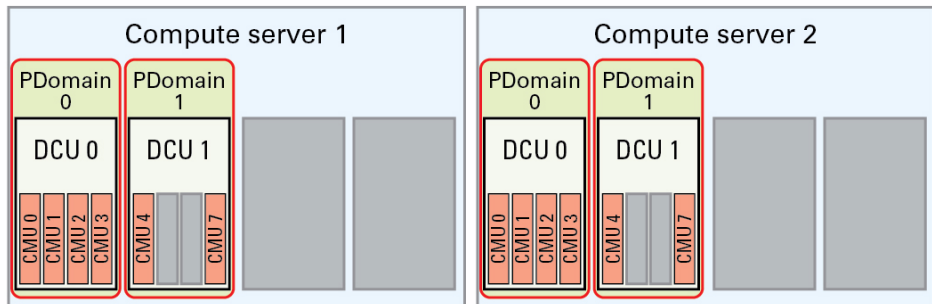
### *Four Half-Populated DCUs (R4\_3 PDomain Configuration)*



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 46
- Two base PDomains on second compute server: “Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)” on page 46

### *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 39
  - DCU 1: “Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)” on page 46
- Second compute server:
  - DCU 0: “Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)” on page 39
  - DCU 1: “Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)” on page 46

## 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 19

## 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 39

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

This configuration has the following characteristics:

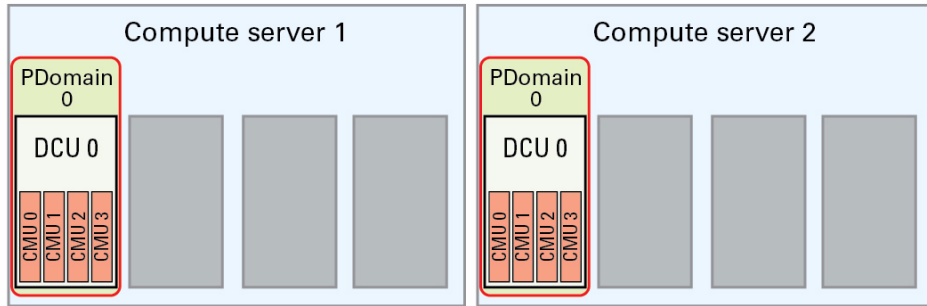
- Two PDomains, with one DCU associated with each PDomain
- The two PDomains split across two compute servers (one PDomain on each compute server)



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

- [“Two Fully-Populated DCUs \(R6\\_1 PDomain Configuration\)”](#) on page 20

## 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 39

---

## Understanding LDom

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 LDom Configurations for Extended Configuration PDomains”](#) on page 21
- [“Understanding LDom Configurations for Base Configuration PDomains”](#) on page 38

# 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

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.	<a href="#">“Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)” on page 21</a>
Choose the LDom configuration for extended configuration PDomains on half-populated DCUs.	<a href="#">“Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)” on page 29</a>

## 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 LDom 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 22](#)

- “E4-1 LDom Configuration” on page 23
- “E4-2 LDom Configuration” on page 24
- “E4-3 LDom Configuration” on page 25
- “E4-4 LDom Configuration” on page 27

### Related Information

- “Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)” on page 39

### *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)

		DCU x								DCU y							
		CMU a		CMU b		CMU c		CMU d		CMU a		CMU b		CMU c		CMU d	
		CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1
Configurations	E4-1	LDom 1															
	E4-2	LDom 1												LDom 2			
	E4-3	LDom 1												LDom 2		LDom 3	
	E4-4	LDom 1								LDom 2		LDom 3		LDom 4			

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 23
- “E4-2 LDom Configuration” on page 24
- “E4-3 LDom Configuration” on page 25
- “E4-4 LDom Configuration” on page 27

### *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	<ul style="list-style-type: none"><li>• HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li><li>• HDD2 through HDD7 in first DCU and HDD0 through HDD7 in second DCU available for additional storage</li></ul>
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

**TABLE:** Networks (E4-1 Configuration)

LDom 1		
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 22](#)

### 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	<ul style="list-style-type: none"> <li>• EMS 1 through EMS 4 in first DCU</li> <li>• EMS 1 and 2 in second DCU</li> </ul>	EMS 3 and EMS 4 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 in first DCU and HDD0 through HDD3 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• HDD6 and HDD7 in second DCU available for additional storage</li> </ul>
IB HCAs	<ul style="list-style-type: none"> <li>• PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>• PCIe slots 2 and 6 in second DCU</li> </ul>	PCIe slots 10 and 14 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>• PCIe slots 1, 4, 5, 7, and 8 in second DCU</li> </ul>	PCI slots 9, 11, 12, 13, 15 and 16 in second DCU
Default CPU/Memory Resources	<ul style="list-style-type: none"> <li>• 100% of resources in first DCU</li> <li>• 50% of resources in second DCU</li> </ul>	50% of resources in second DCU

**TABLE:** Networks (E4-2 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>
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
IB Network: Exadata Private Network (DB Domains)	Active	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	P0 in IB HCAs in slots 10 and 14 in second DCU
	Standby	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	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 22](#)

*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	<ul style="list-style-type: none"> <li>EMS 1 through EMS 4 in first DCU</li> <li>EMS 1 and 2 in second DCU</li> </ul>	EMS 3 in second DCU	EMS 4 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>HDD2 through HDD7 in first DCU and HDD0 through HDD3 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD6 and HDD7 in second DCU as boot and /u01 primary and mirrored drives</li> <li>No drives available for additional storage</li> </ul>
IB HCAs	<ul style="list-style-type: none"> <li>PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>PCIe slots 2 and 6 in second DCU</li> </ul>	PCIe slot 10 in second DCU	PCIe slot 14 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>PCIe slots 1, 4, 5, 7, and 8 in second DCU</li> </ul>	11, 13, and 16 in second DCU	PCIe slots 9, 12, and 15 in second DCU
Default CPU/Memory Resources	<ul style="list-style-type: none"> <li>100% of resources in first DCU</li> <li>50% of resources in second DCU</li> </ul>	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

**TABLE:** Networks (E4-3 Configuration)

		LDom 1	LDom 2	LDom 3
IB Network: Exadata Private Network (DB Domains)	Active	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	P0 in IB HCA in slot 10 in second DCU	P0 in IB HCA in slot 14 in second DCU
	Standby	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	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 22](#)

*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)

	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
EMSs	<ul style="list-style-type: none"> <li>• EMS 1 through EMS 4 in first DCU</li> <li>• EMS 1 in second DCU</li> </ul>	EMS 2 in second DCU	EMS 3 in second DCU	EMS 4 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 in first DCU and HDD0 and HDD1 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD2 and HDD3 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• No drives available for additional storage</li> </ul>
IB HCAs	<ul style="list-style-type: none"> <li>• PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>• PCIe slot 2 in second DCU</li> </ul>	PCIe slot 6 in second DCU	PCIe slot 10 in second DCU	PCIe slot 14 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>• PCIe slots 5 and 8 in second DCU</li> </ul>	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	<ul style="list-style-type: none"> <li>• 100% of resources in first DCU</li> <li>• 25% of resources in second DCU</li> </ul>	25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

**TABLE:** Networks (E4-4 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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 (E4-4 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 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	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCA in slot 2 in second DCU</li> </ul>	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	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCA in slots 2 in second DCU</li> </ul>	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 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

### Related Information

- [“LDom Configurations for Fully-Populated DCUs \(Extended Configuration PDomains\)” on page 22](#)

## 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 LDom 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 30
- [“E2-1 LDom Configuration”](#) on page 31
- [“E2-2 LDom Configuration”](#) on page 33
- [“E2-3 LDom Configuration”](#) on page 34
- [“E2-4 LDom Configuration”](#) on page 36

### **Related Information**

- [“Understanding LDom Configurations for Half-Populated DCUs \(Base Configuration PDomains\)”](#) on page 46

### *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)

		DCU x				DCU y			
		CMU a		CMU b		CMU a		CMU b	
		CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1
Configurations	E2-1	LDom 1							
	E2-2	LDom 1				LDom 2			
	E2-3	LDom 1				LDom 2	LDom 3		
	E2-4	LDom 1				LDom 2	LDom 3	LDom 4	

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 31
- “E2-2 LDom Configuration” on page 33
- “E2-3 LDom Configuration” on page 34
- “E2-4 LDom Configuration” on page 36

### *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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 in first DCU and HDD0 through HDD7 in second DCU available for additional storage</li> </ul>
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 30](#)

## 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	<ul style="list-style-type: none"> <li>EMS 1 through EMS 4 in first DCU</li> <li>EMS 1 and 2 in second DCU</li> </ul>	EMS 3 and EMS 4 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>HDD2 through HDD7 in first DCU and HDD0 through HDD3 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>HDD6 and HDD7 in second DCU available for additional storage</li> </ul>
IB HCAs	<ul style="list-style-type: none"> <li>PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>PCIe slots 2 and 6 in second DCU</li> </ul>	PCIe slots 10 and 14 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>PCIe slots 1, 4, 5, 7, and 8 in second DCU</li> </ul>	PCI slots 9, 11, 12, 13, 15 and 16 in second DCU
Default CPU/Memory Resources	<ul style="list-style-type: none"> <li>100% of resources in first DCU</li> <li>50% of resources in second DCU</li> </ul>	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)

		<b>LDom 1</b>	<b>LDom 2</b>
IB Network: Exadata Private Network (DB Domains)	Active	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	P0 in IB HCAs in slots 10 and 14 in second DCU
	Standby	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	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 30](#)

### *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)

<b>Item</b>	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>
EMSs	<ul style="list-style-type: none"> <li>• EMS 1 through EMS 4 in first DCU</li> <li>• EMS 1 and 2 in second DCU</li> </ul>	EMS 4 in second DCU	EMS 3 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 in first DCU and HDD0 through HDD3 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 in second DCU as boot and /u01 primary and mirrored drives</li> <li>• No drives available for additional storage</li> </ul>

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

Item	LDom 1	LDom 2	LDom 3
IB HCAs	<ul style="list-style-type: none"> <li>• PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>• PCIe slots 2 and 6 in second DCU</li> </ul>	PCIe slot 14 in second DCU	PCIe slot 10 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>• PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>• PCIe slots 1, 4, 5, 7, and 8 in second DCU</li> </ul>	13, 15 and 16 in second DCU	9, 11 and 12 in second DCU
Default CPU/Memory Resources	<ul style="list-style-type: none"> <li>• 100% of resources in first DCU</li> <li>• 50% of resources in second DCU</li> </ul>	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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>
IB Network: Exadata Private Network (DB Domains)	Active	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	P0 in IB HCA in slot 14 in second DCU	P0 in IB HCA in slot 10 in second DCU
	Standby	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCAs in slots 2 and 6 in second DCU</li> </ul>	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 30](#)

### *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)

	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
EMSs	<ul style="list-style-type: none"> <li>EMS 1 through EMS 4 in first DCU</li> <li>EMS 1 in second DCU</li> </ul>	EMS 2 in second DCU	EMS 4 in second DCU	EMS 3 in second DCU
Hard Drives	<ul style="list-style-type: none"> <li>HDD0 and HDD1 in first DCU as boot and /u01 primary and mirrored drives</li> <li>HDD2 through HDD7 in first DCU and HDD0 and HDD1 in second DCU available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD2 and HDD3 in second DCU as boot and /u01 primary and mirrored drives</li> <li>No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD4 and HDD5 in second DCU as boot and /u01 primary and mirrored drives</li> <li>No drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>HDD6 and HDD7 in second DCU as boot and /u01 primary and mirrored drives</li> <li>No drives available for additional storage</li> </ul>
IB HCAs	<ul style="list-style-type: none"> <li>PCIe slots 2, 6, 10, and 14 in first DCU</li> <li>PCIe slot 2 in second DCU</li> </ul>	PCIe slot 6 in second DCU	PCIe slot 14 in second DCU	PCIe slot 10 in second DCU
Empty (free) PCIe slots	<ul style="list-style-type: none"> <li>PCIe slots 1, 4, 5, 7, 8, 9, 11, 12, 13, 15, and 16 in first DCU</li> <li>PCIe slots 1 and 4 in second DCU</li> </ul>	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	<ul style="list-style-type: none"> <li>100% of resources in first DCU</li> <li>25% of resources in second DCU</li> </ul>	25% of resources in second DCU	25% of resources in second DCU	25% of resources in second DCU

**TABLE:** Networks (E2-4 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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	<ul style="list-style-type: none"> <li>• P0 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P0 in IB HCA in slot 2 in second DCU</li> </ul>	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	<ul style="list-style-type: none"> <li>• P1 in IB HCAs in slots 2, 6, 10, and 14 in first DCU</li> <li>• P1 in IB HCA in slots 2 in second DCU</li> </ul>	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 30](#)

## 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.	<a href="#">“Understanding LDom Configurations for Fully-Populated DCUs (Base Configuration PDomains)” on page 39</a>
Choose the LDom configuration for base configuration PDomains on half-populated DCUs.	<a href="#">“Understanding LDom Configurations for Half-Populated DCUs (Base Configuration PDomains)” on page 46</a>

## 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 LDom 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 40](#)
- [“B4-1 LDom Configuration” on page 41](#)
- [“B4-2 LDom Configuration” on page 42](#)
- [“B4-3 LDom Configuration” on page 43](#)
- [“B4-4 LDom Configuration” on page 44](#)

## Related Information

- “Understanding LDom Configurations for Fully-Populated DCUs (Extended Configuration PDomains)” on page 21

### *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)

		DCU <i>x</i>							
		CMU <i>a</i>		CMU <i>b</i>		CMU <i>c</i>		CMU <i>d</i>	
		CMP0	CMP1	CMP0	CMP1	CMP0	CMP1	CMP0	CMP1
Configurations	B4-1	LDom 1							
	B4-2	LDom 1				LDom 2			
	B4-3	LDom 1				LDom 2		LDom 3	
	B4-4	LDom 1		LDom 2		LDom 3		LDom 4	

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

## Related Information

- “B4-1 LDom Configuration” on page 41
- “B4-2 LDom Configuration” on page 42
- “B4-3 LDom Configuration” on page 43

- [“B4-4 LDom Configuration” on page 44](#)

## *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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 available for additional storage</li> </ul>
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
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 40

### *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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 and HDD3 available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• HDD6 and HDD7 available for additional storage</li> </ul>
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

**TABLE:** Networks (B4-2 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>
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 40](#)

### *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)

<b>Item</b>	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>
EMSs	EMS 1 and EMS 2	EMS 3	EMS 4
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 and HDD3 available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>
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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>
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 40](#)

*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)

	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
EMSs	EMS 1	EMS 2	EMS 3	EMS 4
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD2 and HDD3 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>
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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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
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

**TABLE:** Networks (B4-4 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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 40](#)

## 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 LDom 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 47](#)
- [“B2-1 LDom Configuration” on page 48](#)
- [“B2-2 LDom Configuration” on page 49](#)
- [“B2-3 LDom Configuration” on page 50](#)
- [“B2-4 LDom Configuration” on page 51](#)



## Related Information

- “Understanding LDom Configurations for Half-Populated DCUs (Extended Configuration PDomains)” on page 29

### *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 a		CMU b	
		CMP0	CMP1	CMP0	CMP1
Configurations	B2-1	LDom 1			
	B2-2	LDom 1		LDom 2	
	B2-3	LDom 1		LDom 2	LDom 3
	B2-4	LDom 1	LDom 2	LDom 3	LDom 4

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 3	CMU 7	CMU 11	CMU 15

## Related Information

- “B2-1 LDom Configuration” on page 48
- “B2-2 LDom Configuration” on page 49
- “B2-3 LDom Configuration” on page 50
- “B2-4 LDom Configuration” on page 51

## 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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 through HDD7 available for additional storage</li> </ul>
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
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 47](#)

## 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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 and HDD3 available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• HDD6 and HDD7 available for additional storage</li> </ul>
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

**TABLE:** Networks (B2-2 Configuration)

		LDom 1	LDom 2
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 47](#)

### 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	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• HDD2 and HDD3 available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>
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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>
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 47](#)

### *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)

	<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
EMSs	EMS 1	EMS 2	EMS 4	EMS 3
Hard Drives	<ul style="list-style-type: none"> <li>• HDD0 and HDD1 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD2 and HDD3 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD4 and HDD5 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>	<ul style="list-style-type: none"> <li>• HDD6 and HDD7 as boot and /u01 primary and mirrored drives</li> <li>• No hard drives available for additional storage</li> </ul>
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)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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
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

**TABLE:** Networks (B2-4 Configuration)

		<b>LDom 1</b>	<b>LDom 2</b>	<b>LDom 3</b>	<b>LDom 4</b>
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 47](#)

---

## Determining the Cores Available for LDom and Zones

The following sections provide information on determining the number of cores that are available for LDom and zones:

- [“Cores Available for Domains and Zones” on page 53](#)
- [“Example Configuration” on page 56](#)

## Cores Available for Domains and Zones

The number of cores available for LDom, and for zones within those LDom, varies depending on the following factors:

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

Each CMU contains two CMPs, and each CMP contains one twelve-core SPARC M6 processor. The number of CMUs and CMPs that are associated with a specific domain varies, depending on the type of configuration PDomain that you have on each compute server and the type of LDom configuration that you have on the PDomain.

When you first install the operating system instances on an LDom, that LDom is automatically designated as the global zone. When creating zones on Database Domains, the Database Domain is designated as the global zone. The zones created on that Database Domain are designated as nonglobal zones.

A certain number of cores are always set aside for the global zone, and the remaining cores in the Database Domain are available for the nonglobal zones (the zones in the Database Domain). The number of cores that are set aside for the global zone varies, depending on the factors mentioned previously.

The following tables provide information on the number of cores that are available for specific LDom and zones that are created within those LDom. See [“Understanding LDom” on page 20](#) for more information on the different types of PDomain and LDom configurations.

