

# Oracle® SuperCluster T5-8 Zones With Oracle Database on Database Domains Configuration Guide

ORACLE®

Part No: E41620-09  
April 2020



**Part No: E41620-09**

Copyright © 2013, 2020, 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 documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

**Access to Oracle Support**

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

**Référence: E41620-09**

Copyright © 2013, 2020, 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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

**Accès aux services de support Oracle**

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

# Contents

---

- Using This Documentation** ..... 9
  - Product Documentation Library ..... 9
  - Feedback ..... 9
  
- Planning to Set Up Zones on Database Domains** ..... 11
  - Important Cautions ..... 12
  - Naming Conventions ..... 12
  - Understanding Domain Configurations ..... 12
    - Half Rack Domain Configurations ..... 13
    - Full Rack Domain Configurations ..... 14
  - Determining the Cores Available for Domains and Zones ..... 16
    - Cores Available for Domains and Zones ..... 16
    - Example Configuration ..... 17
  - Memory Available for Domains ..... 19
  - Guidelines for Planning Zones ..... 19
  - Guidelines for Planning Zone Clusters ..... 20
  - Guidelines for Planning the Storage Server Disk Group Layout ..... 21
    - Planning the Exadata Storage Server Disk Group Layout ..... 22
    - Understanding an Example Scenario ..... 27
  - Recording Your Existing and Planned Configuration Information ..... 29
    - Recording Your Existing Configuration ..... 29
    - Recording Your Zone Configuration Information ..... 32
    - Recording Your Cluster Configuration Information ..... 36
    - Recording Your Exadata Storage Server Disk Group Layout ..... 38
  
- Understanding Network Considerations** ..... 41
  - Understanding the Networking Setup Before Zone Creation ..... 41
    - Understanding the Network Configuration For the Half Rack ..... 42

Understanding the Network Configuration for the Full Rack .....	47
Understanding the Networking Setup After Zone Creation .....	54
Host Management Network After Zone Creation .....	55
Client Access Network After Zone Creation .....	56
IB Network After Zone Creation .....	56
Network Setup for Clusters .....	57
<b>Preparing to Configure Zones on Database Domains .....</b>	<b>59</b>
▼ Prepare to Configure Zones .....	59
Existing Network Worksheet .....	61
▼ Update the Base Software .....	62
▼ Determine the Repository Location .....	63
▼ Install or Update Packages From the Remote Repository .....	64
▼ Install or Update Files From the Local Repository .....	65
▼ Verify Configuration Tool Installation .....	70
<b>Creating Configuration Files Using OEDA .....</b>	<b>73</b>
Example Configuration .....	73
▼ Verify Exadata Storage Server Disk Space for Additional Zones .....	74
▼ Determine Software Build Level .....	75
Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier .....	76
Oracle Exadata Deployment Assistant Overview .....	78
▼ Set Up the OEDA .....	79
▼ Start the OEDA .....	79
▼ Complete the Customer Details Page .....	81
▼ Complete the Hardware Selection Page .....	83
▼ Complete the Define Customer Networks Page .....	87
▼ Complete the Administration Network Page .....	89
▼ Complete the Client Ethernet Network Page .....	95
▼ Complete the IB Network Page .....	98
▼ Complete the Identify Compute Node OS and Enable Capacity-on-Demand Configuration Page .....	101
▼ Review the Information in the Management and Private Networks Page .....	102
▼ Complete the Define Clusters Page .....	104
▼ Set Zone Default Configurations .....	108
▼ Complete the Cluster Page .....	111

▼ Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page .....	115
▼ Verify Remaining Configuration Information .....	117
▼ Generate the Configuration Files .....	117
Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later .....	120
▼ Locate the Necessary Files to Set Up Zones .....	121
▼ Set Up the OEDA .....	122
▼ Start the OEDA .....	123
▼ Import the Most Recent OEDA Configuration File .....	124
▼ Review Existing Configuration Information .....	126
▼ Review the Information in the Identify Compute Node Operating System Page .....	127
▼ Review the Information in the Management and Private Networks Page ....	129
▼ Complete the Define Clusters Page .....	131
▼ Set Zone Default Configurations .....	137
▼ Complete the Cluster Page .....	140
▼ Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page .....	144
▼ Verify Remaining Configuration Information .....	146
▼ Generate the Configuration Files .....	146
<b>Creating the Template Zone on Each Database Domain .....</b>	<b>149</b>
▼ Create a Template Zone on a Database Domain .....	149
▼ Delete a Template Zone From a Database Domain .....	151
<b>Creating VNETs .....</b>	<b>153</b>
Determining if Additional VNETs Are Needed for a Database Domain .....	153
▼ Determine if Additional VNETs Are Needed (CPU-to-Database Domain Mapping) .....	153
▼ Determine if Additional VNETs Are Needed (Software Commands) .....	154
▼ Locate the Control Domain .....	155
▼ Create VNETs for Zones .....	156
▼ Delete VNETs .....	158
<b>Running the Java OneCommand .....</b>	<b>161</b>
▼ Set Up Public Key Authentication for ZFS Storage Controllers .....	161
▼ Set Up the Java OneCommand .....	162

▼ Run the Java OneCommand .....	163
NTP Issue .....	166
<b>Creating Additional Links on the IB Storage Network for Zones .....</b>	<b>167</b>
▼ Create Additional Links on the IB Storage Network for Zones .....	167
<b>Setting Up VLAN Tagging for Zones .....</b>	<b>173</b>
▼ Set Up VLAN Tagging for Zones .....	173
<b>Index .....</b>	<b>177</b>

## Using This Documentation

---

- **Overview** – Describes how to set up zones with Oracle databases on Database Domains on Oracle SuperCluster T5-8.
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience installing and configuring computer hardware

## Product Documentation Library

Documentation and resources for this product and related products are available at [https://docs.oracle.com/cd/E40166\\_01/index.html](https://docs.oracle.com/cd/E40166_01/index.html).