**TABLE:** Cores Available for Extended Configuration PDomains (Fully-Populated DCU Configurations)

<b>LDom Configuration</b>	<b>Total Number of Cores Available for LDom in This Configuration</b>	<b>Number of Cores Set Aside for Global Zone</b>	<b>Number of Cores Available for Nonglobal Zones</b>
<b>E4-1</b>	LDom 1: 192	LDom 1: 4	LDom 1: 188
<b>E4-2</b>	<ul style="list-style-type: none"> <li>• LDom 1: 144</li> <li>• LDom 2: 48</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 140</li> <li>• LDom 2: 44</li> </ul>
<b>E4-3</b>	<ul style="list-style-type: none"> <li>• LDom 1: 144</li> <li>• LDom 2: 24</li> <li>• LDom 3: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> <li>• LDom 3: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 140</li> <li>• LDom 2: 20</li> <li>• LDom 3: 20</li> </ul>
<b>E4-4</b>	<ul style="list-style-type: none"> <li>• LDom 1: 120</li> <li>• LDom 2: 24</li> <li>• LDom 3: 24</li> <li>• LDom 4: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> <li>• LDom 3: 4</li> <li>• LDom 4: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 116</li> <li>• LDom 2: 20</li> <li>• LDom 3: 20</li> <li>• LDom 4: 20</li> </ul>



**TABLE:** Cores Available for Extended Configuration PDomains (Half-Populated DCU Configurations)

<b>LDom Configuration</b>	<b>Total Number of Cores Available for Domains in This Configuration</b>	<b>Number of Cores Set Aside for Global Zone</b>	<b>Number of Cores Available for Nonglobal Zones</b>
<b>E2-1</b>	LDom 1: 96	LDom 1: 4	LDom 1: 92
<b>E2-2</b>	<ul style="list-style-type: none"> <li>• LDom 1: 72</li> <li>• LDom 2: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 68</li> <li>• LDom 2: 20</li> </ul>
<b>E2-3</b>	<ul style="list-style-type: none"> <li>• LDom 1: 72</li> <li>• LDom 2: 12</li> <li>• LDom 3: 12</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 2</li> <li>• LDom 3: 2</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 68</li> <li>• LDom 2: 10</li> <li>• LDom 3: 10</li> </ul>
<b>E2-4</b>	<ul style="list-style-type: none"> <li>• LDom 1: 60</li> <li>• LDom 2: 12</li> <li>• LDom 3: 12</li> <li>• LDom 4: 12</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 2</li> <li>• LDom 3: 2</li> <li>• LDom 4: 2</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 56</li> <li>• LDom 2: 10</li> <li>• LDom 3: 10</li> <li>• LDom 4: 10</li> </ul>

**TABLE:** Cores Available for Base Configuration PDomains (Fully-Populated DCU Configurations)

<b>LDom Configuration</b>	<b>Total Number of Cores Available for Domains in This Configuration</b>	<b>Number of Cores Set Aside for Global Zone</b>	<b>Number of Cores Available for Nonglobal Zones</b>
<b>B4-1</b>	LDom 1: 96	LDom 1: 4	LDom 1: 92
<b>B4-2</b>	<ul style="list-style-type: none"> <li>• LDom 1: 48</li> <li>• LDom 2: 48</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 44</li> <li>• LDom 2: 44</li> </ul>
<b>B4-3</b>	<ul style="list-style-type: none"> <li>• LDom 1: 48</li> <li>• LDom 2: 24</li> <li>• LDom 3: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> <li>• LDom 3: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 44</li> <li>• LDom 2: 20</li> <li>• LDom 3: 20</li> </ul>
<b>B4-4</b>	<ul style="list-style-type: none"> <li>• LDom 1: 24</li> <li>• LDom 2: 24</li> <li>• LDom 3: 24</li> <li>• LDom 4: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> <li>• LDom 3: 4</li> <li>• LDom 4: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 20</li> <li>• LDom 2: 20</li> <li>• LDom 3: 20</li> <li>• LDom 4: 20</li> </ul>

**TABLE:** Cores Available for Base Configuration PDomains (Half-Populated DCU Configurations)

LDom Configuration	Total Number of Cores Available for Domains in This Configuration	Number of Cores Set Aside for Global Zone	Number of Cores Available for Nonglobal Zones
B2-1	LDom 1: 48	LDom 1: 4	LDom 1: 44
B2-2	<ul style="list-style-type: none"> <li>• LDom 1: 24</li> <li>• LDom 2: 24</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 4</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 20</li> <li>• LDom 2: 20</li> </ul>
B2-3	<ul style="list-style-type: none"> <li>• LDom 1: 24</li> <li>• LDom 2: 12</li> <li>• LDom 3: 12</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 4</li> <li>• LDom 2: 2</li> <li>• LDom 3: 2</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 20</li> <li>• LDom 2: 10</li> <li>• LDom 3: 10</li> </ul>
B2-4	<ul style="list-style-type: none"> <li>• LDom 1: 12</li> <li>• LDom 2: 12</li> <li>• LDom 3: 12</li> <li>• LDom 4: 12</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 2</li> <li>• LDom 2: 2</li> <li>• LDom 3: 2</li> <li>• LDom 4: 2</li> </ul>	<ul style="list-style-type: none"> <li>• LDom 1: 10</li> <li>• LDom 2: 10</li> <li>• LDom 3: 10</li> <li>• LDom 4: 10</li> </ul>

When using the information in the tables, keep in mind that the number of cores that are set aside for the global zone applies only when you are creating nonglobal zones on that Database Domain. In that case, a certain number of cores are reserved for the Database Domain (the global zone) and the remaining cores are available for the zones on that Database Domain (the nonglobal zones). If you have a Database Domain with no zones, then all the cores are available for that Database Domain.

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

## Example Configuration

To better understand the number of cores available for zones on certain types of Database Domains, consider the example configuration as described in [“Example Configuration” on page 96](#), where:

- Your Oracle SuperCluster M6-32 consists of one compute server with two fully-populated DCUs.
- The R1\_1 PDomain configuration is set up on the compute server (see [“Four Fully-Populated DCUs \(R1\\_1 PDomain Configuration\)” on page 7](#) for more information).
- The E4-4 LDom configuration (four LDom) is set up on the R1\_1 PDomain configuration (see [“E4-4 LDom Configuration” on page 27](#) for more information).

- The following LDom types are assigned to each LDom in the E4-4 LDom configuration:
  - **LDom 1:** Database Domain (without zones)
  - **LDom 2:** Solaris 10 Application Domain
  - **LDom 3:** Solaris 11 Application Domain
  - **LDom 4:** Database Domain with zones, with four zones

For the purposes of this exercise, assume the following:

- No zones will be set up on LDom 1 (the second Database Domain) in both PDomains.
- Four zones will be set up on LDom 4 (the first Database Domain) in both PDomains.

Using the information provided in the following sections:

- [“Four Fully-Populated DCUs \(R1\\_1 PDomain Configuration\)” on page 7](#)
- [“E4-4 LDom Configuration” on page 27](#)
- [“Cores Available for Domains and Zones” on page 53](#)

You can determine the following information about the two Database Domains on each PDomain:

- 120 total cores are available for LDom 1 (the first Database Domain).
- 24 total cores are available for LDom 4 (the second Database Domain).

The following sections describe how cores could be allocated in both of these Database Domains:

- [“LDom 4 \(Second Database Domain\)” on page 57](#)
- [“LDom 1 \(First Database Domain\)” on page 57](#)

## LDom 1 (First Database Domain)

This Database Domain has 120 cores available to it, but this Database Domain will not contain any zones. The Database Domain is still considered the global zone in this situation, but because there are no zones in this LDom, all 120 cores are available for this Database Domain, so you could allocate all 120 cores to this Database Domain on both PDomains.

## LDom 4 (Second Database Domain)

This Database Domain will contain four zones, which means that:

- The Database Domain will be the **global** zone.

- The four zones within the Database Domain will be **nonglobal** zones.

Because this Database Domain on each PDomain will contain four zones, and it has 24 cores available to it, it would appear at first that you could create four equal-sized zones with regards to the number of cores, where each zone would have six cores ( $24 / 4 = 6$ ).

However, based on the information provided in [“Cores Available for Domains and Zones” on page 53](#), a certain number of cores will be set aside for the global zone. For this specific domain, four cores will be set aside for the global zone, so you actually have 20 cores available for the four zones in this Database Domain. So, for the zones in this Database Domain, you could allocate five cores for each of the four zones to use all 20 cores available for zones in this Database Domain ( $20 / 4 = 5$ ).

---

## Memory Available for LDom

The amount of memory that is available for specific LDom varies, depending on the following factors:

- The PDomain configuration that you have for your Oracle SuperCluster M6-32 (see [“Understanding PDomains” on page 2](#))
- The number of CMUs that are in each DCU, based on the PDomain configuration
- The amount of memory in each CMU (16 GB or 32 GB)

When determining the amount of memory available for a specific LDom, keep in mind that each CMU contains two CMPs, and each CMP contains 32 DIMM slots.

For example, using the example scenario provided in [“Example Configuration” on page 96](#), you would look at the following sections to gather more information about the LDom in this particular example scenario:

- [“Four Fully-Populated DCUs \(R1\\_1 PDomain Configuration\)” on page 7](#)
- [“E4-4 LDom Configuration” on page 27](#)

Using the information provided in those sections, you could determine these things about the four LDom in this configuration:

- **LDom 1:**
  - 10 CMPs (320 DIMM slots) are associated with this LDom
  - Either 5120 GB (16 GB DIMMs) or 10,240 GB (32 GB DIMMs) of memory is available for this LDom
- **LDoms 2, 3, and 4:**
  - 2 CMPs (64 DIMM slots) are associated with each of these LDom

- Either 1024 GB (16 GB DIMMs) or 2048 GB (32 GB DIMMs) of memory is available for each of these LDomS

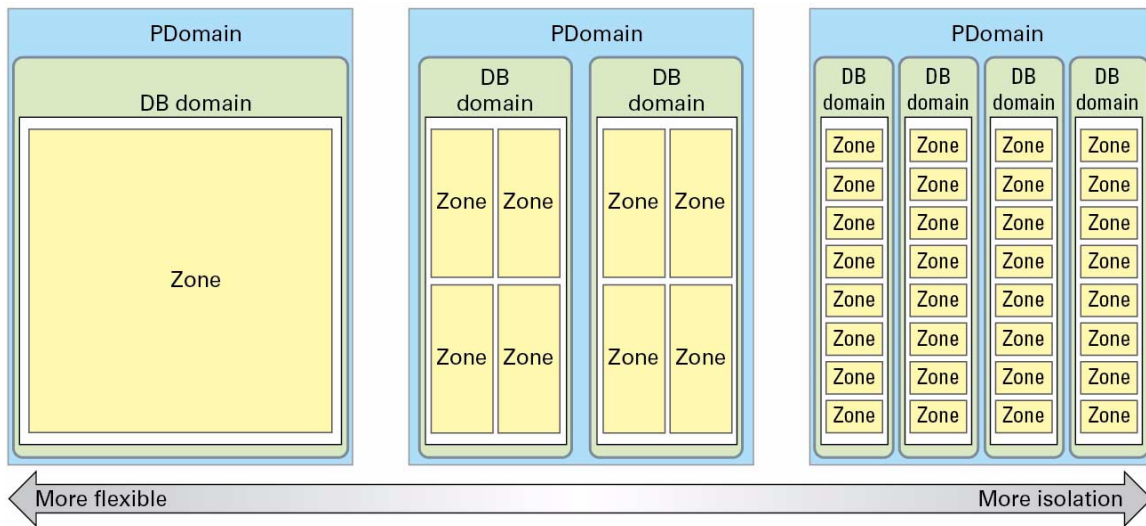
---

## Zones and Cluster Planning for Database Domains

When planning the number and size of the zones on the Database Domains, consider the following points:

- The size of the Database Domain, which will determine the number of cores available for that domain and the maximum number of zones that can be created in that Database Domain
- The number of zones that you want to create in that Database Domain
- The number of cores that you want to assign to each zone in that Database Domain, keeping in mind that two or four cores will always be set aside for the global zone (the Database Domain)
- The clusters that are used for the zones on the Database Domains, where:
  - Each zone that is a member of that cluster should use the same number of cores.
  - Clusters should be created where the zones that are members of a cluster should be on separate PDomains, and the cluster spans across both PDomains so that there is not a single point of failure.

You can have zones that are highly isolated or highly flexible.



## Guidelines for Planning the Number of Zones and Clusters

Following are a list of guidelines regarding zones and clusters on the Database Domains:

- When you first install the operating system instances on a domain, that domain is automatically designated as the global zone. Any zones that you create from that point forward on that domain are nonglobal zones.
- Global zones and nonglobal zones cannot run RAC clusterware in the same domain. When a Database Domain is created and an operating system is installed on the Database Domain, that domain is considered a global zone. The following scenarios describe what is and is not acceptable in that situation:
  - If you have two Database Domains (global zones) and those Database Domains are members of a cluster, you cannot create zones on those Database Domains and have those zones as members of their own clusters. Because those zones would be considered nonglobal zones, you cannot have those nonglobal zones running RAC clusterware within Database Domains that are also running RAC clusterware.
  - If you have two Database Domains (global zones) but those Database Domains are **not** members of a cluster, then you can create zones on those Database Domains and have those zones as members of their own clusters.

- You can create clusters with zones that are of different sizes. However, each zone that is a member of a cluster should be the same size with regards to the number of cores being used by each zone in the cluster. For example, if you are creating a two-node cluster and the first zone in that cluster uses four cores, the other zone in that cluster should also use four cores.
- You can cluster zones together across multiple Database Domains that are of different sizes. For example, you can create a two-node cluster, where there is one zone each on two Database Domains, and those zones are clustered together, even if one of the Database Domains is a different size from the other Database Domain. The private interface connections (the IB connections) are limited to the lower number from the smaller LDom, based on the IB HCAs available to the zone.
- When creating zones on a Database Domain, if the last zone that you are creating on that LDom uses the last remaining core available in that LDom, then that zone is configured with no CPU resource controls. For example, assume you are creating zones on the first LDom in an E4-4 LDom configuration, so four cores would be set aside for the global zone based on the information provided in [“Determining the Cores Available for LDom and Zones” on page 53](#). You could therefore have a maximum of 116 single-core zones on this LDom. In this case, the 116th and final zone in the domain is configured with no CPU resource controls. This zone shares the core and VCPUs with the global zone in this case.
- To minimize remote memory accesses, allocate physical memory to each of the zones in multiples of 16 GB or 32 GB, depending on the DIMM size you have in your system, to align with the size of the memory banks on the server.
- While it is possible to create a cluster with the zones in the same Database Domain or on the same PDomain, that configuration is not a highly-available configuration. Instead, create clusters where the zones are on separate Database Domains, and the cluster spans across multiple PDomains, so that there is not a single point of failure.

---

## Guidelines for Planning the Storage Server Disk Group Layout

An Oracle representative should have set up a cluster on the Database Domains as part of their installation and configuration process. However, when you set up zones on the Database Domains and you set up new clusters with those zones as members, if you want to create clusters that will share the storage servers, the storage servers must be repartitioned. That action erases the clusters that were set up at your initial installation.

Consider these items when planning the storage server disk group layout. For example, how many clusters are you setting up? How many storage servers are available? What level of redundancy do you want for the storage servers for these clusters?

- If three storage servers are available for every cluster with zones on the Database Domains, and you are using normal redundancy, then you do not have to share storage servers across each cluster, and the storage servers do not have to be repartitioned. For high redundancy, five storage servers should be available for every cluster with zones on the Database Domains.

For example, assume you are creating one cluster with zones on the Database Domains and you are using normal redundancy, then you would not have to share storage servers because the storage rack that is part of your Oracle SuperCluster M6-32 has nine storage servers.

Likewise, assume that you have an Oracle Exadata Storage Expansion Full Rack (expansion rack) connected to your Oracle SuperCluster M6-32 (with 18 storage servers). If you are creating eight clusters with zones on the Database Domains and you are using normal redundancy, then you would have 27 storage servers altogether, which means you would have more than three storage servers for every cluster that you are creating. You would not have to share storage servers across the clusters.

- If fewer than three storage servers are available for every cluster with zones on the Database Domains using normal redundancy, or fewer than five storage servers for every cluster if you are using high redundancy, then you must share storage servers across each cluster.

The following topics provide more detail on the process for planning the storage server layout, and an example scenario that describes how certain decisions are made based on an example situations and values.

- [“Planning the Storage Server Disk Group Layout” on page 62](#)
- [“Understanding an Example Scenario” on page 66](#)

## Planning the Storage Server Disk Group Layout

In general, follow this process to determine how to lay out your zones and clusters across the storage servers, where each decision point at the top of the process affects decisions made lower in the decision tree.

1. Determine the amount of existing usable disk space. How many storage servers are available for the clusters that you are creating? How much total usable disk space is available for those clusters?

For more information, see [“Determine the Available Existing Storage Capacity” on page 63](#).



2. Determine the number of clusters that you want for the zones that you are creating on the Database Domains. This, in conjunction with the total amount of existing usable disk space, is used to determine the maximum amount of space that will be available for each cluster.

For more information, see [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster”](#) on page 65.

3. Determine the amount of space to assign to each of the three disk groups in each cluster.
  - DATA
  - RECO
  - DBFS

---

**Note** – There is an option to designate the DBFS disk group set as UNUSED, if necessary.

---

Also determine the names to use for the disk groups for each cluster. For example, if you are creating disk groups for two clusters (osc01r1 and osc01r2), then you might use these names for the three disk groups for each cluster:

- First cluster (osc01r1):
  - DATA\_osc01r1
  - RECO\_osc01r1
  - DBFS\_osc01r1
- Second cluster (osc01r2):
  - DATA\_osc01r2
  - RECO\_osc01r2
  - DBFS\_osc01r2

Record your storage server disk group layout in the worksheet in [“Recording Your Storage Server Disk Group Layout”](#) on page 80.

For more information, see [“Determine the Amount of Space for Each Disk Group in Each Cluster”](#) on page 66.

## Determine the Available Existing Storage Capacity

You must first determine how much existing storage capacity is available before deciding how to carve that storage up for the disk groups for each cluster that you create.

Every storage server contains 12 drives, with the following capacities:

- 1.2 TB High Performance SAS disk
- 4 TB High Capacity SAS disk

In addition, the total number of storage servers varies, depending on the type of Exadata Storage Expansion Rack that you have:

- Storage rack (part of the Oracle SuperCluster M6-32): 9 storage servers
- Optional expansion rack:
  - Quarter rack: 4 storage servers
  - Half rack: 9 storage servers
  - Full rack: 18 storage servers

Finally, the total capacity that you have available depends on the level of redundancy that you want to set for the ASM disk groups (normal redundancy or high redundancy).

The following tables provide capacity information for the Oracle SuperCluster M6-32 storage rack and the expansion rack.

**TABLE:** Storage Server Storage Capacity in the Oracle SuperCluster M6-32 Storage Rack

Capacity Type	4 TB High-Capacity Disks	1.2 TB High-Performance Disks
Raw disk capacity	432 TB	129 TB
Usable mirrored capacity (ASM normal redundancy)	192 TB	58 TB
Usable triple mirrored capacity (ASM high redundancy)	129 TB	39 TB
Raw flash capacity	28.8 TB	28.8 TB

**TABLE:** Storage Server Storage Capacity in the Expansion Rack

Capacity Type	4 TB High-Capacity Disks	1.2 TB High-Performance Disks
Raw disk capacity	<ul style="list-style-type: none"><li>• Quarter Rack: 192 TB</li><li>• Half Rack: 432 TB</li><li>• Full Rack: 864 TB</li></ul>	<ul style="list-style-type: none"><li>• Quarter Rack: 57 TB</li><li>• Half Rack: 129 TB</li><li>• Full Rack: 258 TB</li></ul>
Usable mirrored capacity (ASM normal redundancy)	<ul style="list-style-type: none"><li>• Quarter Rack: 85.3 TB</li><li>• Half Rack: 192 TB</li><li>• Full Rack: 384 TB</li></ul>	<ul style="list-style-type: none"><li>• Quarter Rack: 26 TB</li><li>• Half Rack: 58 TB</li><li>• Full Rack: 116 TB</li></ul>
Usable triple mirrored capacity (ASM high redundancy)	<ul style="list-style-type: none"><li>• Quarter Rack: 57.3 TB</li><li>• Half Rack: 129 TB</li><li>• Full Rack: 258 TB</li></ul>	<ul style="list-style-type: none"><li>• Quarter Rack: 16 TB</li><li>• Half Rack: 39 TB</li><li>• Full Rack: 78 TB</li></ul>
Raw flash capacity	<ul style="list-style-type: none"><li>• Quarter Rack: 12.8 TB</li><li>• Half Rack: 28.8 TB</li><li>• Full Rack: 57.6 TB</li></ul>	<ul style="list-style-type: none"><li>• Quarter Rack: 12.8 TB</li><li>• Half Rack: 28.8 TB</li><li>• Full Rack: 57.6 TB</li></ul>

## Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster

The number of clusters in your system varies, depending on several factors, such as:

- The logical divisions for the databases in your system, where clusters that are created, and zones that are members of those clusters, might have different functions or audiences.
- The number of zones that are available and the number of clusters that can be created with those zones as members, which is partly determined by several factors, such as the number of Database Domains, the size of each Database Domain, and so on.

See [“Zones and Cluster Planning for Database Domains” on page 59](#) for more information on how to decide on the total number of zones and clusters in your system.

Once you determine the number of clusters in your system, you can now determine the maximum amount of space that you can assign to each cluster based on the total amount of usable storage that is available. The total amount of space that you assign to all the clusters in the system cannot exceed the total amount of storage that is available.

For example, assume you have a total amount of 48 TB available as usable storage, and you want to create four clusters total in your system, where each cluster has the same overall size. Then the maximum amount of usable space for each cluster would be 12 TB apiece.

## Determine the Amount of Space for Each Disk Group in Each Cluster

Once you have decided on the total amount of disk space to use for each cluster, you must break that total disk space for each cluster down into disk space that is used for the three disk groups within each cluster:

- DATA
- RECO
- DBFS

The DBFS disk group is normally assigned a substantially smaller amount of total disk space than the other two disk groups. You can also designate the DBFS disk group as UNUSED, if necessary.

For example, if 12 TB of space is available for each cluster, you must carve that 12 TB of space up so that the combined disk space used by all three disk groups is 12 TB or less.

## Understanding an Example Scenario

The best way to illustrate how to decide how to lay out the storage servers is to provide an example scenario, which illustrates the process partly based on the amount of existing storage capacity that is available to you and how you want to carve that existing storage capacity up between the clusters. The following topics provide information on this example scenario:

- [“Determine the Available Existing Storage Capacity \(Example Scenario\)” on page 67](#)
- [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster \(Example Scenario\)” on page 67](#)
- [“Determine the Amount of Space for Each Disk Group in Each Cluster \(Example Scenario\)” on page 68](#)

## Determine the Available Existing Storage Capacity (Example Scenario)

Consider the following example scenario:

- Oracle SuperCluster M6-32 storage rack with nine storage servers
- 12 high-capacity (4 TB) disk drives in each storage server
- Each storage server set at normal (double-mirrored) redundancy

Using the information from the section [“Determine the Available Existing Storage Capacity” on page 63](#), you have **192 TB** of usable storage capacity available for the storage servers in your Oracle SuperCluster M6-32 storage rack.

## Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster (Example Scenario)

In this example scenario, for the number of clusters that are used in this system, assume that you decided to create two two-node clusters for the zones on the Database Domains.

If you have 192 TB of usable disk capacity, and you want to create two clusters for the zones, you could use the existing 192 TB of space in one of several ways:

- Use the entire 192 TB of existing space for the two clusters in one of these ways:
  - Create two equal-sized clusters, where the total amount of space used by the disk groups in each cluster would be 96 TB.
  - Create two clusters of different sizes, where the total disk space used by the disk groups for the two clusters equals 192 TB. For example, one cluster with the total disk space used by the disk groups equaling 150 TB and the other cluster with the total disk space used by the disk groups equaling 42 TB.

Keep in mind that by using the entire 192 TB of existing space for the two clusters in this example scenario, you leave yourself no room for future expansion, either for adding zones and clusters to the system or for expanding the amount of space available to the existing clusters.

- Use a portion of the existing 192 TB of usable space for the two clusters that you will be creating, while reserving some of the usable space for future expansion.

This example scenario creates two equal-sized clusters, so when allotting disk sizes to each disk group within each cluster, verify that the total amount of space used by all three disk groups does not exceed 96 TB for each cluster.

## Determine the Amount of Space for Each Disk Group in Each Cluster (Example Scenario)

You must now determine how much disk space is allotted to the three disk groups in each cluster:

- DATA
- RECO
- DBFS

The amount of disk space that you use for each disk group depends on your database requirements. For this example scenario, assume that you have decided to allot the following amounts of disk space to the three disk groups for each cluster:

- **DATA** – Roughly 2 TB, or 2000 GB
- **RECO** – Roughly 1 TB, or 1000 GB
- **DBFS** – Roughly 200 GB

In this example scenario, the total amount of space that you use for the disk groups for each cluster is roughly 3.2 TB, which is far below the maximum amount of 96 TB that is available for each cluster, based on the calculations from the previous section.

---

## Recording Your Existing and Planned Configuration Information

Use the following worksheets to record the existing configuration that you have for your Oracle SuperCluster M6-32, which affects the number and size of the zones that you can create. Also use the following worksheets to record the decisions about the number of zones and clusters that you are deploying, and the layout that you are using for the storage servers.

- [“Recording Your Zone Configuration Information” on page 68](#)
- [“Recording Your Cluster Configuration Information” on page 79](#)
- [“Recording Your Storage Server Disk Group Layout” on page 80](#)

## Recording Your Zone Configuration Information

Use the tables in the following sections to record the following information for each Database Domain:

- The number of zones that you want to deploy on each Database Domain
- The size of each zone that you want to deploy on each Database Domain (the number of cores)
- The host names and IP addresses that will be assigned to each zone, if you know this information ahead of time, for the following networks:
  - Host management network
  - Client access network
  - IB network
  - VIP

The worksheet for each Database Domain has entries for 20 zones. If you have more than 20 zones on each Database Domain, copy the worksheets to record the information for the additional zones. The number of cores that you can allocate to zones varies, depending on the size of that zone. See [“Determining the Cores Available for LDOMs and Zones”](#) on page 53 for more information.

Enter the zone configuration information in the following worksheets:

- [“Database Domain 1 Zone Configuration Information”](#) on page 69
- [“Database Domain 2 Zone Configuration Information”](#) on page 72
- [“Database Domain 3 Zone Configuration Information”](#) on page 74
- [“Database Domain 4 Zone Configuration Information”](#) on page 77

## Database Domain 1 Zone Configuration Information

Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
1	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
2	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
3	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
4	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
5	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
6	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
8	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
9	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
10	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
11	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
12	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:



<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
13	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
14	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
15	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
16	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
17	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
18	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
19	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
20	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

## Database Domain 2 Zone Configuration Information

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
1	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
2	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
3	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
4	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
5	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
6	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
8	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
9	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
10	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
11	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
12	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
13	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
14	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
15	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
16	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
17	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
18	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
19	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
20	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

## Database Domain 3 Zone Configuration Information

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
1	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
2	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
3	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
4	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
5	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
6	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
8	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
9	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
10	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
11	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
12	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
13	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
14	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
15	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
16	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
17	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
18	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
19	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
20	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

## Database Domain 4 Zone Configuration Information

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
1	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
2	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
3	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
4	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
5	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
6	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
8	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

<b>Zone Number</b>	<b>Number of Cores Assigned to This Zone</b>	<b>Host Names</b>	<b>IP Addresses</b>
9	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
10	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
11	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
12	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
13	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
14	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
15	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
16	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
17	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:



Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
18	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
19	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
20	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

## Recording Your Cluster Configuration Information

Use this table to record the following information for each cluster that you are configuring:

- Name of cluster (for example, *osc01r1* or *osc01r2*)
- Zones that are members of this cluster, and the Database Domains where they reside
- Host name and IP addresses to be used for SCAN for each cluster

RACs should span across multiple PDomains for high availability. If you are creating a cluster, and a zone in a Database Domain is a member of that cluster on one PDomain, then other zones that are members of that cluster should be in a similar Database Domain on another PDomain.

For example, assume you are creating a two-node cluster, where zone 1 in each Database Domain is a member of that cluster. To ensure high availability, create this cluster with the following zones as members of that cluster:

- Zone 1 in Database Domain 1 in PDomain 1
- Zone 1 in Database Domain 1 in PDomain 2

The worksheet has entries for ten clusters. But you can have more than ten clusters in your Oracle SuperCluster M6-32, depending on the number of Database Domains that have zones that are members of each cluster, the size of the zones in each cluster, and so on. Copy the worksheet if you have more than ten clusters in your system.

Cluster Number and Name	Zone Membership Information	SCAN Host Names and IP Addresses for Each Cluster
1	Zone number: Database Domain number:	Host name: IP address:
2	Zone number: Database Domain number:	Host name: IP address:
3	Zone number: Database Domain number:	Host name: IP address:
4	Zone number: Database Domain number:	Host name: IP address:
5	Zone number: Database Domain number:	Host name: IP address:
6	Zone number: Database Domain number:	Host name: IP address:
7	Zone number: Database Domain number:	Host name: IP address:
8	Zone number: Database Domain number:	Host name: IP address:
9	Zone number: Database Domain number:	Host name: IP address:
10	Zone number: Database Domain number:	Host name: IP address:

## Recording Your Storage Server Disk Group Layout

Use the following table to record your storage server disk group layout for each cluster. Note that the worksheet has entries for ten clusters, but you could have more than ten clusters in your Oracle SuperCluster M6-32, depending on the number of Database Domains that have zones that are members of each cluster, the size of the zones in each cluster, and so on. Copy the worksheet if you have more than ten clusters in your system.

Assign a name to each disk group that designates it as a specific type of disk group for a specific cluster, as described in [“Planning the Storage Server Disk Group Layout” on page 62](#). For example, if the name of a particular cluster is osc01r1, then name the disk groups for that cluster something similar to the following:

- DATA\_osc01r1

- RECO\_osc01r1
- DBFS\_osc01r1

Cluster Number and Name	Name of Disk Groups	Size of Disk Groups	Total Size of All Disk Groups for This Cluster
1	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
2	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
3	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
4	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
5	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
6	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
7	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
8	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
9	DATA: RECO: DBFS:	DATA: RECO: DBFS:	
10	DATA: RECO: DBFS:	DATA: RECO: DBFS:	



# Preparing to Configure Zones on Database Domains

---

These topics describe how to prepare to configure zones on the Database Domains.

- [“Prepare to Configure Zones” on page 83](#)
- [“Update the Base Software” on page 85](#)
- [“Determine the Repository Location” on page 85](#)
- [“Install or Update Packages From the Remote Repository” on page 86](#)
- [“Install or Update Files From the Local Repository” on page 88](#)
- [“Verify Configuration Tool Installation” on page 91](#)

---

## ▼ Prepare to Configure Zones

Complete these tasks before setting up zones and clusters on the Database Domains with zones in the Oracle SuperCluster M6-32.

### 1. **Back up the data on your storage servers.**

You might have to repartition the drives in the storage servers that are accessed by the zones that you create in this document. See [“Guidelines for Planning the Storage Server Disk Group Layout” on page 61](#) for more information. Refer to the storage server documentation for more information.

### 2. **Plan the zone and cluster layout.**

You must make the following decisions on the zones and the clusters before beginning the processes in this document:

- The number of zones that you want to create on the Database Domains with zones
- The size of each zone that you want to create
- The number of clusters that you want to create, and the number of zones that you want to have as members of each cluster

- The starting IP addresses for the following networks for each zone and cluster that you are creating:
  - Administration network
  - Client access network
  - Private IB network
  - Backup/Data Guard network (optional)

### 3. Plan the storage server layout.

You must make the following decisions on the storage server layout before beginning the processes in this document:

- The number of storage servers and the amount of disk space available for the zones and clusters that you will be creating
- The size of the following disk groups for each cluster:
  - DATA
  - RECO
  - DBFS

---

**Note** – You can opt to designate the DBFS disk group as UNUSED, where the Oracle Cluster Registry (OCR) and voting disks are part of the DATA disk group.

---

### 4. Gather the necessary networking information and configuration files for your existing configuration.

At the time of your initial installation, the Oracle representatives should have provided a document with all the following information:

- Name of the Oracle SuperCluster M6-32
- IP addresses and host names for Oracle ILOM for the following components:
  - PDomain SPP with Database Domains that you are setting up zones on
  - Storage servers
- IP addresses and host names for the 1GbE host management administration network for the following components:
  - Storage servers
  - Switches (IB leaf and spine switches, and the Cisco switch)
  - PDUs
- IP addresses and host names for the IB network for the storage servers

In addition, as part of the initial installation when your Oracle SuperCluster M6-32 was first set up, the Oracle installer should have saved a configuration file from OEDA that contains all the configuration information for your system. You need this configuration file in order to set up zones on Database Domains with zones.

5. **Update the base software.**

See “Update the Base Software” on page 85.

---

## ▼ Update the Base Software

1. **Log in to the Oracle SuperCluster M6-32.**
2. **Log in to MOS.**  
<http://support.oracle.com>
3. **Select the Patches & Updates tab.**
4. **In the Patch Search panel, select Product or Family (Advanced).**
5. **In the Product field, search for and select Oracle SuperCluster M6-32.**  
Select any of the Oracle SuperCluster M6-32 options.
6. **In the Release field, select the Oracle SuperCluster M6-32 1.6 release or later.**
7. **Click the Search button.**
8. **Select the Quarterly Full Stack Download Patch For SuperCluster (*release-date*) (Patch) option.**
9. **Click the Read Me button to get installation instructions for the patch.**
10. **Click the Download button to download the ZIP files for the patch.**

---

## ▼ Determine the Repository Location

Your Oracle SuperCluster M6-32 accesses the files that you need to set up zones on the Database Domains using one of the following repositories:

- Files are accessed remotely on Oracle’s IPS repository
- Files are accessed locally on the ZFS storage appliance in the Oracle SuperCluster M6-32 storage rack

This procedure allows you to determine how you access the files that you need to set up the zones on the Database Domains.

1. **Log in to a Database Domain.**

## 2. Type.

```
# pkg publisher
```

## 3. In the output for that command, locate the URI field at the end of the output.

- If you see one of the the following outputs in the URI field, then your system is set up to access files remotely, from Oracle’s IPS repository. Go to [“Install or Update Packages From the Remote Repository”](#) on page 86.

```
https://pkg.oracle.com/solaris/exa-family
```

```
https://pkg.oracle.com/solaris/support
```

- If you see one of the following outputs in the URI field, where *repository-location* is not `pkg.oracle.com`, but rather a local IPS repository that serves packages over `http`, then your system is set up to access files locally, from the ZFS storage appliance in the Oracle SuperCluster M6-32 storage rack. Go to [“Install or Update Files From the Local Repository”](#) on page 88.

```
file:///repository-location
```

```
http://repository-location
```

---

# ▼ Install or Update Packages From the Remote Repository

## 1. On the first Database Domain to contain zones, determine if you have the necessary packages to set up zones on the Database Domains.

```
# pkg info supercluster supercluster/ssc-exavm supercluster/iscsi  
supercluster/vnet supercluster/ssctuner
```

- If all packages listed in the command are installed, then you have the correct packages. Go to [Step 2](#).
- If you see that one or more of the packages listed in the command are *not* installed, then you must get the packages that you need. Go to [Step 4](#).



## 2. Verify that the `exa-family` publisher is configured.

```
# pkg publisher exa-family
```

If the output shows that a publisher for `exa-family` is not configured, add the `exa-family` publisher.

### a. Go to the following location.

<https://pkg-register.oracle.com/>

### b. In the Certificate Requests page, select the Oracle Exadata Database Machine option.

### c. In the Certificate Information page, download the key and certificate, and follow the instructions on the page to install the certificate.

## 3. Update those packages to the latest versions.

### a. Update the packages from Oracle's IPS repository on this Database Domain.

```
# pkg update pkg://exa-family/system/platform/supercluster
# pkg update pkg://exa-family/system/platform/supercluster/iscsi
# pkg update pkg://exa-family/system/platform/supercluster/ssc-exavm
# pkg update pkg://exa-family/system/platform/supercluster/ssctuner
# pkg update pkg://exa-family/system/platform/supercluster/vnet
```

### b. Repeat [Step a](#) for every Database Domain to contain zones.

When you have updated the packages on every Database Domain to contain zones, go to ["Verify Configuration Tool Installation"](#) on page 91.

## 4. Install the necessary packages.

```
# pkg install pkg://exa-family/system/platform/supercluster
# pkg install pkg://exa-family/system/platform/supercluster/iscsi
# pkg install pkg://exa-family/system/platform/supercluster/ssc-exavm
# pkg install pkg://exa-family/system/platform/supercluster/ssctuner
# pkg install pkg://exa-family/system/platform/supercluster/vnet
```

## 5. Repeat [Step 2](#) through [Step 4](#) for every Database Domain to contain zones.

## 6. Verify that the configuration tools are installed.

Go to ["Verify Configuration Tool Installation"](#) on page 91.

---

## ▼ Install or Update Files From the Local Repository

This procedure describes how to install or update files from the local repository on the ZFS storage appliance in the Oracle SuperCluster M6-32 storage rack.

---

**Note** – For more information on setting up a repository on the ZFS storage appliance, refer to MOS note 1503899.1.

---

1. Type `ssh` to get in to the active ZFS storage controller through the 1GbE host management network.
2. Locate the project for the local repository on the ZFS storage appliance.

```
stor-contr:> shares
stor-contr:shares> show
```

The project for the repository should have an obvious name, such as `IPS-repos`, and should be shown towards the end of the output from the `show` command.

3. Select that project for the repository.

```
stor-contr:shares> select repository
```

For example:

```
stor-contr:shares> select IPS-repos
```

4. List the file systems in the repository area.

```
stor-contr:shares IPS-repos> show
```

5. For every file system that is displayed in the `show` command, select each file system in this area and set the `root` permissions to 755.

For example, assume that you see the following file systems in the output for the `show` command:

- `solaris`
- `s11repo`

- `exa-family`

Then type the following commands to set the root permissions to 755 for each of those file systems, ending with the `end` command after you have set the root permissions for the last file system in the list:

```
stor-contr:shares IPS-repos> filesystem solaris
stor-contr:shares IPS-repos/solaris (uncommitted)> set root_permissions=755
stor-contr:shares IPS-repos/solaris (uncommitted)> commit
stor-contr:shares IPS-repos/solaris (uncommitted)> done
stor-contr:shares IPS-repos> filesystem s11repo
stor-contr:shares IPS-repos/s11repo (uncommitted)> set root_permissions=755
stor-contr:shares IPS-repos/s11repo (uncommitted)> commit
stor-contr:shares IPS-repos/s11repo (uncommitted)> done
stor-contr:shares IPS-repos> filesystem exa-family
stor-contr:shares IPS-repos/exa-family (uncommitted)> set root_permissions=755
stor-contr:shares IPS-repos/exa-family (uncommitted)> commit
stor-contr:shares IPS-repos/exa-family (uncommitted)> done
stor-contr:shares IPS-repos> end
```

## 6. On the first Database Domain to contain zones, determine if you have the necessary packages to set up zones on the Database Domains.

```
# pkg info supercluster/ssc-exavm supercluster/iscsi supercluster/vnet  
supercluster/ssctuner
```

- If you see that one or more of the packages listed in the command are *not* installed, then you must get the packages that you need. Go to [Step 7](#).
- If all four packages listed in the command above are installed, then you have the files that you need. Follow these instructions to determine if you have the latest versions of the packages:

- a. Log in to MOS.

<http://support.oracle.com>

- b. Select the Patches & Updates tab.

- c. In the Patch Search panel, select Product or Family (Advanced).

- d. In the Product field, search for and select Oracle SuperCluster M6-32.

Select any of the Oracle SuperCluster M6-32 options.

- e. In the Release field, select all releases.

- f. Choose Description contains in the pull-down menus, then type `exa-family` in the search field.

- g. Click the Search button.



d. Update the local repository on the ZFS storage appliance.

```
# mkdir -p tmp-dir
# mount -F hsfs iso-name.iso tmp-dir
# pkgrecv -s /tmp-dir/repo -d path-to-repo-on-ZFSSA/repo '*'
# umount tmp-dir
# pkgrepo rebuild -s path-to-repo-on-ZFSSA/repo
```

where:

- *tmp-dir* is the name of the temporary directory that you use as a mount point.
- *iso-name.iso* is the name of the ISO image.
- *path-to-repo-on-ZFSSA* is the path to the repository on the ZFS storage appliance. This path appears after the `file://` output from the `pkg publisher` command that you entered in the section [“Determine the Repository Location”](#) on page 85.

The single quotes at the end of the second line (\*) are quotes formed when you select the unshifted double-quotes (") key on your keyboard (single quotes that are either straight up and down or pointing left).

10. Install the packages onto each Database Domain with zones where you are setting up zones.
  - a. Log in to the first Database Domain with zones.
  - b. Install the necessary packages.

```
# pkg install pkg://exa-family/system/platform/supercluster
# pkg install pkg://exa-family/system/platform/supercluster/iscsi
# pkg install pkg://exa-family/system/platform/supercluster/ssc-exavm
# pkg install pkg://exa-family/system/platform/supercluster/ssctuner
# pkg install pkg://exa-family/system/platform/supercluster/vnet
```

11. Repeat these steps for every Database Domain.
12. When you have installed the packages on every Database Domain that will contain zones, go to [“Verify Configuration Tool Installation”](#) on page 91.

---

## ▼ Verify Configuration Tool Installation

1. Log in to the first Database Domain to contain zones.
2. Download the latest Oracle OneCommand utility.

- a. **Locate the My Oracle Support note that provides information on the latest version of the patch with the Oracle OneCommand utility.**

Refer to the following My Oracle Support note and locate the Oracle Exadata Deployment Assistant (OEDA) section in this note:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=888828.1>

- b. **Select the most recent patch from the table.**



---

**Caution** – You must download patch 19766190 or later.

---

Typically, the most recent patch is shown at the top of the table, and will have a higher number than other patches. For example, Patch 19766190 would be a more recent version of the patch than Patch 19571935.

- c. **Select the appropriate option from the Platform or Language field for the operating system on the machine that you are using for OEDA and Java OneCommand.**

The options are:

- Apple Mac OS X (Intel) (64-bit)
- Microsoft Windows (32-bit)
- Linux x86-64
- Oracle Solaris on x86-64 (64-bit)
- Oracle Solaris on SPARC (64-bit)

- d. **Click the Download button to download the patch.**

- e. **Unzip the patch file.**

Several files become available.