## Feedback

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



# Planning to Set Up Zones on Database Domains

---

Once the Oracle SuperCluster T5-8 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://docs.oracle.com/cd/E23823\\_01/index.html](http://docs.oracle.com/cd/E23823_01/index.html)
- 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://docs.oracle.com/cd/E23824\\_01/](http://docs.oracle.com/cd/E23824_01/)

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

- “Important Cautions” on page 12
- “Naming Conventions” on page 12
- “Understanding Domain Configurations” on page 12
- “Determining the Cores Available for Domains and Zones” on page 16
- “Memory Available for Domains” on page 19
- “Guidelines for Planning Zones” on page 19
- “Guidelines for Planning Zone Clusters” on page 20
- “Guidelines for Planning the Storage Server Disk Group Layout” on page 21
- “Recording Your Existing and Planned Configuration Information” on page 29

## 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.

---



---

**Caution** - No more than eight zones with an Oracle database are supported on a Database Domain.

---

## Naming Conventions

This document uses these naming conventions:

- *sscnumber* for the Oracle SuperCluster T5-8
- *rnumber* for the cluster number
- *cnumber* for the SPARC T5-8 server (compute node) number
- *dbnumber* for the Database Domain number
- *znumber* for the zone number

## Understanding Domain Configurations

Several configurations are available for both the Half Rack version of the Oracle SuperCluster T5-8 (Half Rack) and the Full Rack version of the Oracle SuperCluster T5-8 (Full Rack):

- [“Half Rack Domain Configurations” on page 13](#)

- [“Full Rack Domain Configurations” on page 14](#)

## Half Rack Domain Configurations

Every Oracle SuperCluster T5-8 contains two SPARC T5-8 servers. For the Half Rack, each SPARC T5-8 server contains *two* processor modules.

The following domain configurations are available for each SPARC T5-8 server in the Half Rack:

- **Config H1-1:** One domain (one Large Domain)
- **Config H2-1:** Two domains (two Medium Domains)
- **Config H3-1:** Three domains (one Medium Domain, two Small Domains)
- **Config H4-1:** Four domains (four Small Domains)

The following figure provides more information on these configuration options.

		Processor Module 0				Processor Module 3			
		CPU 0		CPU 1		CPU 6		CPU 7	
		pci@300 pci@340		pci@380 pci@3c0		pci@600 pci@640		pci@680 pci@6c0	
Configurations	H1-1	Large Domain							
	H2-1	Medium Domain				Medium Domain			
	H3-1	Medium Domain				Small Domain	Small Domain		Small Domain
	H4-1	Small Domain	Small Domain		Small Domain	Small Domain		Small Domain	
Slot no.	1	3	9	11	6	8	14	16	
	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	

The number of InfiniBand host channel adaptors (IB HCAs) that are assigned to any domain differs, depending on the size of the domain. In addition, the specific IB HCAs that are assigned to a domain differs, depending on that domain's location in the SPARC T5-8 server.

For example, if the H3-1 domain configuration is set up on one of the SPARC T5-8 servers, these IB HCAs, in these PCIe slots, are assigned to each domain:

- **Medium Domain:** Two IB HCAs, in PCIe slots 3 and 11
- **First Small Domain:** One IB HCA, in PCIe slot 8
- **Second Small Domain:** One IB HCA, in PCIe slot 16

## Full Rack Domain Configurations

Every Oracle SuperCluster T5-8 contains two SPARC T5-8 servers. For the Full Rack, each SPARC T5-8 server contains *four* processor modules.

The following domain configurations are available for each SPARC T5-8 server in the Full Rack:

- **Config F1-1:** One domain (one Giant Domain)
- **Config F2-1:** Two domains (two Large Domains)
- **Config F3-1:** Three domains (one Large Domain, two Medium Domains)
- **Config F4-1:** Four domains (four Medium Domains)
- **Config F4-2:** Four domains (one Large Domain, two Small Domains, one Medium Domain)
- **Config F5-1:** Five domains (two Medium Domains, two Small Domains, one Medium Domain)
- **Config F5-2:** Five domains (one Large Domain, four Small Domains)
- **Config F6-1:** Six domains (one Medium Domain, four Small Domains, one Medium Domain)
- **Config F7-1:** Seven domains (one Medium Domain, six Small Domains)
- **Config F8-1:** Eight domains (eight Small Domains)

The following figure provides more information on these configuration options.

		Processor Module 0		Processor Module 1		Processor Module 2		Processor Module 3								
		CPU 0	CPU 1	CPU 2	CPU 3	CPU 4	CPU 5	CPU 6	CPU 7							
		pci@300 pci@340	pci@380 pci@3c0	pci@400 pci@440	pci@480 pci@4c0	pci@500 pci@540	pci@580 pci@5c0	pci@600 pci@640	pci@680 pci@6c0							
Configurations	F1-1	Giant Domain														
	F2-1	Large Domain				Large Domain										
	F3-1	Large Domain				Medium Domain		Medium Domain								
	F4-1	Medium Domain		Medium Domain		Medium Domain		Medium Domain								
	F4-2	Large Domain				Small Domain	Small Domain	Medium Domain								
	F5-1	Medium Domain		Medium Domain		Small Domain	Small Domain	Medium Domain								
	F5-2	Large Domain				Small Domain	Small Domain	Small Domain	Small Domain							
	F6-1	Medium Domain		Small Domain	Small Domain	Small Domain	Small Domain	Medium Domain								
	F7-1	Medium Domain		Small Domain												
	F8-1	Small Domain														
Slot no.	1	3	9	11	2	4	10	12	5	7	13	15	6	8	14	16
	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA	10 GbE NIC	IB HCA

The number of InfiniBand host channel adaptors (IB HCAs) that are assigned to any domain differs, depending on the size of the domain. In addition, the specific IB HCAs that are assigned to a domain differs, depending on that domain's location in the SPARC T5-8 server.