3. **Change directories to the `/opt/oracle.supercluster` directory.**

```
# cd /opt/oracle.supercluster
```

4. **Verify that the necessary tools and scripts are available in the `bin` and `zonetools` subdirectories.**

- To verify the tools and scripts are in the `bin` subdirectory, type.

```
# ls bin
```

You should see the following tools and scripts:

- `addvnet-wrapper.sh`

- `addvnet.sh`
- `applyconfig_ssc.sh`
- `dcli`
- `gen_cellaffinity_ssc.sh`
- `ib_set_node_desc_ssc.sh`
- `iscsi-lun.sh`
- `iscsi-zpool.sh`
- `m6_bonds_ssc.sh`
- `rmvnet.sh`
- `setcoremem`
- `setup_ssh_eq.sh`
- `zonemanifest.sh`
- To verify that the tools and scripts are in the `zonetools` subdirectory, type.

```
# ls zonetools
```

Verify that you have the `ssc_exavm` script in the `zonetools` directory.

5. Repeat [Step 1](#) through [Step 4](#) for all other Database Domains to contain zones.
6. Use the OEDA to create the configuration files to be used by the `exavmgen` tool and the OneCommand utility later on in this process.

Go to [“Preparing to Configure Zones on Database Domains”](#) on page 83.





# Creating Configuration Files (OEDA)

---

These topics describe how to create configuration files using the `config.sh` tool (OEDA).

- “Example Configuration” on page 96
- “Verify Storage Server Disk Space for Additional Zones” on page 97
- “Locate the Necessary Files to Set Up Zones” on page 98
- “Set Up the OEDA” on page 100
- “Start the OEDA” on page 101
- “Import the Most Recent OEDA Configuration File” on page 102
- “Review Existing Configuration Information” on page 104
- “Review the Information in the Identify Compute Node Operating System Page” on page 105
- “Review the Information in the Management and Private Networks Page” on page 107
- “Complete the Define Clusters Page” on page 109
- “Set Zone Default Configurations” on page 114
- “Complete the Cluster Page” on page 117
- “Complete the Cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks Page” on page 121
- “Verify Remaining Configuration Information” on page 123
- “Generate the Configuration Files” on page 123

---

# Example Configuration

The following configuration is used as an example throughout this exercise:

- Your Oracle SuperCluster M6-32 consists of one compute server with two fully-populated DCUs
- The R1\_1 PDomain configuration is set up on the compute server (see [“Four Fully-Populated DCUs \(R1\\_1 PDomain Configuration\)”](#) on page 7 for more information)
- The E4-4 LDom configuration (four LDom) is set up on the R1\_1 PDomain configuration (see [“E4-4 LDom Configuration”](#) on page 27 for more information)
- The following LDom types are assigned to each LDom in the E4-4 LDom configuration:
  - **LDom 1** — Database Domain (without zones)
  - **LDom 2** — Solaris 10 Application Domain
  - **LDom 3** — Solaris 11 Application Domain
  - **LDom 4** — Database Domain with zones, with four zones
- Five cluster instances are set up on the R1\_1 PDomain configuration:
  - One cluster instance for the two Database Domains (without zones)
  - One cluster instance for each zone on the Database Domain with zones (four cluster instances total for the zones)
- The storage rack has storage servers with 1.2 TB high performance disks

For each cluster, the example configuration assumes that you have zones within the Database Domain with zones, or Database Domains (without zones), that reside on different PDomains. So for the zones on Database Domains with zones, these zones are members of the first cluster:

- First zone in the Database Domain with zones in the *first* PDomain (DCU 0 and DCU 1)
- First zone in the Database Domain with zones in the *second* PDomain (DCU 2 and DCU 3)

Throughout these procedures, it is assumed that this configuration was created during your initial installation, and that you are adding zones to the existing Database Domains with zones.

---

## ▼ Verify Storage Server Disk Space for Additional Zones

Before beginning the process of creating additional zones on existing Database Domains with zones, you must first verify that you have the appropriate amount of disk space available on the storage servers:

- If you are not sharing storage servers across clusters and you have entire storage servers available for the additional zones that you are adding to the existing Database Domains with zones, then you do not have to do anything more here. You can assign the entire storage servers to the new zones as part of this procedure.
- If you are sharing storage servers across clusters, you must determine how much free space is available on the storage servers that you are using for these new zones before proceeding.

### 1. Determine the amount of free space available on the storage servers.

```
CellCLI> list celldisk attributes name, freespace, freespacemap
```

You should see output similar to the following:

```
CD_00_etc25celadm01 366.6875G ((offset=162.046875G,size=366.6875G))
CD_01_etc25celadm01 366.6875G ((offset=162.046875G,size=366.6875G))
CD_02_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_03_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_04_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_05_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_06_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_07_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_08_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_09_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_10_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_11_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
FD_00_etc25celadm01 0
FD_01_etc25celadm01 0
FD_02_etc25celadm01 0
FD_03_etc25celadm01 0
FD_04_etc25celadm01 0
FD_05_etc25celadm01 0
FD_06_etc25celadm01 0
FD_07_etc25celadm01 0
FD_08_etc25celadm01 0
```

```
FD_09_etc25celadm01 0
FD_10_etc25celadm01 0
FD_11_etc25celadm01 0
FD_12_etc25celadm01 0
FD_13_etc25celadm01 0
FD_14_etc25celadm01 0
FD_15_etc25celadm01 0
```

2. Look at the *CD\_number* entries to determine the amount of available space on each of the storage servers.
3. Use the information on the amount of available space on each of the storage servers to determine how much space you can use for the disk groups for each new cluster you are creating for the new zones.

You need enough space for these disk groups for each new cluster that you are creating:

- DATA
- RECO
- DBFS (if necessary)

4. Locate the necessary files to set up your system.

See [“Locate the Necessary Files to Set Up Zones”](#) on page 98.

---

## ▼ Locate the Necessary Files to Set Up Zones

1. Locate the folder that contains the most recent configuration files.
  - If this is the first time after the initial installation of your system that you are creating additional zones on Database Domains with zones, locate the folder that contains the configuration files that your Oracle installer saved at the time of the initial installation of your system.
  - If you have used these instructions previously to create additional zones on Database Domains with zones, locate the folder that contains the most recent configuration files from that previous configuration.

The configuration files are typically in a folder named `ExadataConfigurations`.

## 2. Locate the necessary files within the folder that contains the most recent configuration files.

The folder should contain the following files:

- Individual configuration files for every zone that was set up on Database Domains with zones (shown as *name-prefix-zone-name-zoneConfig.xml*)
- Individual configuration files for every cluster that was set up, either for zones or for Database Domains without zones (shown as *customer-name-name-prefix-cluster-clucluster-number.xml*)
- A single master file that contains the configuration information for all the zones, Database Domains, and clusters that were set up on your system (shown as *customer-name-name-prefix.xml*)

For example, assume you had zones and clusters created during your initial installation as described in [“Example Configuration” on page 96](#), and you provided the following information during your initial installation:

- Customer name: Acme
- Name prefix: osc01

Then you should see configuration files similar to the following in the configuration folder:

- Individual configuration files for the eight zones that were set up on Database Domains with zones:
  - *osc01zdbadm01-zoneConfig.xml*: Configuration file for the first zone
  - *osc01zdbadm02-zoneConfig.xml*: Configuration file for the second zone
  - *osc01zdbadm03-zoneConfig.xml*: Configuration file for the third zoneAnd so on for the configuration files for the remaining five zones.
- Individual configuration files for the five clusters that were set up:
  - *Acme-osc01-cluster-clu1.xml*: Configuration file for the first cluster (the first cluster with zones)
  - *Acme-osc01-cluster-clu2.xml*: Configuration file for the second cluster (the second cluster with zones)
  - *Acme-osc01-cluster-clu3.xml*: Configuration file for the third cluster (the third cluster with zones)
  - *Acme-osc01-cluster-clu4.xml*: Configuration file for the fourth cluster (the fourth cluster with zones)
  - *Acme-osc01-cluster-clu5.xml*: Configuration file for the fifth cluster (the cluster for the Database Domains without zones)

- A single master configuration file that contains all the configuration information for your system:
  - `Acme-osc01.xml`

You will be using one of these configuration files when you use OEDA or Java OneCommand to set up zones on the Database Domains.

---

## ▼ Set Up the OEDA



---

**Caution** – Ensure that you backup all existing databases before running the Java OneCommand.

---



---

**Caution** – You must use the latest version of OEDA and the Java OneCommand patch (patch 19766190 or later). Refer to the Oracle Exadata Deployment Assistant (OEDA) section in MOS Note 888828.1 for details.

---



---

**Caution** – Any version of Java OneCommand prior to patch 19766190 can destroy the storage cell disks and griddisk if an undo option is performed on certain steps (Create Cell disk, for example). This can cause complete destruction of all the griddisks on the storage cells. In addition, re-running the griddisk creation step or mistakenly specifying a non-unique diskgroup in OEDA will result in the destruction of existing griddisks. Be aware that older versions of Java OneCommand also destroy cell disks and griddisks with the Create Cell Disks step.

---

### 1. Prepare to configure zones.

See [“Preparing to Configure Zones on Database Domains”](#) on page 83.

### 2. Ensure that the machine running the OEDA has Oracle JRE 1.6 or later.

### 3. Locate the OEDA ZIP files (Ocmd) for your operating system.

The OEDA files are part of the Ocmd ZIP file that you downloaded earlier in this process. See [“Verify Configuration Tool Installation”](#) on page 91 for more information.

There are different OEDA ZIP files for these operating systems:

- Linux
- MacOS
- Solaris (SPARC)

- Solaris (Intel)
  - Windows
4. **Unzip the `Ocmd` ZIP file for your operating system on the first Database Domain that will deploy zones.**
  5. **Go into the directory that was created after you unzipped the `Ocmd` ZIP file for your operating system.**
  6. **Start the OEDA.**  
Go to [“Start the OEDA” on page 101.](#)

---

## ▼ Start the OEDA

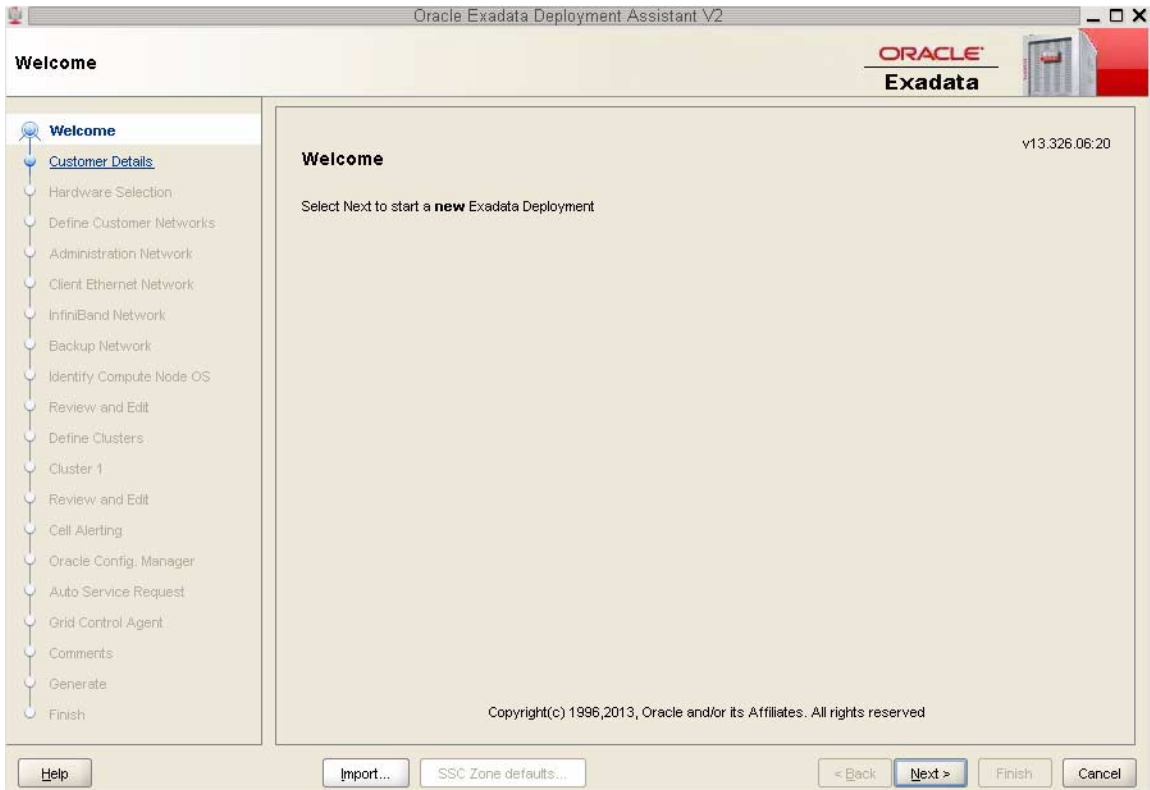
1. **Start the OEDA.**
  - On Linux, MacOS, or the Oracle Solaris OS, type:

```
config.sh
```

- On Microsoft Windows, type:

```
config.cmd
```

The OEDA Welcome window is displayed.



## 2. Import the OEDA XML input file.

Go to [“Import the Most Recent OEDA Configuration File”](#) on page 102.

---

# ▼ Import the Most Recent OEDA Configuration File

### 1. At the bottom of the Welcome page, click the Import button.

A pop-up window is displayed, with the default location at the OneCommand directory.

### 2. Open the folder that contains the most recent configuration files for your system.



- If this is the first time after the initial installation of your system that you are creating additional zones on Database Domains with zones, locate the folder that contains the configuration files that your Oracle installer saved at the time of the initial installation of your system.
- If you have used these instructions previously to create additional zones on Database Domains with zones, locate the folder that contains the most recent configuration files from that previous configuration.

See “Locate the Necessary Files to Set Up Zones” on page 98 for more information.

**3. Within that folder, import the single master configuration file that contains all the configuration information for your system.**

This file is typically named *customer-name.name-prefix.xml*. For example, assume you provided the following information during your initial installation:

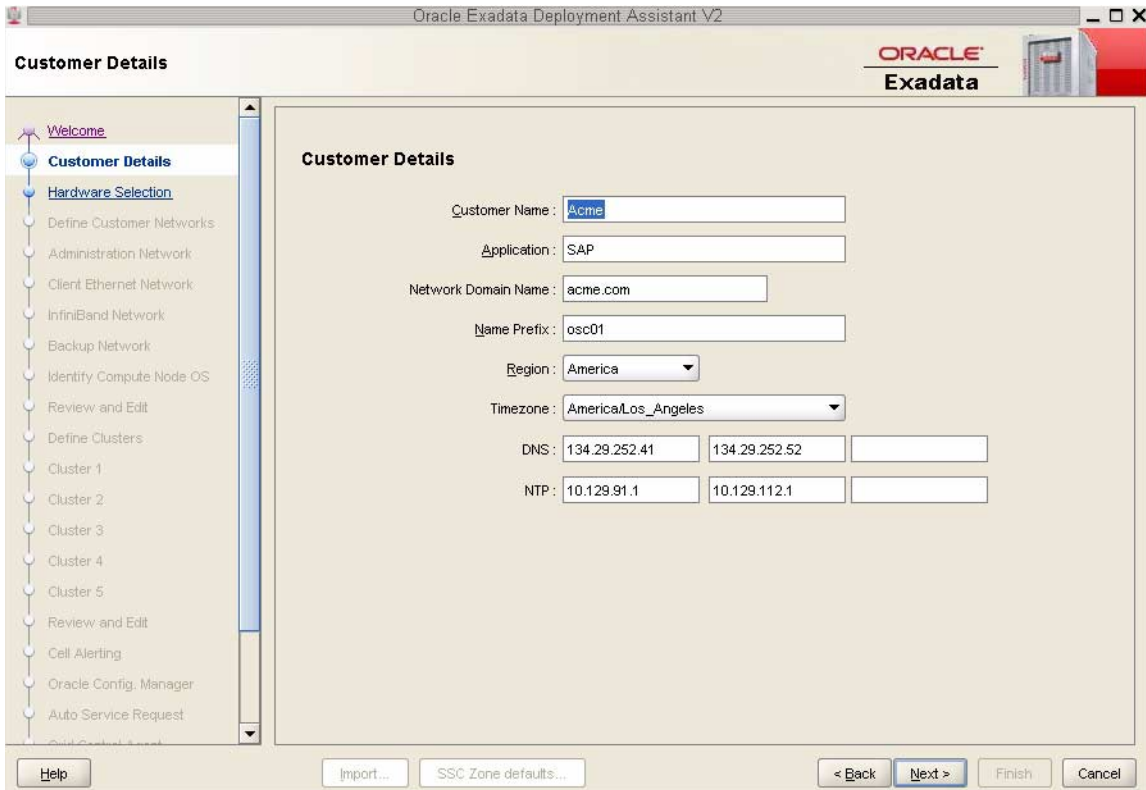
- **Customer name:** Acme
- **Name prefix:** osc01

Then the configuration file for this cluster would be named *Acme.osc01.xml*.

A pop-up window appears, confirming that the configuration files were imported successfully. Click OK.

**4. In the Welcome window, click Next.**

The Customer Details page is displayed.



## 5. Verify the existing configuration information.

Go to [“Review Existing Configuration Information”](#) on page 104.

---

# ▼ Review Existing Configuration Information

After you import the most recent master configuration file, the fields in each of the pages in OEDA are populated with information that you provided for that configuration, such as the IP addresses and host names for each network for all of the components and domains in your Oracle SuperCluster M6-32. If you are adding zones to a Database Domain with zones on a system that was correctly configured, do not modify any of the information in the first set of screens in OEDA.

### 1. Review the existing configuration information.

Go through the following screens in OEDA, reviewing the configuration information that is provided and clicking Next at the bottom of each page.

---

**Note** – Do not change any of the information in the following screens.

---

- Customer Details page
- Hardware Selection page
- Define Customer Networks page
- Administration Network page
- Client Ethernet Network page
- InfiniBand Network page
- Backup/Data Guard page

Stop when the Identify Compute Node Operating System page is displayed.

### 2. Review the information in the Identify Compute Node Operating System page.

Go to [“Review the Information in the Identify Compute Node Operating System Page”](#) on page 105.

---

## ▼ Review the Information in the Identify Compute Node Operating System Page

### 1. Review the information in the Identify Compute Node Operating System page.

---

**Note** – Do not change any of the information in the Identify Compute Node Operating System page.

---

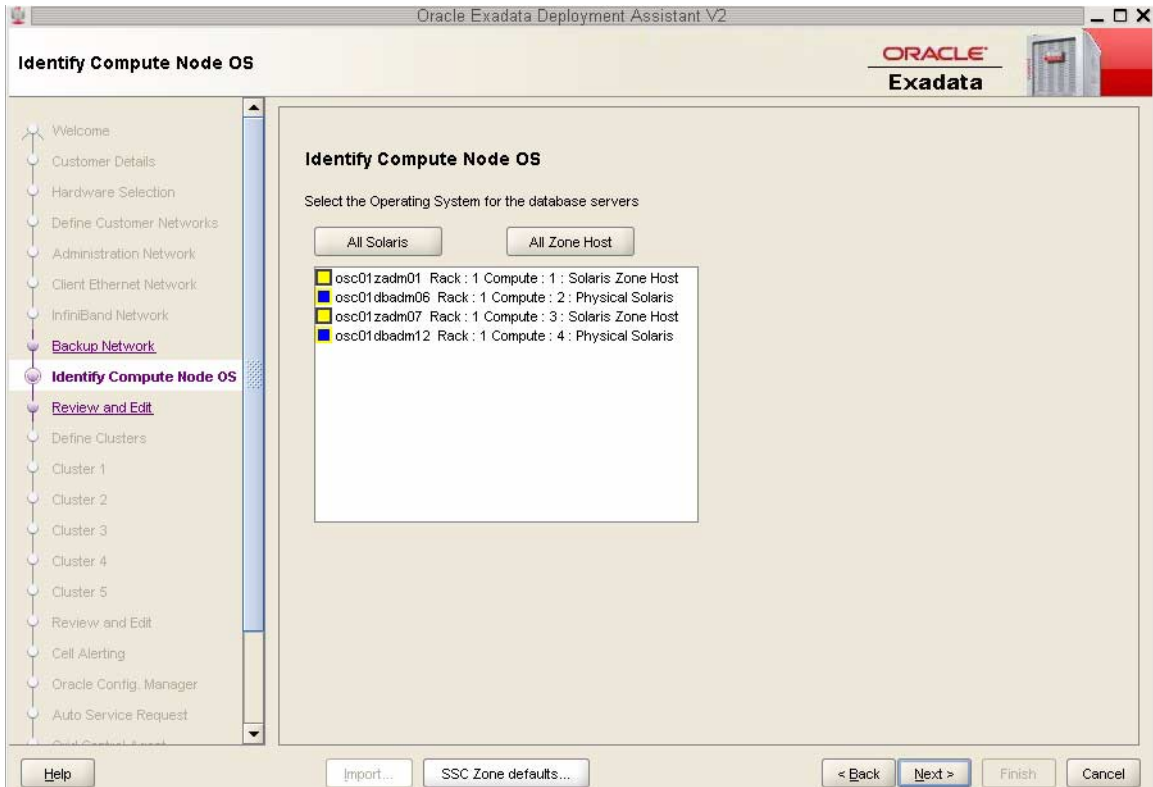
The Identify Compute Node Operating System page should display the total number of Database Domains with zones and Database Domains (without zones) on your Oracle SuperCluster M6-32 in your existing configuration. This page does not display information on any Applications Domains in your existing configuration.

For example, assume you have these LDoms as part of your existing configuration, as described in [“Example Configuration”](#) on page 96:

- One Database Domain (without zones) on each PDomain (2 total)
- One Database Domain with zones on each PDomain (2 total)

- One Oracle Solaris 10 Application Domain on each PDomain (2 total, not displayed in this page)
- One Oracle Solaris 11 Application Domain on each PDomain (2 total, not displayed in this page)

Then you should see the following screen:



In this screen, these domains are displayed in this manner:

- **Database Domain with zones:** Shown as Solaris Zone Host
- **Database Domain (without zones):** Shown as Physical Solaris

---

**Note** – The zones within the Database Domains with zones will not be displayed on this page. Only the Database Domains with zones that house the zones are displayed.

---

## 2. Click Next.

The Review and Edit page is displayed.

3. **Review the information in the Management and Private Networks page.**

Go to [“Review the Information in the Management and Private Networks Page”](#) on page 107.

---

## ▼ Review the Information in the Management and Private Networks Page

1. **Review the information in the Management and Private Networks page.**

---

**Note** – Do not change any of the information in the Management and Private Networks page.

---

The Management and Private Networks page displays management and IB network address and host name information for the following components in your existing configuration:

- Database Domains with zones (shown as Compute Node # - Zone Host)
- Database Domains (without zones) (shown as Compute Node #)
- Exadata Storage Servers in the Oracle SuperCluster M6-32
- Switches and PDUs in the Oracle SuperCluster M6-32

Oracle Exadata Deployment Assistant

**Review and Edit**

ORACLE  
Exadata

**Review and Edit Management and Private Networks**

This page captures node specific data for the Management, ILOM and Private Networks for the Compute Nodes, Storage Cells and the switches used in the Rack.  
The Client, VIP, SCAN and backup network names/IP address are collected later in the interview process.

Re-Generate Data

<b>Exadata SuperCluster T5 Compute Node 1 - Zone Host</b>		Rack 1 - Rack Location 10
Admin Name:	<input type="text" value="osc01zadm01.acme.com"/>	Admin IP: <input type="text" value="10.204.74.100"/>
Ilom Name:	<input type="text" value="osc01-node1-ilom.acme.com"/>	Ilom IP: <input type="text" value="10.204.74.132"/>
Priv Name:	<input type="text" value="osc01zadm01-priv1.acme.com"/>	Priv IP: <input type="text" value="192.168.10.1"/>
<b>Exadata SuperCluster T5 Compute Node 2</b>		Rack 1 - Rack Location 27
Admin Name:	<input type="text" value="osc01dbadm06.acme.com"/>	Admin IP: <input type="text" value="10.204.74.103"/>
Ilom Name:	<input type="text" value="osc01-node1-ilom.acme.com"/>	Ilom IP: <input type="text" value="10.204.74.132"/>
Priv Name:	<input type="text" value="osc01db06-priv1.acme.com"/>	Priv IP: <input type="text" value="192.168.10.5"/>
<b>Exadata SuperCluster T5 Compute Node 3 - Zone Host</b>		Rack 1 - Rack Location 43
Admin Name:	<input type="text" value="osc01zadm07.acme.com"/>	Admin IP: <input type="text" value="10.204.74.108"/>
Ilom Name:	<input type="text" value="osc01-node2-ilom.acme.com"/>	Ilom IP: <input type="text" value="10.204.74.133"/>
Priv Name:	<input type="text" value="osc01zadm07-priv1.acme.com"/>	Priv IP: <input type="text" value="192.168.10.23"/>
<b>Exadata SuperCluster T5 Compute Node 4</b>		Rack 1 - Rack Location 44
Admin Name:	<input type="text" value="osc01dbadm12.acme.com"/>	Admin IP: <input type="text" value="10.204.74.111"/>
Ilom Name:	<input type="text" value="osc01-node2-ilom.acme.com"/>	Ilom IP: <input type="text" value="10.204.74.133"/>
Priv Name:	<input type="text" value="osc01db12-priv1.acme.com"/>	Priv IP: <input type="text" value="192.168.10.27"/>
<b>Exadata Cell Node HP 1.2TB 1</b>		Rack 1 - Rack Location 2
Admin Name:	<input type="text" value="osc01celadm01.acme.com"/>	Admin IP: <input type="text" value="10.204.74.116"/>
Ilom Name:	<input type="text" value="osc01celadm01-ilom.acme.com"/>	Ilom IP: <input type="text" value="10.204.74.134"/>
Priv Name:	<input type="text" value="osc01cel01-priv1.acme.com"/>	Priv IP: <input type="text" value="192.168.10.45"/>

Help    Import...    SSC Zone defaults...    Save    < Back    Next >    Finish    Cancel

2. Click Next after reviewing the information in this page.

The Define Clusters page is displayed.

3. Complete the Define Clusters page.

Go to “Complete the Define Clusters Page” on page 109.

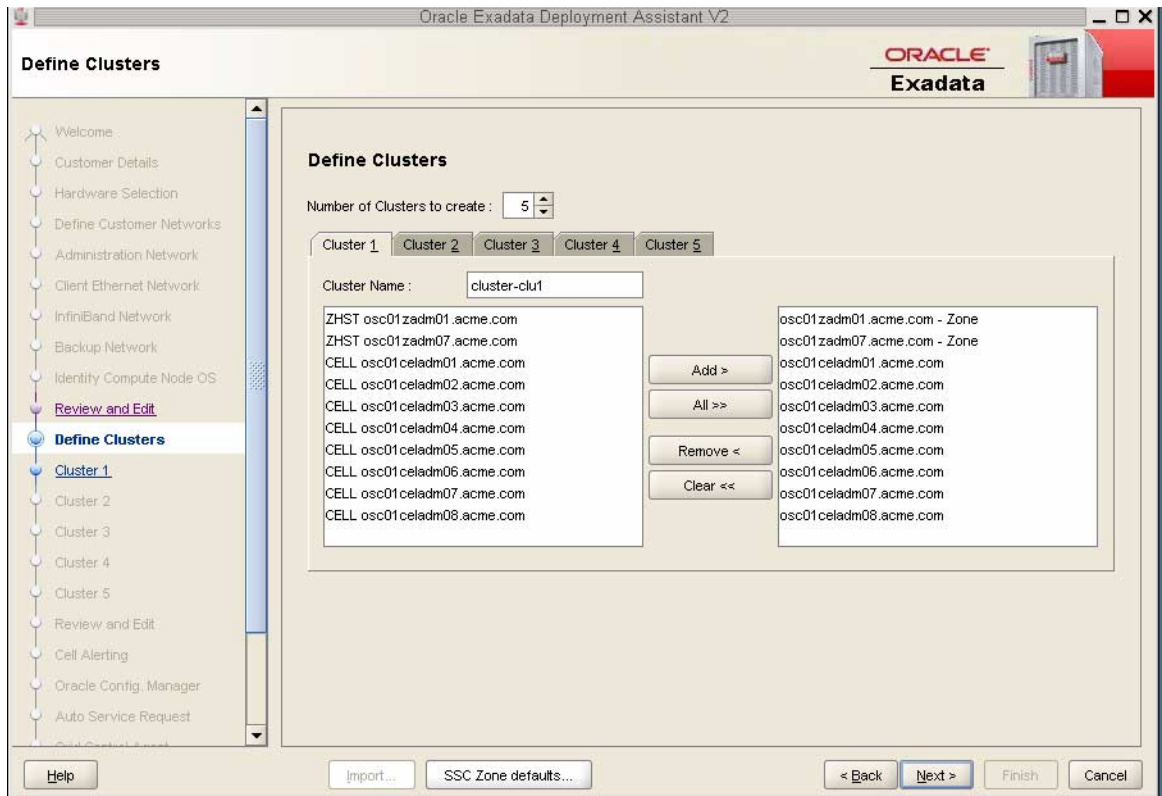
## ▼ Complete the Define Clusters Page

### 1. Review the information in the Define Clusters page.

The Define Clusters page should show your current configuration. For example, assume you have these clusters as part of your existing configuration, as described in “Example Configuration” on page 96:

- Four clusters that were set up for the Database Domains with zones (four zones within each Database Domain with zones, so four clusters for the zones in the Database Domains with zones)
- One cluster that was set up for the Database Domains (without zones)

Then you should see the following screen:



Where:

- Cluster 1 should display the cluster that was set up for the Database Domains (without zones)

- Clusters 2 to 5 should display the clusters that were set up for the Database Domains with zones (four zones within each Database Domain with zones, so four clusters for the zones in the Database Domains with zones)

For Clusters 2 to 5, the screens for each cluster show the information for the Database Domains with zones that house the zones, not the zones themselves.

In this example configuration, the screens for Clusters 2 to 5 show the following information:

- `osc01zadm01.acme.com`
- `osc01zadm07.acme.com`

Those are the Database Domains with zones that house the zones that make up the clusters. Information on the individual zones within each Database Domain with zones appear in a future screen.

Because Cluster 1 is for the two Database Domains (without zones), that screen would show those two Database Domains and not the Database Domains with zones that are shown in Clusters 2 to 5.

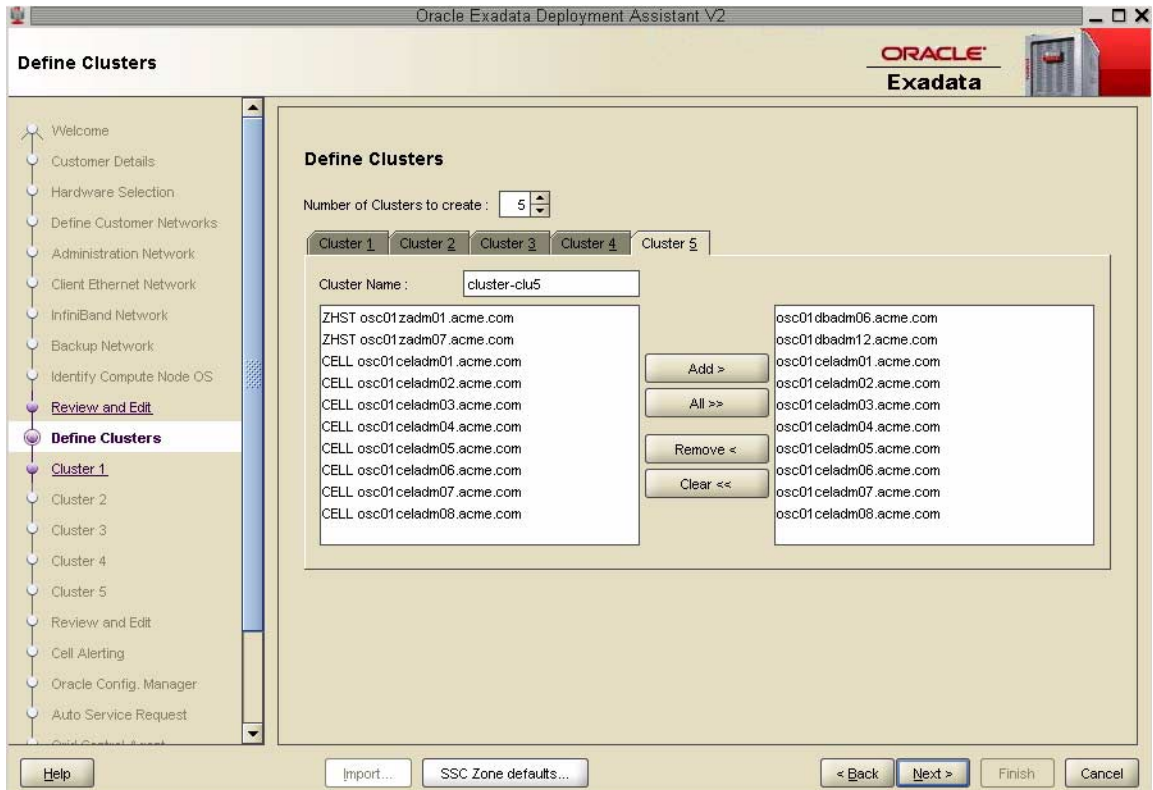
In this example configuration, the screen for Cluster 1 shows the following information:

- `osc01dbadm06.acme.com`
- `osc01dbadm12.acme.com`

**2. Click on the tab for the last cluster shown in this page and note the cluster name for that cluster.**

Using the previous example, click on the tab for Cluster 5 and note that the cluster name is `acme-osc01-clu5`, as shown in the following figure.





**3. Determine the number of zones and clusters that you want to add to the Database Domains with zones.**

Using the example configuration outlined in “Example Configuration” on page 96, assume you want to add the following zones and clusters:

- Two additional zones to each Database Domain with zones (4 zones total)
- Two additional clusters, one cluster for each new zone set up on each Database Domain with zones

**4. In the Number of Clusters to create area, increase the number of clusters by the number of new clusters that you want to add.**

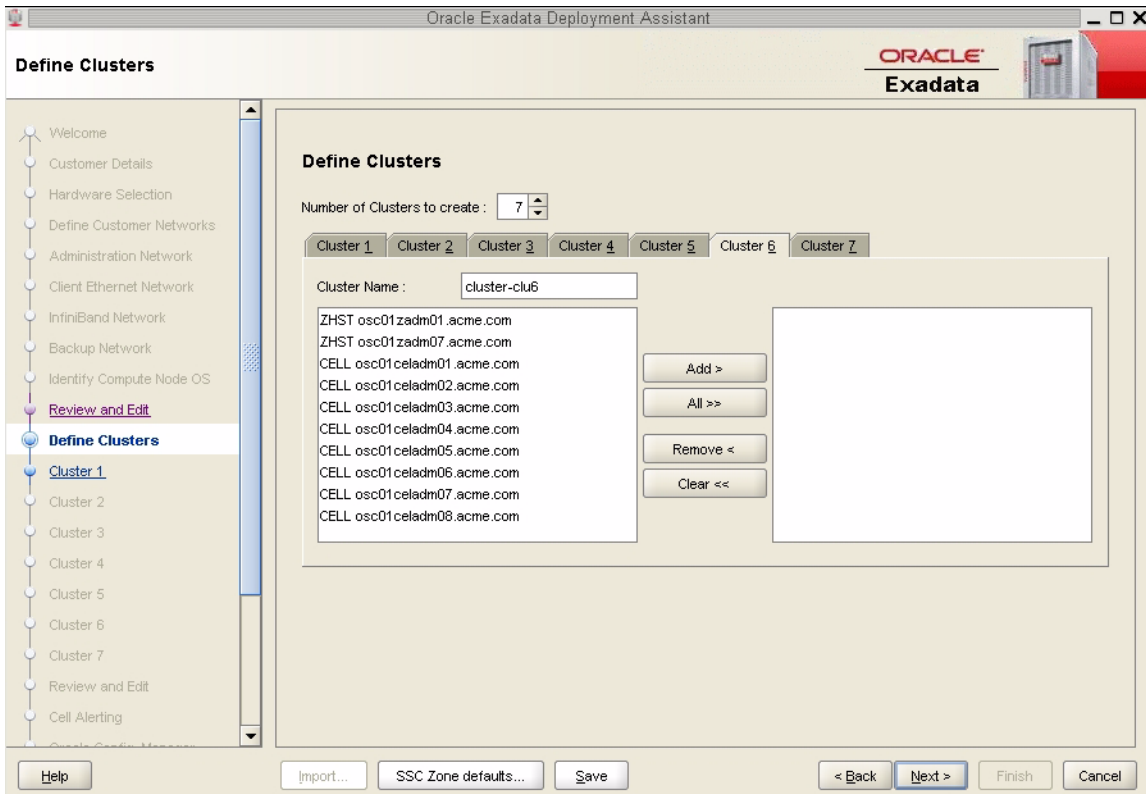
For example, if you are adding two new clusters, you would increase the number of clusters to seven total clusters:

- Five existing clusters
- Two new clusters

**5. For each cluster, set a unique name for the cluster and select the Exadata Storage Servers that will be part of each cluster.**

- a. Click on the tab for the first new cluster that you want to add.

In this example, you would click on Cluster 6. Note that that cluster was automatically named `cluster-clu6` in the Cluster Name field, as shown in the following figure.



- b. Set a unique name within your organization for the cluster.

By default, a cluster name of `cluster-clunumber` will be assigned to the cluster (such as `cluster-clu6`). You may run into an issue with Enterprise Manager if you have multiple systems with the same default cluster names across the systems, so you must change the default cluster name to a unique name within your organization for the cluster (for example, *organization-systemname-clusternumber*, such as `acme-osc01-clu6`).

- c. Select the Database Domains with zones that house the zones that you want to create and click Add.

The Database Domains with zones are shown as ZHST in the left pane in the Define Clusters page. The Database Domains with zones that you selected appear in the right pane.

- d. Select the Exadata Storage Servers that are part of this cluster and click Add.

---

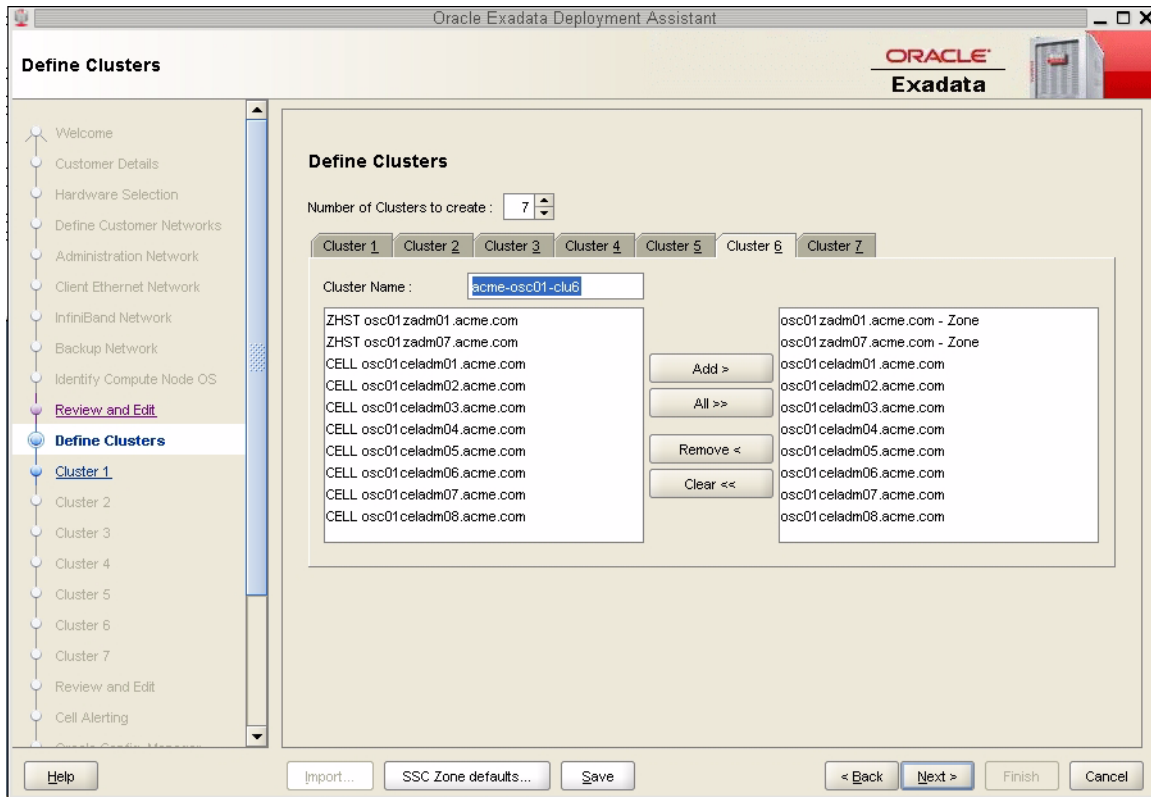
**Note** – Oracle Support Best Practices dictate that you should assign all Exadata Storage Servers to each cluster rather than split the Exadata Storage Servers between the clusters.