For example, if the F3-1 domain configuration is set up on one of the SPARC T5-8 servers, these IB HCAs, in these PCIe slots, are assigned to each domain:

- **Large Domain:** Four IB HCAs, in PCIe slots 3, 11, 4 and 12
- **First Medium Domain:** Two IB HCAs, in PCIe slots 7 and 15
- **Second Medium Domain:** Two IB HCAs, in PCIe slots 8 and 16

## Determining the Cores Available for Domains and Zones

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

- [“Cores Available for Domains and Zones” on page 16](#)
- [“Example Configuration” on page 17](#)

### Cores Available for Domains and Zones

Because the number of processor modules in each SPARC T5-8 server varies, depending on whether you have a Half Rack or Full Rack, the number of cores available for each SPARC T5-8 server in your Oracle SuperCluster T5-8 also varies:

- For the Half Rack, each SPARC T5-8 server contains *two* processor modules. There are two sockets or PCIe root complex pairs on each processor module. Sixteen cores are associated with each socket, for a total of 64 cores (four sockets or PCIe root complex pairs) for each SPARC T5-8 server in the Half Rack.
- For the Full Rack, each SPARC T5-8 server contains *four* processor modules. There are two sockets or PCIe root complex pairs on each processor module. Sixteen cores are associated with each socket, for a total of 128 cores (eight sockets or PCIe root complex pairs) for each SPARC T5-8 server in the Full Rack.

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

A certain number of cores are always set aside for the global zone (the Database Domain), 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 size of the domain.

The following table provides information on the number of cores that are available for specific domains and zones that are created within those domains. See [“Understanding Domain Configurations” on page 12](#) for more information on the different types of domains.

**TABLE 1** Cores Available for the Domains and Zones

	Total Number of Cores Available for This Domain	Number of Cores Set Aside for Global Zone	Number of Cores Available for Nonglobal Zones
<b>Giant Domain (Full Rack Only)</b>	128	4	124

	Total Number of Cores Available for This Domain	Number of Cores Set Aside for Global Zone	Number of Cores Available for Nonglobal Zones
<b>Large Domain</b>	64	4	60
<b>Medium Domain</b>	32	4	28
<b>Small Domain</b>	16	2	14

When using the information in the table, keep in mind that the number of cores that are set aside for the global zone applies only when you are creating zones (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 on each SPARC T5-8 server. 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 that would be available for zones on certain types of Database Domains, consider the example configuration as described in [“Example Configuration” on page 73](#), where:

- The system is the Full Rack version of the Oracle SuperCluster T5-8
- The F4-2 LDom configuration (four LDom) is set up on both SPARC T5-8 servers in the Oracle SuperCluster T5-8 (see [“Full Rack Domain Configurations” on page 14](#) for more information on this specific type of LDom configuration)
- The following LDom types are assigned to each LDom in both SPARC T5-8 servers:
  - **LDom 1:** Database Domain
  - **LDom 2:** Solaris 10 Application Domain
  - **LDom 3:** Solaris 11 Application Domain
  - **LDom 4:** Database Domain

For the purposes of this exercise, assume the following:

- Four zones will be set up on LDom 1 (the first Database Domain) in both SPARC T5-8 servers.
- No zones will be set up on LDom 4 (the second Database Domain) in both SPARC T5-8 servers.

Using the information provided in [“Full Rack Domain Configurations” on page 14](#) and [“Cores Available for Domains and Zones” on page 16](#), you can determine the following information about the two Database Domains on each SPARC T5-8 server in this Full Rack:

- LDom 1 (the first Database Domain) has the following characteristics:
  - This domain is a Large Domain.
  - 64 cores are available for this domain.
- LDom 4 (the second Database Domain) has the following characteristics:
  - This domain is a Medium Domain.
  - 32 cores are available for this domain.

The following sections describe how cores could be allocated in both of these Database Domains in both SPARC T5-8 servers:

- [“LDom 1 \(First Database Domain\)” on page 18](#)
- [“LDom 4 \(Second Database Domain\)” on page 18](#)

## LDom 1 (First 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 SPARC T5-8 server will contain four zones, and it has 64 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 16 cores ( $64 / 4 = 16$ ).

However, based on the information provided in [“Cores Available for Domains and Zones” on page 16](#), a certain number of cores will be set aside for the global zone. This domain is a Large Domain, which means that four cores will be set aside for the global zone, so you actually have 60 cores available for the four zones in this Database Domain. So, for the zones in this Database Domain, you could allocate 15 cores for each of the four zones to use all 60 cores available for zones in this Database Domain ( $60 / 4 = 15$ ).

## LDom 4 (Second Database Domain)

This Database Domain has 32 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 domain, all 32 cores are available for this Database Domain, so you could allocate all 32 cores to this Database Domain on both SPARC T5-8 servers.

## Memory Available for Domains

The following memory configurations are available for the Oracle SuperCluster T5-8:

- **Half Rack:** 1 TB (1024 GB) of total memory in each SPARC T5-8 server
- **Full Rack:** 2 TB (2048 GB) of total memory in each SPARC T5-8 server

The amount of memory that is available for specific domains in each SPARC T5-8 server varies, depending on the type of domain:

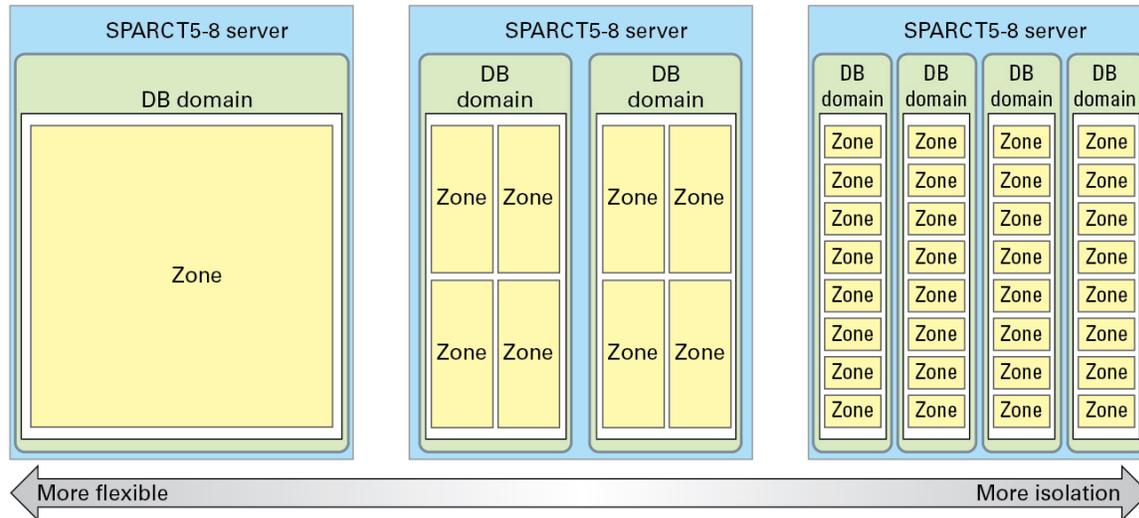
- **Giant Domain (Full Rack only):** 2 TB of memory
- **Large Domain (Full Rack or Half Rack):** 1 TB of memory
- **Medium Domain (Full Rack or Half Rack):** 512 GB of memory
- **Small Domain (Full Rack or Half Rack):** 256 GB of memory

## Guidelines for Planning Zones

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

- The number of zones that you want to create in that Database Domain. The maximum number of zones per domain is 8.
- The size of the Database Domain, which will determine the number of cores available for that domain and the zones that can be created 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)

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



## Guidelines for Planning Zone Clusters

Following are a list of guidelines regarding zone clusters:

- When you first install the operating system instances on a domain, that domain is automatically designated as the global zone. Any domains or zones that you create from that point forward 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 non-global zones, you cannot have those non-global 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.
- 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 Database Domains, and the cluster spans across both SPARC T5-8 servers so that there is not a single point of failure.
- 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 on the Database Domains on the two SPARC T5-8 servers, and those zones are clustered together, even if one of the Database Domains is a Medium Domain and the other Database Domain is a Small Domain. The private interface connections (the IB connections) are limited to the lower number from the smaller domain, based on the IB HCAs available to the zone.

## Guidelines for Planning the Storage Server Disk Group Layout

An Oracle representative should have set up Database Domain clusters 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 Exadata Storage Servers, the Exadata Storage Servers will have to be repartitioned. That action erases the clusters that were set up at your initial installation.

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

- If three Exadata 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 Exadata Storage Servers across each cluster, and the Exadata Storage Servers do not have to be repartitioned. For high redundancy, five Exadata Storage Servers should be available for every cluster with zones on the Database Domains.

For example, assume you have a Half Rack, which has four Exadata Storage Servers. You are creating one cluster with zones on the Database Domains and you are using normal redundancy, then you would not have to share Exadata Storage Servers.

Likewise, assume you have a Full Rack (with eight Exadata Storage Servers) and an Oracle Exadata Storage Expansion Full Rack connected to your Oracle SuperCluster T5-8 (with 18 Exadata Storage Servers). You are creating eight clusters with zones on the Database Domains and you are using normal redundancy, then you would have three Exadata Storage Servers for every cluster that you are creating. You would not have to share Exadata Storage Servers across the clusters.

- If fewer than three Exadata Storage Servers are available for every cluster with zones on the Database Domains using normal redundancy, or fewer than five Exadata Storage Servers

for every cluster if you are using high redundancy, then you must share Exadata Storage Servers across each cluster.

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

- [“Planning the Exadata Storage Server Disk Group Layout” on page 22](#)
- [“Understanding an Example Scenario” on page 27](#)

## Planning the Exadata Storage Server Disk Group Layout

In general, follow this process to determine how to lay out your zones and clusters across the Exadata 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 Exadata Storage Servers are available for the clusters that you are creating? How much total usable disk space is available for those clusters?
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.
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 Exadata Storage Server disk group layout in the worksheet in [“Recording Your Exadata Storage Server Disk Group Layout”](#) on page 38.

The following topics provide more detailed information on each of these decisions points.

- [“Determine the Available Existing Storage Capacity”](#) on page 23
- [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster”](#) on page 26
- [“Determine the Amount of Space for Each Disk Group in Each Cluster”](#) on page 27

## 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.

### Storage Capacity for X3-2 and X4-2 Exadata Storage Servers

Every Exadata Storage Server contains 12 drives. The capacity of those drives varies, depending on the type of drives that are installed:

- First Exadata Storage Server type (X3-2):
  - 600 GB High Performance SAS disk
  - 3 TB High Capacity SAS disk
- Second Exadata Storage Server type (X4-2):
  - 1.2 TB High Performance SAS disk
  - 4 TB High Capacity SAS disk

In addition, the total number of Exadata Storage Servers varies, depending on the type of Oracle SuperCluster T5-8 and Exadata Storage Expansion Rack that you have:

Rack Type	Quarter Rack	Half Rack	Full Rack
Oracle SuperCluster T5-8	N/A	4	8

Rack Type	Quarter Rack	Half Rack	Full Rack
Oracle Exadata Storage Expansion Rack	4	9	18

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 T5-8 and the Oracle Exadata Storage Expansion Rack.