---

**Note** – If you are sharing Exadata Storage Servers across clusters, you must first confirm that the Exadata Storage Servers that you are adding to this cluster have enough available space to use for these new zones. See [“Verify Storage Server Disk Space for Additional Zones”](#) on page 97 for more information.

---

The following figure shows an example of a new cluster with all the Exadata Storage Servers assigned to that cluster.



6. Repeat [Step 5](#) to assign the Exadata Storage Servers to all of the new clusters that you will be creating.

7. When you have assigned Database Domains with zones and Exadata Storage Servers to all new clusters, click Next.

The Cluster page for the first new cluster is displayed.

8. Set the zone default configurations for each cluster that you want to create.

Go to [“Set Zone Default Configurations”](#) on page 114.

---

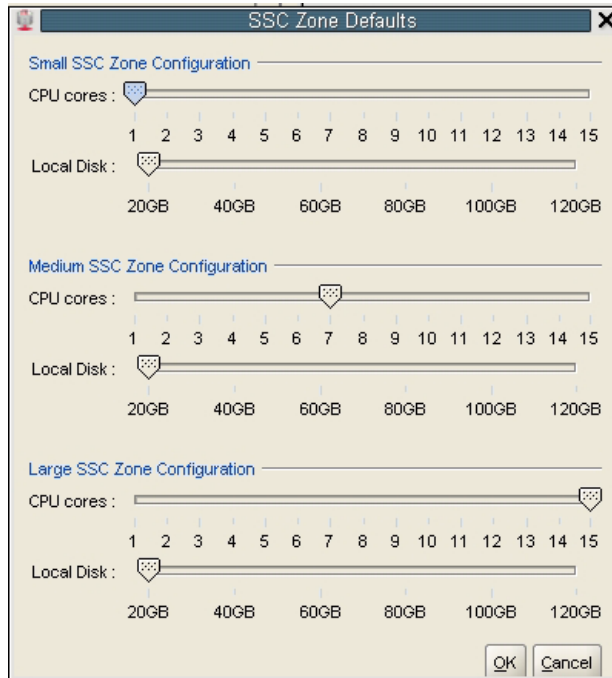
## ▼ Set Zone Default Configurations

Before you begin to enter information about each cluster that you want to create, you must first assign values to the default zone configurations that you use when setting up the clusters.

There are three choices for the zone configurations (small, medium, or large) when you begin entering information for each of the clusters that you are creating. Each of these choices assigns a different number of CPU cores and local disk size to the zones in a particular cluster. This section describes how you can change the number of CPU cores and the size of the local disk to the value that you want for each of these zone configuration choices.

1. In the Cluster page for the first cluster, click SSC Zone Defaults at the bottom of the page.

The SSC Zone Defaults page appears.



**2. Review the default values for the number of CPU cores and the disk space to be allocated for the small, medium, and large zone configurations.**

The following default values are used for the zone configurations:

- Small zone configuration:
  - 1 core
  - 20 GB of disk space
- Medium zone configuration:
  - 7 cores
  - 20 GB of disk space
- Large zone configuration:
  - 12 cores
  - 20 GB of disk space

**3. Modify the values for the number of CPU cores and the disk space to be allocated for the small, medium, and large zone configurations.**

For each of the zone configuration sizes, move the slider to set the appropriate value for the number of CPU cores and the amount of disk space to be allocated for zones in a cluster with this configuration size.

Note the following important points when deciding the values to assign to these options:

- The following values are available for both options:
  - **Number of CPU cores:** Between 1 and 12 cores
  - **Amount of disk space:** Between 20 GB and 120 GB of disk space
- When deciding on the number of cores to be allocated, keep in mind that two to four cores will be reserved for the global zone (the Database Domain) and the remaining cores are available for the zones in that Database Domain. See [“Determining the Cores Available for LDoms and Zones” on page 53](#) for more information.
- When deciding on the amount of disk space to be allocated, the amount you select defines the disk size only for the Oracle Home directory. Select an amount greater than 20 GB, because that amount of disk space is typically not sufficient for the Oracle Home directory.

In addition, keep in mind that an additional 45 GB of disk space is automatically allocated for the root file system, so the actual size of the disk space being used is the amount you select plus the additional 45 GB allocated for the root file system.

So, for example, if you select 100 GB of disk space, the zone LUN is 145 GB (100 GB of disk space for the Oracle Home directory and Oracle binaries, and 45 GB of space for the root file system).

Use the options in this window to define your own small, medium, and large zone configurations, which are used as individual templates when you assign values to zones within the clusters in the Cluster page for each cluster in the section [“Complete the Cluster Page” on page 117](#). The values you assign to each of these types of zone configurations are completely user-definable, so you could assign the same values to all three types of zone configurations if you wish.

For example, assume that you want the following values for clusters with a medium-sized zone configuration:

- 8 CPU cores
- 80 GB of disk space, where:
  - 80 GB of disk space will be used for the Oracle Home directory
  - An additional 45 GB of disk space will be automatically assigned to the root file system

In the Medium SSC Zone Configuration section, move the slider to assign those values in the CPU cores and Local Disk areas. Doing so means that these values are used whenever you select the medium-sized zone configuration for any cluster that you enter information for in subsequent screens.

**4. When you have set up the values that you want for each sized zone configuration, click OK.**

**5. Complete the Cluster page for each cluster that you want to create.**

Go to [“Complete the Cluster Page” on page 117](#).

---

## ▼ Complete the Cluster Page



---

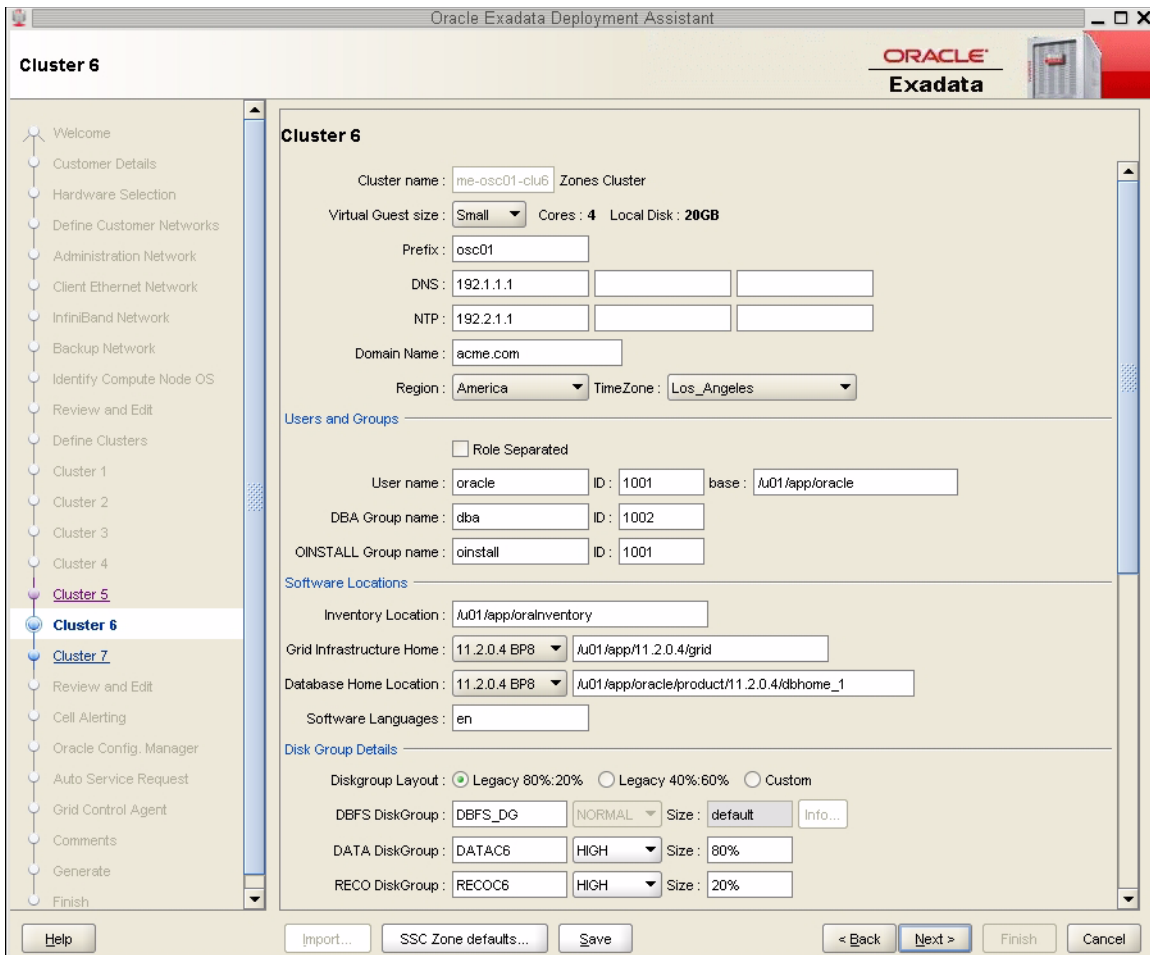
**Caution** – Take extreme care when specifying the DATA and RECO diskgroup names. The names must be unique for the cluster you are defining. Check existing griddisks on the storage cells and ensure they are not already in use. If an existing griddisk is specified for a new cluster, the griddisk will be dropped and recreated by older versions of Java Onecommand (prior to patch 19766190), possibly resulting in loss of production data.

---

In the left pane, there is a cluster configuration page for each new cluster that you are creating as part of this exercise.

1. **Review the cluster configuration page for the first new cluster that you are creating.**

For example, using the example configuration outlined in [“Example Configuration” on page 96](#), review the cluster configuration page for Cluster 6, which is the first new cluster that you are creating. Skip over the cluster configuraton pages for the first five clusters, which were already set up as part of your initial installation, until you get to the cluster configuration page for Cluster 6.



2. Determine the number of cores that are available in the Database Domains with zones or Database Domains (without zones) for this cluster.

3. Determine the number of cores that you want to allocate to the members of this cluster.

- If this cluster is for zones in Database Domains with zones, determine the number of cores that you want to allocate to the zones to be members of this cluster. Keep in mind the other zones that are also on these Database Domains with zones when allocating cores for zones in a particular cluster, and note that two to four cores must be set aside for the global zone. See [“Determining the Cores Available for LDOMs and Zones”](#) on page 53 for more information.
- If this cluster is for Database Domains (without zones), all cores that are available for this domain are allocated to this Database Domain (without zones).



**4. Assign the appropriate zone configuration size for the zones in this cluster.**

The zone configuration size can be small, medium, or large. The number of CPU cores and the amount of local disk space allocated to the zones in this cluster for each zone configuration size depends on any changes that you might have made when you completed the procedures in “Set Zone Default Configurations” on page 114.

---

**Note** – You must make a selection in the Virtual Guest size area, even if the zone configuration size that you want is shown on the screen by default. The values for the number of cores and the amount of local disk space do not update to the values that you assigned in the section “Set Zone Default Configurations” on page 114, unless you make a selection in the Virtual Guest size area.

---

**5. In the Software Locations area, select the database versions for this cluster in the following areas:**

- Grid Infrastructure Home
- Database Home Location

**6. In the Disk Group Details area, provide the following information:**

- A selection in the Diskgroup Layout area:
  - Legacy 80%:20% : Assigns the following sizes to the disk groups:
    - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)
    - DATA Diskgroup: 80% size
    - RECO Diskgroup: 20% size
  - Legacy 40%:60% : Assigns the following sizes to the disk groups:
    - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)
    - DATA Diskgroup: 40% size
    - RECO Diskgroup: 60% size
  - Custom : Allows you to assign your own sizes to the disk groups
- Unique name for the disk groups (DBFS, DATA, and RECO) for this cluster. For example, for the first new cluster, which is cluster 6:
  - DBFS Diskgroup: DBFS\_DGC6
  - DATA Diskgroup: DATAC6
  - RECO Diskgroup: RECOC6

- Level of redundancy for the disk groups (DBFS, DATA, and RECO) for this cluster:
  - UNUSED (DBFS disk group only)
  - NORMAL
  - HIGH
- Size of the disk groups (DBFS, DATA, and RECO) for this cluster

---

**Note** – If you are sharing Exadata Storage Servers across clusters, follow the instructions in [“Verify Storage Server Disk Space for Additional Zones” on page 97](#) to determine the amount of free space that you have in the Exadata Storage Servers for each disk group.

---

**7. Make cluster-specific changes to the Database Name field, if necessary.**

For example, because this is the sixth cluster, you might use `dbm06` for the Database Name.

**8. Change the Start IP address to the first available, unused IP address for each of the networks.**

There are sections in the cluster configuration page for these three networks:

- Admin Network
- Client Network
- Private Network (the IB network)

For each network, change the Start IP address to the first available, unused IP address.

A certain number of IP addresses are assigned to each of the new clusters for each network, determined by the Pool Size field next to the Start IP field for each network. Keep that range of IP addresses in mind when creating the new clusters.

For example, if the first available starting IP address for the client network for the first new cluster is `172.16.8.150`, and the Pool Size field next to the Start IP field for that network shows 7, then IP addresses `172.16.8.150` to `172.16.8.156` are assigned to the client access network for this cluster. Therefore, `172.16.8.157` is the next available IP address for the client network for the next cluster, if necessary.

**9. Review the rest of the information in this page and make changes as necessary.**

**10. Click Next when you have completed the page for this cluster.**

The Cluster page for the next cluster that you want to create is displayed.

**11. Repeat [Step 1](#) through [Step 10](#) for each cluster.**

12. **Click Next when you have set up the last new cluster.**

The cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks page is displayed.

13. **Complete the cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks page.**

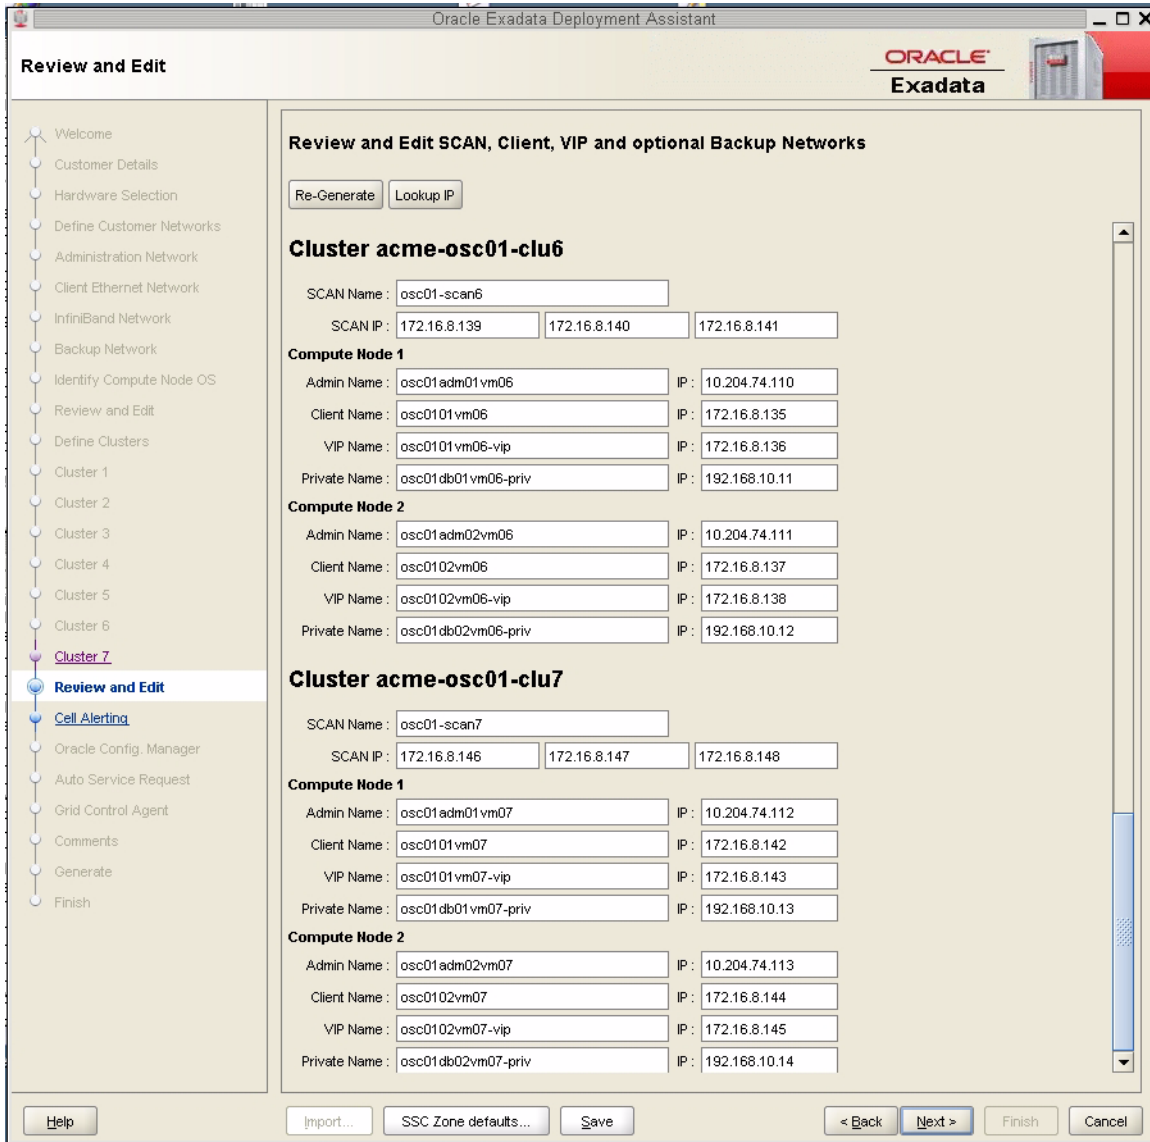
Go to [“Complete the Cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks Page”](#) on page 121.

---

## ▼ Complete the Cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks Page

1. **Review the information in the cluster Review and Edit SCAN, Client, VIP, and Optional Backup Networks page, and make any necessary changes.**

The new clusters are displayed, with two new zones in in each Database Domain with zones as part of those clusters, as shown in this figure.



## 2. Click Next.

The Cell Alerting page is displayed.

---

**Note** – If an error message appears when you click Next in the Review and Edit SCAN, Client, VIP and Optional Backup Networks page, indicating that there are duplicate IP addresses, correct those duplicate IP addresses in the Review and Edit page and click Next again to go to the Cell Alerting page.

---

**3. Verify that the remaining configuration information is correct.**

Go to [“Verify Remaining Configuration Information”](#) on page 123.

---

## ▼ Verify Remaining Configuration Information

**1. Verify that the configuration information in the following screens are correct:**

- Cell Alerting page
- Oracle Config. Manager page
- Auto Service Request
- Grid Control Agent

**2. In the Comments page, provide additional information that might be useful for this deployment, then click Next.**

**3. Generate the configuration files.**

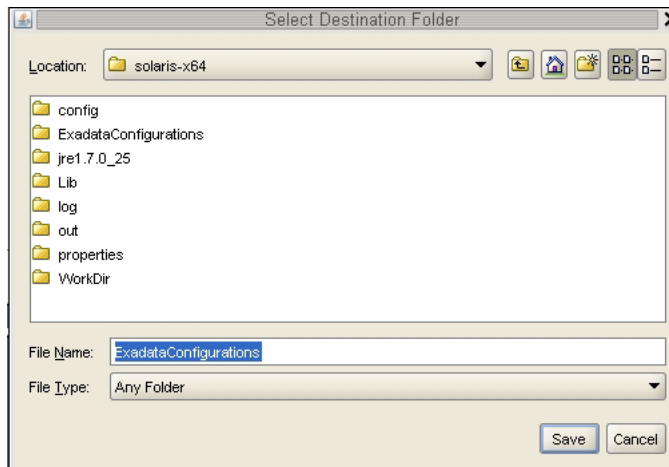
Go to [“Generate the Configuration Files”](#) on page 123.

---

## ▼ Generate the Configuration Files

**1. In the Generate page, click Next.**

A pop-up window appears.



**2. Navigate to the appropriate directory, enter the name for the configuration folder for this deployment, and click Save.**

A confirmation window appears with a link to the configuration folder that you just saved.