**TABLE 2** Exadata Storage Server Storage Capacity in the Oracle SuperCluster T5-8, High-Capacity Version

Capacity Type	3 TB Disks (X3-2)	4 TB Disks (X4-2)
Raw disk capacity	<ul style="list-style-type: none"> <li>■ Half Rack: 144 TB</li> <li>■ Full Rack: 288 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 192 TB</li> <li>■ Full Rack: 384 TB</li> </ul>
Usable mirrored capacity (ASM normal redundancy)	<ul style="list-style-type: none"> <li>■ Half Rack: 64 TB</li> <li>■ Full Rack: 128 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 85.3 TB</li> <li>■ Full Rack: 170.6 TB</li> </ul>
Usable triple mirrored capacity (ASM high redundancy)	<ul style="list-style-type: none"> <li>■ Half Rack: 43 TB</li> <li>■ Full Rack: 86 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 57.3 TB</li> <li>■ Full Rack: 114.6 TB</li> </ul>
Raw flash capacity	<ul style="list-style-type: none"> <li>■ Half Rack: 6.4 TB</li> <li>■ Full Rack: 12.8 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 12.8 TB</li> <li>■ Full Rack: 25.6 TB</li> </ul>

**TABLE 3** Exadata Storage Server Storage Capacity in the Oracle SuperCluster T5-8, High-Performance Version

Capacity Type	600 GB Disks (X3-2)	1.2 TB Disks (X4-2)
Raw disk capacity	<ul style="list-style-type: none"> <li>■ Half Rack: 28 TB</li> <li>■ Full Rack: 56 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 57 TB</li> <li>■ Full Rack: 114 TB</li> </ul>
Usable mirrored capacity (ASM normal redundancy)	<ul style="list-style-type: none"> <li>■ Half Rack: 13 TB</li> <li>■ Full Rack: 26 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 26 TB</li> <li>■ Full Rack: 52 TB</li> </ul>
Usable triple mirrored capacity (ASM high redundancy)	<ul style="list-style-type: none"> <li>■ Half Rack: 8 TB</li> <li>■ Full Rack: 16 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 16 TB</li> <li>■ Full Rack: 32 TB</li> </ul>
Raw flash capacity	<ul style="list-style-type: none"> <li>■ Half Rack: 6.4 TB</li> <li>■ Full Rack: 12.8 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Half Rack: 12.8 TB</li> <li>■ Full Rack: 25.6 TB</li> </ul>

**TABLE 4** Exadata Storage Server Storage Capacity in the Oracle Exadata Storage Expansion Rack, High-Capacity Version

Capacity Type	3 TB Disks (X3-2)	4 TB Disks (X4-2)
Raw disk capacity	<ul style="list-style-type: none"> <li>■ Quarter Rack: 144 TB</li> <li>■ Half Rack: 324 TB</li> <li>■ Full Rack: 648 TB</li> </ul>	<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>

Capacity Type	3 TB Disks (X3-2)	4 TB Disks (X4-2)
Usable mirrored capacity (ASM normal redundancy)	<ul style="list-style-type: none"> <li>■ Quarter Rack: 64 TB</li> <li>■ Half Rack: 144 TB</li> <li>■ Full Rack: 288 TB</li> </ul>	<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>
Usable triple mirrored capacity (ASM high redundancy)	<ul style="list-style-type: none"> <li>■ Quarter Rack: 43 TB</li> <li>■ Half Rack: 97 TB</li> <li>■ Full Rack: 193.5 TB</li> </ul>	<ul style="list-style-type: none"> <li>■ Quarter Rack: 57.3 TB</li> <li>■ Half Rack: 129 TB</li> <li>■ Full Rack: 384 TB</li> </ul>
Raw flash capacity	<ul style="list-style-type: none"> <li>■ Quarter Rack: 6.4 TB</li> <li>■ Half Rack: 14.4 TB</li> <li>■ Full Rack: 28.8 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>

**TABLE 5** Exadata Storage Server Storage Capacity in the Oracle Exadata Storage Expansion Rack, High-Performance Version

Capacity Type	600 GB Disks (X3-2)	1.2 TB Disks (X4-2)
Raw disk capacity	<ul style="list-style-type: none"> <li>■ Quarter Rack: 28 TB</li> <li>■ Half Rack: 64 TB</li> <li>■ Full Rack: 128 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: 13 TB</li> <li>■ Half Rack: 29 TB</li> <li>■ Full Rack: 58 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: 8 TB</li> <li>■ Half Rack: 19 TB</li> <li>■ Full Rack: 39 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: 6.4 TB</li> <li>■ Half Rack: 14.4 TB</li> <li>■ Full Rack: 28.8 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>

## Storage Capacity for X5-2L and E6-2L Exadata Storage Servers

The number and capacity of the drives installed in the Exadata Storage Servers vary, depending on the type of drives that are installed:

- X5-2L Exadata Storage Server:
  - 8 x 1.6 TB Extreme Flash drives
  - 12 x 8 TB High Capacity drives
- X6-2L Exadata Storage Server:
  - 8 x 3.2 TB Extreme Flash drives
  - 12 x 8 TB High Capacity drives

In addition, 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 two types of Exadata Storage Servers.