**3. Click Finish.**

The OEDA closes.

**4. Verify that the new zone and cluster files are in the configuration folder for this deployment.**

If you navigate to the configuration folder with the new configuration files for this deployment, you should now see the following files:

- Four new individual configuration files for the four zones that you added to the Database Domains with zones (shown as *name-prefix-zone-name-zoneConfig.xml*)
- Two new individual configuration files for the two clusters that you added (shown as *customer-name-name-prefix-cluster-clucluster-number.xml*)
- A single updated master file that contains the configuration information for all the zones, Database Domains, and clusters that are set up on your system (shown as *customer-name-name-prefix.xml*)

---

**Note** – The *zone-name* portion of the configuration files for the new zones that you added might differ from the *zone-name* portion of the configuration files for the zones that were set up on your system previously.

---

**5. Create the template zone on each Database Domain.**

Go to [“Creating the Template Zone on Each Database Domain”](#) on page 125.

# Creating the Template Zone on Each Database Domain

---

Creating the template zone on each Database Domain expedites the process of creating multiple zones. When you create the template zone on each Database Domain, the core Oracle Solaris packages are installed on each Database Domain. When zones are created later on, the core Oracle Solaris packages can be cloned from the template zone onto each of the zones that are created on the Database Domains. This process can significantly reduce the amount of time it takes to install core Oracle Solaris packages on each zone.

- [“Create a Template Zone on a Database Domain” on page 125](#)
- [“Delete a Template Zone From a Database Domain” on page 128](#)

---

## ▼ Create a Template Zone on a Database Domain

1. **Log in to the first Database Domain in your Oracle SuperCluster M6-32 where you are setting up zones.**
2. **Determine if the template zone is already in place on the first Database Domain where you are setting up zones.**

The template zone may already be in place from the initial installation of your system. Type the following command to determine if the template zone is already in place:

```
zoneadm list -cv
```

- If you see the following output, where a file is listed as *DB-domain\_T* (for example, *osc01dbadmin01-01\_T*), then the template file is already in place from the initial installation.

```
ID NAME STATUS PATH BRAND IP
0 global running / solaris shared
5 osc01dbadmin01.us.oracle.com running
/zoneHome/osc01dbadmin01.us.oracle.com solaris excl
- osc01dbadmin01-01_T installed /zoneHomeT/osc01dbadmin01-01_T solaris excl
```

You do not have to create a template zone on the Database Domain in this case. Go to [“Creating VNETs” on page 131](#).

- If you do not see a file listed as *DB-domain\_T*, then you do not have the template file in place already. Go to [Step 3](#).

### 3. Locate the `ssc_exavm` script.

See [“Verify Configuration Tool Installation” on page 91](#) for the location of the `ssc_exavm` script.



#### 4. Create the template zone on the Database Domain where you are setting up zones.

The output shown below is abbreviated from the actual output.

```
./ssc_exavm -template
INFO: Logging all actions in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/osc01dbadmin01-01-20121129093
659.log and traces in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/osc01dbadmin01-01-20121129093
659.trc and o/p in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/osc01dbadmin01-01-20121129093
659.out
INFO: Begin Template Creation
INFO: Checking for osc01dbadmin01-01_T
INFO: Checking for osc01dbadmin01-01_T
INFO: Creating a template Virtual Guest for future use
INFO: 2012-11-29 09:37:14 : Executing zfs create rpool/osc01dbadmin01-01_T

INFO: Executing
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/vmtrash/8461.SysCall.0.cmd

INFO: Running System command ...
R R
SUCCESS: Done
INFO: 2 : Completed Command Execution
INFO: 2012-11-29 09:37:16 : Completed Command Execution
INFO: 2012-11-29 09:37:16 : Executing zfs set quota=6G rpool/osc01dbadmin01-01_T

INFO: Executing
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/vmtrash/8461.SysCall.1.cmd

... (OUTPUT ABBREVIATED) ...

INFO: Running System command ...
R
SUCCESS: Done
INFO: 1 : Completed Command Execution
INFO: 2012-11-29 09:51:25 : Completed Command Execution
SUCCESS: Completed creation of Virtual guests

INFO: Virtual Guest creation Summary :

INFO: Completed : osc01dbadmin01-01_T Data file :: none
```

The name for the template zone for this particular domain is *DB-domain\_T* (for example, *osc01dbadmin01-01\_T*).

5. Repeat [Step 1](#) through [Step 4](#) for every Database Domain where you are setting up zones.

You must create a separate template zone on each Database Domain where you are setting up zones. Each template zone has its own domain-specific name.

6. Determine if you need to create additional VNETs for certain Database Domains.

Go to [“Creating VNETs”](#) on page 131.

---

## ▼ Delete a Template Zone From a Database Domain

Follow these procedures if you want to delete the template zone from a Database Domain for any reason.

1. Locate the `ssc_exavm` script.

See [“Verify Configuration Tool Installation”](#) on page 91 for the location of the `ssc_exavm` script.

2. Log in to the Database Domain that has the template zone that you want to delete.

3. Locate the template zone on this Database Domain that you want to delete.

```
# zoneadm list -cv
ID NAME                STATUS    PATH                                BRAND  IP
0  global                running   /                                    solaris shared
-  osc01dbadmin01-01_T  installed /zoneHomeT/osc01dbadmin01-01_T  solaris excl
```

The template zone is shown as *DB-domain\_T* (for example, in the example output above, `osc01dbadmin01-01_T`).

4. Destroy the template zone on this Database Domain.

The output shown is abbreviated from the actual output.





# Creating VNETs

---

The following procedures describe how to create or delete VNETs.

- [“Determining if Additional VNETs Are Needed for a Database Domain”](#) on page 131
- [“Locate the Control Domain”](#) on page 133
- [“Create VNETs for Zones”](#) on page 134
- [“Delete VNETs”](#) on page 136

---

## Determining if Additional VNETs Are Needed for a Database Domain

You must create additional VNETs in certain specific situations, where you are setting up zones on certain Database Domains that do not have physical 1GbE host management ports associated with those domains, but rather use virtual network devices for the host management. You can use either of the following methods to determine if you need to create additional VNETs for the Database Domain where you will be setting up zones:

- [“Determine if Additional VNETs Are Needed \(Using Specific LDom Information\)”](#) on page 132
- [“Determine if Additional VNETs Are Needed \(Using Software Commands\)”](#) on page 132

## ▼ Determine if Additional VNETs Are Needed (Using Specific LDom Information)

You must create additional VNETs if you are setting up zones on any Database Domain that is not the first LDom (LDom 1) in your Oracle SuperCluster M6-32. See [“Understanding LDom” on page 20](#) to determine if the LDom that you are setting up zones on is LDom 1 or not.

For example, looking at the information in [“Understanding LDom” on page 20](#), the following would apply:

- If you are setting up zones on an LDom in the E4-1 PDomain configuration, you would not have to create additional VNETs for that LDom because there is only one LDom in that domain, so it is LDom 1 by default.
- If you are setting up zones on an LDom in the E4-2 PDomain configuration, you would not have to create additional VNETs for the first LDom (LDom 1), but you would have to create additional VNETs for the second LDom (LDom 2).
- **Using the information provided above, determine if you need to create additional VNETs.**
  - If none of these situations apply for you, you do not have to create additional VNETs. Go to [“Running the Java OneCommand” on page 139](#).
  - If you are setting up zones on Database Domains that are listed above, then you must create additional VNETs for the zones that you will be setting up on these Database Domains. Go to [“Create VNETs for Zones” on page 134](#).

## ▼ Determine if Additional VNETs Are Needed (Using Software Commands)

Another way to determine whether or not you need to create additional VNETs for certain Database Domains is by using these software commands.

1. **Log in to the Database Domain where you are setting up zones.**
2. **Type.**

```
# ipmpstat -g
```

You might see output similar to the following:

GROUP	GROUPNAME	STATE	FDT	INTERFACES
bondmgt0	bondmgt0	ok	--	net6 net2

3. Note the information displayed in the Interfaces column from the output.

In the example above, net6 and net2 are the interfaces displayed in the column.

4. Type.

```
# dladm show-phys | grep vnet
```

- If you do not see any output (if no output was returned after entering the command above), then you do not have to create additional VNETs. Go to [“Running the Java OneCommand” on page 139](#).
- If you see output similar to the following, then you must create additional VNETs for the zones that you are setting up on these Database Domains.

net2	Ethernet	up	0	unknown	vnet0
net6	Ethernet	up	0	unknown	vnet1

where the net# interfaces displayed in the `ipmpstat -g` command is shown with VNETs associated with those net# interfaces (in the above example, with vnet0 associated with the net2 interface and vnet1 associated with the net6 interface). Go to [“Create VNETs for Zones” on page 134](#).

---

## ▼ Locate the Control Domain

The tool for creating VNETs must be run on the control domain.

1. Log in to the first domain in your system.

2. Type.

```
# virtinfo -a
```

Output similar to the following appears:

```
Domain role: LDoms control I/O service root
Domain name: primary
Domain UUID: 96ef6675-114f-e6b1-d29b-90dd48693efb
Control domain: etc25dbadm01
Chassis serial#: AK00083251
```

3. Locate the line beginning with `Control domain` in the output, highlighted in the example output above.

The domain shown in this line is your control domain. If you are logged in to this particular domain when you enter the `virtinfo -a` command, you also see the value `LDoms control I/O service root` in the `Domain role` line for this domain.

---

## ▼ Create VNETs for Zones

1. Locate the `addvnet-wrapper.sh` script.

See “[Preparing to Configure Zones on Database Domains](#)” on page 83 for the location of the `addvnet-wrapper.sh` script.

2. On the first PDomain, list the Database Domains to locate the Database Domain that you are adding VNETs to.

```
ldm list
```

3. Log in to the control domain.

See “[Locate the Control Domain](#)” on page 133 for instructions on locating the control domain.

4. From the control domain, add VNETs to the guest domain on this PDomain.

```
addvnet-wrapper.sh DB-domain num-zone
```

where:

- *DB-domain* is the name of the guest domain.
- *num-zone* is the number of zones that you are setting up on the Database Domains in this particular PDomain.

For example, assume you have the following setup:

- The name of the guest domain is `osc01dbadmin02-01`.
- You are setting up four zones on the Database Domains on this PDomain.

Then you would type the following command:

```
addvnet-wrapper.sh osc01dbadmin02-01 4
Adding VNETs on osc01dbadmin02-01 ----| Okay
Info: found active ldom configuration F8_1_1_ML11242013214449
Info: will add new ldom configuration F8_1_1_ML11242013214449
```



```
Adding ldom configuration -----/ Okay
```

```
INFO: ssh into osc01dbadmin02-01 for post-vnet operations
INFO: Renamed link net24 to zonenet24
INFO: Renamed link net25 to zonenet25
INFO: Renamed link net26 to zonenet26
INFO: Renamed link net27 to zonenet27
```

- Viewed from the Hypervisor perspective, using the `ldm` command, the VNETs that you added start with `z`, followed by a nonnegative integer, followed by `mgnet`, and ending with either a `0` or a `1` (for example, `z0mgnet1`).
- Viewed from the Oracle Solaris perspective, using the `dladm` command, the network created is in the form of `zonenetxx`, where `xx` are consecutive numbers (such as `zonenet12`, `zonenet13`, and so on). Once a zone is deployed, the network picked for that particular zone or instance is mapped to `zonzonetxx` (such as `zonzonet12`, `zonzonet13`, and so on) to indicate that they are locked and used for the respective zone.

If you destroy a zone using the instructions in “[Delete a Template Zone From a Database Domain](#)” on page 128, the net pool is freed and renamed back to `zonenetxx`, so that any future addition of zones using zone tools can pick those VNETs.

##### 5. Check the guest domains to ensure the VNETs are created on the specified domain.

```
ldm ls -l -p DB-domain | grep '^VNET'
```

where *DB-domain* is the name of the Database Domain with zones that you added the VNETs to.

For example, if the name of the Database Domain with zones that you added the VNETs to is `osc01dbadmin02-01`, type the following:

```
ldm ls -l -p osc01dbadmin02-01 | grep '^VNET'
VNET | name=mngnet0 | dev=network@0 | service=mng-primary-net1@primary | mac-addr=
00:14:4f:fa:f8:e8 | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=0
VNET | name=mgnet1 | dev=network@1 | service=mng-primary-net00@primary | mac-addr=
00:14:4f:fb:a0:cd | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=1
VNET | name=z0mgnet0 | dev=network@2 | service=mng-primary-net1@primary | mac-addr=
00:14:4f:fb:da:06 | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=2
VNET | name=z0mgnet1 | dev=network@3 | service=mng-primary-net1@primary | mac-addr=
00:14:4f:f8:a8:06 | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=3
VNET | name=z1mgnet0 | dev=network@4 | service=mng-primary-net1@primary | mac-addr=
00:14:4f:fb:83:cc | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=4
VNET | name=z1mgnet1 | dev=network@5 | service=mng-primary-net1@primary | mac-addr=
00:14:4f:fb:26:ab | mode= | pvid=1 | vid= | mtu=1500 | linkprop= | id=5
```

6. Repeat [Step 2](#) through [Step 5](#) to add VNETs to the control domain on each PDomain.

7. Run the Java OneCommand.

Go to [“Running the Java OneCommand”](#) on page 139.

---

## ▼ Delete VNETs

Follow these procedures if you need to delete any VNETs that you added using the `advnet-wrapper.sh` script, and you created VNETs accidentally that are not needed.

---

**Note** – Do not delete VNETs that are used by zones. Doing so results in a failed zone.

---

1. Locate the `rmvnet.sh` script.

See [“Preparing to Configure Zones on Database Domains”](#) on page 83 for the location of the `rmvnet.sh` script.

2. Identify the VNETs that you want to remove.

```
ldm ls -l -p DB-domain | grep '^VNET'
```

where *DB-domain* is the name of the Database Domain that you added the VNETs to.

For example, if the name of the Database Domain that you added the VNETs to is `osc01dbadmin02-01`, you would type the following:

```
ldm ls -l -p osc01dbadmin02-01 | grep '^VNET'
VNET|name=mgnet0|dev=network@0|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fa:f8:e8|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=0
VNET|name=mgnet1|dev=network@1|service=mng-primary-net00@primary|mac-addr=
00:14:4f:fb:a0:cd|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=1
VNET|name=z0mgnet0|dev=network@2|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:da:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=2
VNET|name=z0mgnet1|dev=network@3|service=mng-primary-net1@primary|mac-addr=
00:14:4f:f8:a8:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=3
VNET|name=z1mgnet0|dev=network@4|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:83:cc|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=4
VNET|name=z1mgnet1|dev=network@5|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:26:ab|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=5
```

### 3. Remove the VNETs.

- If you want to remove specific VNETs, type.

```
rmvnet.sh -n vnet-name DB-domain
```

where:

- *vnet-name* is the name of the VNET that you want to remove. If you want to remove more than one VNET, type the name of each VNET that you want to remove, separated by commas.
- *DB-domain* is the name of the Database Domain from which you are removing the VNETs.

For example, assume you have the following setup:

- The names of the VNETs that you want to remove are `z1mgnet0` and `z1mgnet1`.
- The name of the Database Domain that you are removing the VNETs from is `osc01dbadmin02-01`.

Then you would type the following:

```
rmvnet.sh -n z1mgnet0,z1mgnet1 osc01dbadmin02-01
```

- If you want to remove several VNETs that have common characteristics, type.

```
rmvnet.sh -p 'beginning-term+' DB-domain
```

where:

- *beginning-term* is the term at the beginning of the VNET names for the VNETs that you want to remove (for example, all VNETs with names that begin with `z1`).
- *domain-name* is the name of the Database Domain from which you are removing the VNETs.

For example, assume you have the following setup:

- You want to remove all VNETs with names that begin with `z1` (in the output from [Step 2](#), that would be `z1mgnet0` and `z1mgnet1`).
- The name of the Database Domain that you are removing the VNETs from is `osc01dbadmin02-01`.

Then you would type the following:

```
rmvnet.sh -p 'z1+' osc01dbadmin02-01
```

#### 4. Verify that the VNETs have been removed.

```
ldm ls -l -p DB-domain | grep '^VNET'
```

where *DB-domain* is the name of the Database Domain that you removed the VNETs from.

For example, if the name of the Database Domain that you removed the VNETs from is `osc01dbadmin02-01`, type the following:

```
ldm ls -l -p osc01dbadmin02-01 | grep '^VNET'  
VNET|name=mgnet0|dev=network@0|service=mng-primary-net1@primary|mac-addr=  
00:14:4f:fa:f8:e8|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=0  
VNET|name=mgnet1|dev=network@1|service=mng-primary-net00@primary|mac-addr=  
00:14:4f:fb:a0:cd|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=1  
VNET|name=z0mgnet0|dev=network@2|service=mng-primary-net1@primary|mac-addr=  
00:14:4f:fb:da:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=2  
VNET|name=z0mgnet1|dev=network@3|service=mng-primary-net1@primary|mac-addr=  
00:14:4f:f8:a8:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=3
```

The VNETs that you removed should not appear in the list.

# Running the Java OneCommand

---

These topics describe how to run the Java OneCommand.

- “Set Up Public Key Authentication for ZFS Storage Controllers” on page 139
- “Set Up the Java OneCommand” on page 141
- “Run the Java OneCommand” on page 142
- “NTP Issue” on page 144

---

## ▼ Set Up Public Key Authentication for ZFS Storage Controllers

Before running the Java OneCommand, you must set up public key authentication on the ZFS storage controllers for every LDom where you are setting up zones on the Database Domains.

### 1. Log in to the first ZFS storage controller using the CLI.

To log in remotely using the CLI, use an `ssh` client. If you have not configured other users to administer the appliance, you must log in as `root`. When you log in, the CLI prompt consists of the hostname, followed by a colon, followed by a greater-than sign:

```
% ssh root@zfs-controller1-hostname  
Password:  
Last login: Mon Oct 13 15:43:05 2009 from kiowa.sf.fishpo  
zfs-controller-hostname:>
```

where *zfs-controller1-hostname* is the host name for the first ZFS storage controller.

2. Set up the public key authentication on the first ZFS storage controller for the first Database Domain that contains zones.

```
zfs-controller1-hostname:> configuration preferences keys
zfs-controller1-hostname:configuration preferences keys> create
zfs-controller1-hostname:configuration preferences key (uncommitted)> set type=DSA
zfs-controller1-hostname:configuration preferences key (uncommitted)> set key=
"DSA-key-text"
                                key = DSA-key-text == (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> set comment=
"fw-log1"
                                comment = fw-log1 (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> commit
zfs-controller1-hostname:configuration preferences keys> show
Keys:
NAME          MODIFIED          TYPE    COMMENT
key-000      10/12/2009 10:54:58    DSA     fw-log1
```

The *DSA-key-text* is the key text itself (usually hundreds of characters), without spaces.

3. Repeat [Step 2](#) for every Database Domain to contain zones.
4. Log out of the first ZFS storage controller.

```
zfs-controller1-hostname:> exit
```

---

**Note** – You do not have to repeat this procedure for the second ZFS storage controller. When you set up the public key authentication on the first ZFS storage controller, it appears when you log in to the second ZFS storage controller.

---

5. Log in to the first Database Domain in your system, if you are not logged in already.
6. Verify that you set up the public key authentication correctly.

```
% ssh root@zfs-controller1-hostname hostname
```

where *zfs-controller1-hostname* is the host name for the first ZFS storage controller. You should be able to log in to the ZFS storage controller successfully and see the CLI prompt.

7. Repeat [Step 5](#) and [Step 6](#) for every Database Domain in your system.

---

## ▼ Set Up the Java OneCommand

1. **Ensure that the machine running the OEDA has Oracle JRE 1.6 or later.**

2. **Locate the `WorkDir` directory.**

By default, the `WorkDir` directory is located in the directory where you downloaded the Java OneCommand patch, as described in “[Preparing to Configure Zones on Database Domains](#)” on page 83.

3. **Copy all the necessary binaries in this `WorkDir` directory.**

Place binaries such as Oracle Database, Oracle Grid Infrastructure, and patches into this `WorkDir` directory. Refer to the following My Oracle Support note for more information on the necessary binaries:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=888828.1>

Following are examples of the binaries and patches to place in this `WorkDir` directory:

- `p13390677_112040_SOLARIS64_1of7.zip`
- `p13390677_112040_SOLARIS64_2of7.zip`
- `p13390677_112040_SOLARIS64_3of7.zip`
- `p17628025_112040_SOLARIS64.zip`
- `p6880880_112000_SOLARIS64.zip`

4. **Locate the Java OneCommand file.**

By default, the Java OneCommand file (`install.sh`) is located in the directory where you downloaded the Java OneCommand patch, as described in “[Preparing to Configure Zones on Database Domains](#)” on page 83.