**TABLE 6** X5-2L Single Storage Server Capacities

Capacity Type	8 x 1.6 TB Extreme Flash Drives	12 x 8 TB High Capacity Drives
Raw disk capacity	12.8 TB	96 TB
Usable mirrored capacity (ASM normal redundancy)	5 TB	40 TB
Usable triple mirrored capacity (ASM high redundancy)	4.3 TB	30 TB

**TABLE 7** X6-2L Single Storage Server Capacities

Capacity Type	8 x 3.2 TB Extreme Flash Drives	12 x 8 TB High Capacity Drives
Raw disk capacity	25.6 TB	96 TB
Usable mirrored capacity (ASM normal redundancy)	10 TB	40 TB
Usable triple mirrored capacity (ASM high redundancy)	8.6 TB	30 TB

## 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 on each SPARC T5-8 server, the size of each Database Domain, and so on.

See [“Guidelines for Planning Zones” on page 19](#) 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 Exadata 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 27](#)
- [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster \(Example Scenario\)” on page 28](#)
- [“Determine the Amount of Space for Each Disk Group in Each Cluster” on page 28](#)

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

Consider the following example scenario:

- A Half Rack version of the Oracle SuperCluster T5-8, with four Exadata Storage Servers
- 12 high-capacity (3 TB) disk drives in each Exadata Storage Server
- Each Exadata Storage Server set at normal (double-mirrored) redundancy

Using the information from the section [“Determine the Available Existing Storage Capacity” on page 23](#), you have **64 TB** of usable storage capacity available for the Exadata Storage Servers in your Oracle SuperCluster T5-8.

## **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 64 TB of usable disk capacity, and you want to create two clusters for the zones, you could use the existing 64 TB of space in one of several ways:

- Use the entire 64 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 32 TB.
  - Create two clusters of different sizes, where the total disk space used by the disk groups for the two clusters equals 64 TB. For example, one cluster with the total disk space used by the disk groups equaling 12 TB and the other cluster with the total disk space used by the disk groups equaling 52 TB.

Keep in mind that by using the entire 64 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 64 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 32 TB for each cluster.

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

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 32 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 T5-8, 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 Exadata Storage Servers.

- [“Recording Your Existing Configuration” on page 29](#)
- [“Recording Your Zone Configuration Information” on page 32](#)
- [“Recording Your Cluster Configuration Information” on page 36](#)
- [“Recording Your Exadata Storage Server Disk Group Layout” on page 38](#)

## Recording Your Existing Configuration

Several zone setup decisions and procedures are affected by the configuration choices that you made when you first purchased your Oracle SuperCluster T5-8. Record your existing configuration in the following tables, and use this information as a reference for certain zone setup decisions and procedures.

For each SPARC T5-8 server, provide the following information:

- The size of the Database Domains that you have on that server. Use this information to determine the maximum number of cores that you can use for each zone that you create on the Database Domains on that server.

For example, if you are setting up zones on a Small Domain in a Half Rack, then the maximum number of cores that you can assign to each zone created on that domain is sixteen. See [“Guidelines for Planning Zones” on page 19](#) for more information.

- The number of IB HCAs that are associated with each Database Domain. This information is used as part of the configuration process with the Oracle Exadata Deployment Assistant (OEDA).
- The specific location of each Database Domain. This information is also used to determine whether or not you must manually create VNETs for the host management network for certain Database Domains. See [“Host Management Network After Zone Creation” on page 55](#) for more information on such situations.
- The 1GbE host management host name and IP address for each domain.

The following sections provide worksheets to record your existing configuration:

- [“SPARC T5-8 Server 1 Configuration Information” on page 30](#)
- [“SPARC T5-8 Server 2 Configuration Information” on page 31](#)

## SPARC T5-8 Server 1 Configuration Information

Configuration Information	Entry
What type of domains do you have (Giant, Large, Medium, or Small)? How many IB HCAs are associated with that domain? What is the 1GbE host management host name and IP address for each domain?	
Domain 1	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 2	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 3	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 4	Domain type: Number of IB HCAs associated with this domain:

Configuration Information	Entry
	1GbE host management host name: 1GbE host management IP address:
Domain 5 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 6 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 7 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 8 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:

## SPARC T5-8 Server 2 Configuration Information

Configuration Information	Entry
What type of domains do you have (Giant, Large, Medium, or Small)? How many IB HCAs are associated with that domain? What is the 1GbE host management host name and IP address for each domain?	
Domain 1	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 2	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name:

Configuration Information	Entry
	1GbE host management IP address:
Domain 3	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 4	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 5 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 6 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 7 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:
Domain 8 (Full Rack)	Domain type: Number of IB HCAs associated with this domain: 1GbE host management host name: 1GbE host management IP address:

## Recording Your Zone Configuration Information

Use the tables in the following sections to record the following information for each SPARC T5-8 server:

- The number of zones that you want to deploy on each Database Domain (1 - 8)

- The size of each zone that you want to deploy on each Database Domain (the number of cores and VCPUs)
- 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 SPARC T5-8 server has entries for 8 zones. If you have more than 30 zones on each SPARC T5-8 server, copy the worksheets to record the information for the additional zones.

Because of the number of cores that are reserved for the global domain (the Database Domain) and the fact that you should have a minimum of one core for each zone, you can therefore have from 1 to 60 zones on each Database Domain in a Half Rack and from 1 to 124 zones on each Database Domain in a Full Rack. The number of cores that you have allocated to those zones varies, depending on the size of that zone. See [“Guidelines for Planning Zones” on page 19](#) for more information.

Also refer to the information that you entered in the [“Recording Your Existing Configuration” on page 29](#) worksheets for each SPARC T5-8 server to determine the maximum number of zones that you can have on each Database Domain, and the maximum number of cores that are available for those zones, depending on the sizes of the domains on that server.

For example, if you have a H4-1 domain configuration on a SPARC T5-8 server, then the following applies, based on the information provided in [“Half Rack Domain Configurations” on page 13](#) and [“Cores Available for Domains and Zones” on page 16](#):

- Each domain in the H4-1 domain configuration is a Small Domain.
- 16 cores are available for each Small Domain.
- For each Small Domain, 2 cores must be set aside for the global zone (the Database Domain). The remaining 14 cores are therefore available for the zones in each of these Small Domains.

Because you must have at least one core assigned to each zone, the maximum number of zones that you can have on any of the four domains is 14. You can then choose any of the following zone configurations for each domain:

- One zone, consisting of all 14 available cores for zones in this domain
- Two zones, consisting of 7 cores apiece
- Seven zones, consisting of 2 cores apiece
- Fourteen zones, consisting of 1 core apiece

Enter the zone configuration information in the following worksheets:

- [“SPARC T5-8 Server 1 Zone Configuration Information” on page 34](#)
- [“SPARC T5-8 Server 2 Zone Configuration Information” on page 35](#)

## SPARC T5-8 Server 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:
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:	Host management: Client access: IB:

Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
		VIP:	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:

### SPARC T5-8 Server 2 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:
4	Cores:	Host management: Client access: IB:	Host management: Client access: IB:

Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
		VIP:	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:

## 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

Clusters should span across multiple SPARC T5-8 servers for high availability. If you are creating a cluster, and a zone in a Database Domain is a member of that cluster on one SPARC T5-8 server, then other zones that are members of that cluster should be in a similar Database Domain on the second SPARC T5-8 server.

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, this cluster should be created with the following zones as members of that cluster:

- Zone 1 in Database Domain 1 in SPARC T5-8 server 1
- Zone 1 in Database Domain 1 in SPARC T5-8 server 2

The worksheet has entries for ten clusters. But you can have more than ten clusters in your Oracle SuperCluster T5-8, 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 Exadata Storage Server Disk Group Layout

Use the following table to record your Exadata 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 T5-8, 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 Exadata Storage Server Disk Group Layout” on page 22](#). For example, if the name of a particular cluster is osc01r1, then the names of the disk groups for that cluster should be 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:	

Cluster Number and Name	Name of Disk Groups	Size of Disk Groups	Total Size of All Disk Groups for This Cluster
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:	



# Understanding Network Considerations

---

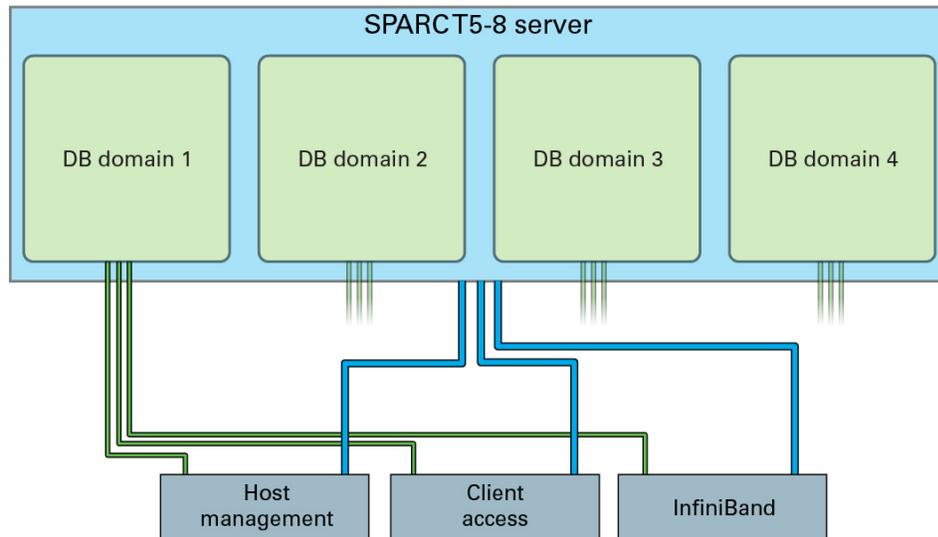
These topics describe how the network is set up before and after you create zones on the Database Domains.

- [“Understanding the Networking Setup Before Zone Creation” on page 41](#)
- [“Understanding the Networking Setup After Zone Creation” on page 54](#)
- [“Network Setup for Clusters” on page 57](#)

## Understanding the Networking Setup Before Zone Creation

Before you begin the process of setting up zones on the Database Domains, you should first understand how the networks were set up when you had your Oracle SuperCluster T5-8 installed. Getting a clear understanding of the existing network setup is helpful when understanding how the network setup will change when you configure zones on the Database Domains.

When your Oracle SuperCluster T5-8 was installed, the SPARC T5-8 servers and the domains created within each server were assigned host names and IP addresses that enabled for access to the host management, client access, and IB networks. The following figure shows a graphical representation of this initial setup.



The networks are configured differently, depending on the type of rack and the type of domains on each SPARC T5-8 server.