5. **Locate the `iscsi.zpool.sh` script.**

See “[Verify Configuration Tool Installation](#)” on page 91 for the location of that script.

6. **Copy the `iscsi.zpool.sh` script into the following directory on the Database Domain where you are running the Java OneCommand:**

`/opt/oracle.supercluster/bin`

7. **Run the Java OneCommand.**

Go to “[Run the Java OneCommand](#)” on page 142.

---

## ▼ Run the Java OneCommand



---

**Caution** – Ensure that you backup all existing databases before running the Java OneCommand.

---



---

**Caution** – You must use the latest version of OEDA and the Java OneCommand patch (patch 19766190 or later). Refer to the Oracle Exadata Deployment Assistant (OEDA) section in MOS Note 888828.1 for details.

---



---

**Caution** – Any version of Java OneCommand prior to patch 19766190 can destroy the storage cell disks and griddisk if an undo option is performed on certain steps (Create Cell disk, for example). This can cause complete destruction of all the griddisks on the storage cells. In addition, re-running the griddisk creation step or mistakenly specifying a non-unique diskgroup in OEDA will result in the destruction of existing griddisks. Be aware that older versions of Java OneCommand also destroy cell disks and griddisks with the Create Cell Disks step.

---

1. **Locate the configuration file that was generated at the end of the OEDA process when you added zones to the Database Domains.**

The configuration file should be named *customer-name.name-prefix.xml*, such as `Acme.osc01.xml`.



## 2. Display the list of functions to be performed by the Java OneCommand.

```
./install.sh -cf customer-name.name-prefix.xml -l
```

For example:

```
./install.sh -cf Acme.osc01.xml -l
```

Information similar to the following is displayed:

```
1. Validate Configuration File
2. Create Virtual Machine
3. Setup Required Files
4. Create Users
5. Setup Cell Connectivity
6. Verify Infiniband and Calibrate Cells
7. Create Cell Disks
8. Create Grid Disks
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Create Installation Summary
17. Resecure Machine
```

The output that you see and the number assigned to each step in the output varies depending on the configuration choices you made in OEDA.

## 3. Note the specific steps that you will not run in Java OneCommand.

In the next step, you will run the appropriate steps in Java OneCommand. **Do not run any of the following steps in Java OneCommand in the next step:**

- Verify InfiniBand and calibrate cells (step 6 in the example output above)
- Create cell disks (step 7 in the example output above)
- Configure cell alerting (not shown in the example output above)

4. Run the appropriate steps in the Java OneCommand, using the configuration file that was generated at the end of the OEDA process as input.

```
./install.sh -cf customer-name.name-prefix.xml -s step-number
```

where *step-number* is the number of each step that you are running, in order. For example:

```
./install.sh -cf Acme.osc01.xml -s 1
```

Run all of the steps in the Java OneCommand **except for the following**:

- Verify InfiniBand and calibrate cells
- Create cell disks
- Configure cell alerting

You can also run a series of steps by using the `-r` option. For example, if you wanted to run steps 1 through 5, enter the following:

```
./install.sh -cf Acme.osc01.xml -r 1-5
```

You can destroy a zone from a Database Domain using the instructions in [“Delete a Template Zone From a Database Domain” on page 128](#) if necessary (for example, if Java OneCommand fails at step 2 for some reason). However, use caution when destroying zones from a Database Domain.

5. **Create additional links on the IB storage network for the zones that you set up on the Database Domains.**

Go to [“Creating Additional Links on the IB Storage Network for Zones” on page 147](#).

---

## NTP Issue

When your Oracle SuperCluster M6-32 was first installed, the NTP service was configured and enabled in the global zone. Only one NTP daemon can run on one system. So if you have a system with a combination of global and nonglobal zones, only one NTP can run across those global and nonglobal zones.

When you create zones on Database Domains, the NTP service within the zones on the Database Domains is disabled automatically by the zone configuration scripts, and they must be left disabled. The zone configuration scripts also ensure that the Cluster Time Synchronization Services daemon (`ctssd`) only runs in observer mode once the database software is installed and configured.

---

**Note** – Do not remove the `/etc/inet/ntp.conf` file from any zones on Database Domains. Doing so causes `ctssd` to try to go into active mode, which results in a failure to start the Grid Infrastructure in the zones on Database Domains.

---



# Creating Additional Links on the IB Storage Network for Zones

You must create additional links on the IB storage network for the zones that you set up on the Database Domains. This topic provides those instructions:

- [“Create Additional Links on the IB Storage Network for Zones” on page 147](#)

## ▼ Create Additional Links on the IB Storage Network for Zones

1. Log in to the Database Domain that contains the zones that you created.
2. Type:

```
# dladm show-ib
```

Output similar to the following appears:

LINK	HCAGUID	PORTGUID	PORT	STATE	PKEYS
net24	21280001FCBEBE	21280001FCBEBE	2	up	FFFF
<b>net22</b>	<b>21280001FCC0EC</b>	<b>21280001FCC0EE</b>	<b>2</b>	<b>up</b>	<b>8503,8504,FFFF</b>
<b>net23</b>	<b>21280001FCBEBE</b>	<b>21280001FCBEBD</b>	<b>1</b>	<b>up</b>	<b>8503,8504,FFFF</b>
net21	21280001FCC0EC	21280001FCC0ED	1	up	FFFF

3. Determine which network links are associated with the 8503 PKEY.

In the example output shown in the previous step, the lines that show the network links associated with the 8503 PKEY are shown in **bold**. The network links are shown in the LINK column (net22 and net23 in the example output).

#### 4. Add links to the zones in this Database Domains:

```
# zonecfg -z zone_management_hostname
zonecfg:zone_management_hostname> add anet
zonecfg:zone_management_hostname:anet> set linkname=first_net_link
zonecfg:zone_management_hostname:anet> set lower-link=first_net_link
zonecfg:zone_management_hostname:anet> set pkey=8503
zonecfg:zone_management_hostname:anet> end
zonecfg:zone_management_hostname> add anet
zonecfg:zone_management_hostname:anet> set linkname=second_net_link
zonecfg:zone_management_hostname:anet> set lower-link=second_net_link
zonecfg:zone_management_hostname:anet> set pkey=8503
zonecfg:zone_management_hostname:anet> end
zonecfg:zone_management_hostname> exit
```

where:

- *zone\_management\_hostname* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone
- *first\_net\_link* is the first network link shown in the output from [Step 2](#)
- *second\_net\_link* is the second network link shown in the output from [Step 2](#)

For example:

```
# zonecfg -z zone_management_hostname
zonecfg:zone_management_hostname> add anet
zonecfg:zone_management_hostname:anet> set linkname=net22
zonecfg:zone_management_hostname:anet> set lower-link=net22
zonecfg:zone_management_hostname:anet> set pkey=8503
zonecfg:zone_management_hostname:anet> end
zonecfg:zone_management_hostname> add anet
zonecfg:zone_management_hostname:anet> set linkname=net23
zonecfg:zone_management_hostname:anet> set lower-link=net23
zonecfg:zone_management_hostname:anet> set pkey=8503
zonecfg:zone_management_hostname:anet> end
zonecfg:zone_management_hostname> exit
```

#### 5. Reboot the zone for the changes to take effect:

```
# zoneadm -z zone_management_hostname reboot
```

For example:

```
# zoneadm -z osc01zdbadm02 reboot
```

## 6. Log in to the zone:

```
# zlogin zone_management_hostname
```

For example:

```
# zlogin osc01zdbadm02
```

## 7. Create the IP addresses in the zone for the new links:

```
# ipadm create-ip first_net_link  
# ipadm create-ip second_net_link
```

where:

- *first\_net\_link* is the first network link shown in the output from [Step 2](#)
- *second\_net\_link* is the second network link shown in the output from [Step 2](#)

For example:

```
# ipadm create-ip net22  
# ipadm create-ip net23
```

## 8. Find an unused IP address from the same subnet as the `stor_ipmp0` interface:

a. Enter the following command:

```
# ipadm show-addr
```

Output similar to the following appears:

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net13/v4	static	ok	169.254.182.77/24
<b>stor_ipmp0/v4</b>	<b>static</b>	<b>ok</b>	<b>192.168.28.2/22</b>
bondeth0/v4	static	ok	10.129.119.1/20
bondib0/v4	static	ok	192.168.10.1/22
bondib1/v4	static	ok	192.168.10.2/22
bondmgt0/v4	static	ok	10.129.104.1/20
lo0/v6	static	ok	::1/128

b. Locate the `stor_ipmp0` line in the output, highlighted in the example output above.

In the ADDR column for this line, the subnet is shown as 192.168.28.2, with a subnet mask of 22.

- c. Find an unused IP address from this `stor_ipmp0` subnet to use for the new IPMP group.

9. Create the new IPMP group inside the zone:

```
# ipadm create-ipmp -i first_net_link,second_net_link ipmp_group
# ipadm create-addr -T static -a ip_address/22 ipmp_group/v4
```

where:

- *first\_net\_link* is the first network link shown in the output from [Step 2](#)
- *second\_net\_link* is the second network link shown in the output from [Step 2](#)
- *ipmp\_group* is the name of the IPMP group that you are creating
- *ip\_address* is an unused IP address from the same subnet as the `stor_ipmp0` interface that you picked in [Step 8](#)

For example:

```
# ipadm create-ipmp -i net22,net23 stor_ipmp0
# ipadm create-addr -T static -a 192.168.28.101/22 stor_ipmp0/v4
```

10. Designate the second network interface as the standby interface:

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

where *second\_net\_link* is the second network link shown in the output from [Step 2](#).

For example:

```
# ipadm set-ifprop -p standby=on -m ip net23
```

11. Verify that the second network interface was set up correctly as the standby interface:

```
# ipmpstat -g
```

The second network interface should appear in parenthesis in the output. For example:

```
# ipmpstat -g
GROUP      GROUPNAME    STATE  FDT  INTERFACES
stor_ipmp0 stor_ipmp0   ok     --   net22 (net23)
```

12. Repeat [Step 4](#) through [Step 11](#) for each zone that you created in this Database Domain.



# Setting Up VLAN Tagging for Zones

---

This topic provides instructions for setting up VLAN tagging for zones:

- [“Set Up VLAN Tagging for Zones” on page 151](#)

---

## ▼ Set Up VLAN Tagging for Zones

1. Log in to the nonglobal zone where you want to set up VLAN tagging.
2. Execute the following commands in the nonglobal zone to stop and disable the automatic startup of the Oracle Clusterware (CRS/HAS) software:

```
# crsctl stop crs -f
# crsctl disable crs
```

3. Log out of the nonglobal zone.
4. Log in to the Database Domain (the global zone) that contains the zones that you created.
5. Halt the zones that are in this Database Domain:

```
# zoneadm -z zone_management_hostname halt
```

where *zone\_management\_hostname* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone. For example:

```
# zoneadm -z osc01zdbadm02 halt
```

6. Identify the client network devices to be used for this VLAN:

a. Type:

```
# ipmpstat -g
```

Output similar to the following appears. Locate the line in the output for the bondeth0 interface, highlighted in the following example output:

GROUP	GROUPNAME	STATE	FDT	INTERFACES
bondmgt0	bondmgt0	ok	--	net0 (net1)
bondib1	bondib1	ok	--	bondib1_0 (bondib1_1)
bondib0	bondib0	ok	--	bondib0_0 (bondib0_1)
<b>bondeth0</b>	<b>bondeth0</b>	<b>ok</b>	<b>--</b>	<b>net4 (net7)</b>
stor_ipmp0	stor_ipmp0	ok	--	stor_ipmp0_0 (stor_ipmp0_1)

b. Note the entries for this line under the INTERFACES column.

In the example output above, those entries would be net4 and net7.

c. Type:

```
# dladm show-phys | grep ixgbe
```

Output similar to the following appears. Locate the lines in the output for the two entries that you noted in [Step b](#), highlighted in the following example output:

net5	Ethernet	unknown	0	unknown	ixgbe3
net1	Ethernet	up	1000	full	ixgbe1
net6	Ethernet	unknown	0	unknown	ixgbe4
net0	Ethernet	up	1000	full	ixgbe0
<b>net4</b>	<b>Ethernet</b>	<b>up</b>	<b>10000</b>	<b>full</b>	<b>ixgbe2</b>
<b>net7</b>	<b>Ethernet</b>	<b>up</b>	<b>10000</b>	<b>full</b>	<b>ixgbe5</b>

## 7. Modify the zone configuration:

```
# zonecfg -z zone_management_hostname
zonecfg: zone_management_hostname> select anet linkname=net_device_1
zonecfg: zone_management_hostname:anet> set vlan-id=VLAN_ID
zonecfg: zone_management_hostname:anet> set lower-link=net_device_1
zonecfg: zone_management_hostname:anet> end
zonecfg: zone_management_hostname> select anet linkname=net_device_2
zonecfg: zone_management_hostname:anet> set vlan-id=VLAN_ID
zonecfg: zone_management_hostname:anet> set lower-link=net_device_2
zonecfg: zone_management_hostname:anet> end
zonecfg: zone_management_hostname> verify
zonecfg: zone_management_hostname> commit
zonecfg: zone_management_hostname> exit
```

where:

- *zone\_management\_hostname* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone
- *VLAN\_ID* is a VLAN ID that you provide
- *net\_device\_1* is the first client network device that you saw in the output in [Step b](#)
- *net\_device\_2* is the second client network device that you saw in the output in [Step b](#)

For example:

```
# zonecfg -z osc01zdbadm02
zonecfg: osc01zdbadm02> select anet linkname=net4
zonecfg: osc01zdbadm02:anet> set vlan-id=13
zonecfg: osc01zdbadm02:anet> set lower-link=net4
zonecfg: osc01zdbadm02:anet> end
zonecfg: osc01zdbadm02> select anet linkname=net7
zonecfg: osc01zdbadm02:anet> set vlan-id=13
zonecfg: osc01zdbadm02:anet> set lower-link=net7
zonecfg: osc01zdbadm02:anet> end
zonecfg: osc01zdbadm02> verify
zonecfg: osc01zdbadm02> commit
zonecfg: osc01zdbadm02> exit
```

8. Reboot the zone for the changes to take effect:

```
# zoneadm -z zone_management_hostname reboot
```

For example:

```
# zoneadm -z osc01zdbadm02 reboot
```

9. Ping the default gateway to verify that the client network is up.
10. Log out of the Database Domain (the global zone).
11. Log in to the nonglobal zone where you set up VLAN tagging.
12. Execute the following commands in the nonglobal zone to enable and start the automatic startup of the Oracle Clusterware (CRS/HAS) software:

```
# crsctl enable crs  
# crsctl start crs
```

# Index

---

## A

addvnet.sh script  
    creating VNETs, 132

## B

B2-1 LDom configuration, 48  
B2-2 LDom configuration, 49  
B2-3 LDom configuration, 50  
B2-4 LDom configuration, 51  
B4-1 LDom configuration, 41  
B4-2 LDom configuration, 42  
B4-3 LDom configuration, 43  
B4-4 LDom configuration, 44  
bounded PDomains, description, 4

## C

configurations  
    LDoms, 20  
    PDomains, 3

## D

determining  
    amount of space for each disk group, 66  
    available storage capacity, 63  
    maximum amount of available space for  
        RACs, 65  
    number of RACs, 65  
disk groups  
    considerations when planning layout, 61  
    determining amount of disk space, 66  
    planning layout, 62  
    recording planned layout, 80

## E

E2-1 LDom configuration, 31  
E2-2 LDom configuration, 33

E2-3 LDom configuration, 34

E2-4 LDom configuration, 36

E4-1 LDom configuration, 23

E4-2 LDom configuration, 24

E4-3 LDom configuration, 25

E4-4 LDom configuration, 27

Exadata Storage Servers

    considerations when planning disk group  
        layout, 61

    example layout scenario, 66

    planning disk group layout, 62

    recording planned disk group layout, 80

    rules and restrictions for disk group layout, 62

## F

fully-populated DCUs

    LDom configurations for, 22

## I

iSCSI LUNs

    locating iscsi-lun.sh script, 93

iscsi-lun.sh script

    locating, 93

## L

LDom configurations

    B2-1, 48

    B2-2, 49

    B2-3, 50

    B2-4, 51

    B4-1, 41

    B4-2, 42

    B4-3, 43

    B4-4, 44

    base PDomains

        fully-populated DCUs, 40

        half-populated DCUs, 47

- E2-1, 31
- E2-2, 33
- E2-3, 34
- E2-4, 36
- E4-1, 23
- E4-2, 24
- E4-3, 25
- E4-4, 27
- extended PDomains
  - fully-populated DCUs, 22
  - half-populated DCUs, 30

## N

- nonbounded PDomains, description, 4

## O

- Oracle Exadata Deployment Assistant
  - creating the configuration files, 100, 141
  - downloading, 91
- Oracle OneCommand utility
  - downloading, 91
- overview
  - base configuration PDomains, 38
  - extended configuration PDomains, 21
  - PDomains, 3
    - base configurations, 11
    - extended configurations, 5
  - SPs, 3

## P

- PDomains
  - base configurations
    - overview, 11, 38
    - R3, 13
      - R3\_1, 14
      - R3\_2, 14
      - R3\_3, 15
    - R4, 15
      - R4\_1, 16
      - R4\_2, 16
      - R4\_3, 17
      - R4\_4, 18
    - R5, 18
      - R5\_1, 19
    - R6, 19
      - R6\_1, 20
  - bounded PDomains, description, 4
  - configuration guidelines, 4

- extended configurations
  - overview, 5, 21
    - R1, 6
      - R1\_1, 7
      - R1\_2, 7
      - R1\_3, 8
    - R2, 8
      - R2\_1, 9
      - R2\_2, 10
      - R2\_3, 11
  - nonbounded PDomains, description, 4
  - overview, 3
  - preparing to set up zones, 83
  - process for setting up zones on Database Domains
    - tasks, 83, 95

## R

- R1 extended configuration PDomains, 6
- R1\_1 extended configuration PDomain, 7
- R1\_2 extended configuration PDomain, 7
- R1\_3 extended configuration PDomain, 8
- R2 extended configuration PDomain, 8
- R2\_1 extended configuration PDomain, 9
- R2\_2 extended configuration PDomain, 10
- R2\_3 extended configuration PDomain, 11
- R3 base configuration PDomain, 13
- R3\_1 base configuration PDomain, 14
- R3\_2 base configuration PDomain, 14
- R3\_3 base configuration PDomain, 15
- R4 base configuration PDomain, 15
- R4\_1 base configuration PDomain, 16
- R4\_2 base configuration PDomain, 16
- R4\_3 base configuration PDomain, 17
- R4\_4 base configuration PDomain, 18
- R5 base configuration PDomain, 18
- R5\_1 base configuration PDomain, 19
- R6 base configuration PDomain, 19
- R6\_1 base configuration PDomain, 20
- RACs
  - determining maximum allowable disk space, 65
  - determining number of RACs, 65
  - planning number of zones and RACs, 56
  - recording planned configuration, 79
  - rules and recommendations, 60
- recording

- planned Exadata Storage Server disk group
  - layout, 80
- planned RAC configuration, 79
- planned zone configuration, 68
- repository
  - determining location of, 85
  - installing or updating files from local repository, 88
  - installing or updating files from remote repository, 86
- restrictions
  - Exadata Storage Server disk group layout, 62
  - setting up zones and RACs, 60
- rmvnet.sh script
  - deleting VNETs, 136
  - locating, 93
- rules
  - Exadata Storage Server disk group layout, 62
  - setting up zones and RACs, 60

## **S**

- SPs
  - overview, 3
- ssc\_exavm tool
  - creating template zone, 126
  - deleting template zone, 128

## **T**

- template zone
  - creating, 126
  - deleting, 128
  - understanding, 125

## **V**

- VNETs
  - creating, 132
  - deleting, 136
  - locating rmvnet.sh script, 93

## **Z**

- zonemanifest.sh script, locating, 93
- zones
  - planning number of zones and RACs, 56
  - recording planned zone configuration, 68
  - rules and recommendations, 60