- [“Understanding the Network Configuration For the Half Rack” on page 42](#)
- [“Understanding the Network Configuration for the Full Rack” on page 47](#)

## Understanding the Network Configuration For the Half Rack

The following sections provide more detailed information on how the networks are configured at that initial installation:

- [“Host Management Network Before Zone Creation \(Half Rack\)” on page 42](#)
- [“Client Access Network Before Zone Creation \(Half Rack\)” on page 43](#)
- [“IB Network Before Zone Creation \(Half Rack\)” on page 45](#)

### Host Management Network Before Zone Creation (Half Rack)

There are four host management ports (NET0 to NET3) on each SPARC T5-8 server. The method that you use to connect to the host management network through these four host management ports varies, depending on the type of domain that you have set up on each SPARC T5-8 server:

- **Large Domain:** Two out of the four 1GbE host management ports are part of one IPMP group for this domain:
  - NET0
  - NET3
- **Medium Domains:** Two 1GbE host management ports are part of one IPMP group for each Medium Domain. Following are the 1GbE host management ports associated with each Medium Domain, depending on how many domains are on the SPARC T5-8 server in the Half Rack:
  - First Medium Domain: NET0-1
  - Second Medium Domain, if applicable: NET2-3
- **Small Domains:** The number and type of 1GbE host management ports that are assigned to each Small Domain varies, depending on the CPU that the Small Domain is associated with:
  - CPU0: NET0 to 1
  - CPU1: NET1, NET3 (through virtual network devices)
  - CPU6: NET0, NET2 (through virtual network devices)
  - CPU7: NET2-3

See [“Half Rack Domain Configurations” on page 13](#) for more information on which CPU would be associated with certain Small Domains.

## Client Access Network Before Zone Creation (Half Rack)

This required 10GbE network connects the SPARC T5-8 servers to your existing client network and is used for client access to the servers. 10GbE NICs installed in the PCIe slots are used for connection to this network. The number of 10GbE NICs and the PCIe slots that they are installed in varies depending on the type of configuration that is set up on the SPARC T5-8 server.

In each case, IPMP grouping provides redundancy to the 10GbE client access network, where two physical ports are part of one IPMP group and a single data address accesses these two physical ports. For the situations where the two physical ports are on separate dual-port 10GbE NICs, that data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Large Domain:** All four PCI root complex pairs, and therefore four 10GbE NICs, are associated with the logical domain on the server in this configuration. However, only two of the four 10GbE NICs are used with this domain. One port is used on each dual-ported 10GbE NIC. The two ports on the two separate 10GbE NICs are part of one IPMP group. One port from the dual-ported 10GbE NICs is connected to the 10GbE network in this case, with the remaining unused ports and 10GbE NICs unconnected.

The following 10GbE NICs and ports are used for connection to the client access network for this configuration:

- PCIe slot 1, port 0 (active)
- PCIe slot 14, port 1 (standby)

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Medium Domains:** Two PCI root complex pairs, and therefore two 10GbE NICs, are associated with the Medium Domain on the SPARC T5-8 server in the Half Rack. One port is used on each dual-ported 10GbE NIC. The two ports on the two separate 10GbE NICs are part of one IPMP group.

The following 10GbE NICs and ports are used for connection to the client access network for this configuration:

- First Medium Domain:
  - PCIe slot 1, port 0 (active)
  - PCIe slot 9, port 1 (standby)
- Second Medium Domain, if applicable:
  - PCIe slot 6, port 0 (active)
  - PCIe slot 14, port 1 (standby)

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Small Domains:** One PCI root complex pair, and therefore one 10GbE NIC, is associated with the Small Domain on the SPARC T5-8 server in the Half Rack. Both ports are used on each dual-ported 10GbE NIC, and both ports on each 10GbE NIC are part of one IPMP group. Both ports from each dual-ported 10GbE NIC are connected to the 10GbE network for the Small Domains.

The following 10GbE NICs and ports are used for connection to the client access network for the Small Domains, depending on the CPU that the Small Domain is associated with:

- CPU0:
  - PCIe slot 1, port 0 (active)
  - PCIe slot 1, port 1 (standby)
- CPU1:
  - PCIe slot 9, port 0 (active)
  - PCIe slot 9, port 1 (standby)
- CPU6:
  - PCIe slot 6, port 0 (active)

- PCIe slot 6, port 1 (standby)
- CPU7:
  - PCIe slot 14, port 0 (active)
  - PCIe slot 14, port 1 (standby)

See [“Half Rack Domain Configurations” on page 13](#) for more information on which CPU would be associated with certain Small Domains.

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if the connection to one of the two ports on the 10GbE NIC fails.

## IB Network Before Zone Creation (Half Rack)

Connection to the IB network is achieved through the IB host channel adapters (IB HCAs) installed in express module slots in each SPARC T5-8 server. There are four IB HCAs in each SPARC T5-8 server.

The IB HCAs, and the ports on each HCA, that are assigned to each domain vary, depending on the type of domains on each SPARC T5-8 server:

- **Large Domain:** The HCAs and the ports on those HCAs that are used to connect to the IB network vary, depending on the type of domain:
  - Database Domain:
    - **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the domain.
    - For a Large Domain in a Half Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 16 (standby).
    - **Exadata private network:** Connections through P0 (active) and P1 (standby) on all IB HCAs associated with the domain.
    - For a Large Domain in a Half Rack, connections are made through all four IB HCAs, with P0 on each as the active connection and P1 on each as the standby connection.
  - Application Domain:
    - **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the domain.
    - For a Large Domain in a Half Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 16 (standby).

- **Oracle Solaris Cluster private network:** Connections through P0 (active) on the IB HCA associated with the *second* CPU in the domain and P1 (standby) on the IB HCA associated with the *third* CPU in the domain.

For a Large Domain in a Half Rack, these connections are through P0 on the IB HCA installed in slot 11 (active) and P1 on the IB HCA installed in slot 8 (standby).

- **Medium Domains:** The HCAs and the ports on those HCAs that are used to connect to the IB network will vary, depending on the type of domain:

- Database Domain:

- **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *second* CPU in the domain.

For the first Medium Domain in a Half Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 11 (standby).

- **Exadata private network:** Connections through P0 (active) and P1 (standby) on all IB HCAs associated with the domain.

For the first Medium Domain in a Half Rack, connections are made through both IB HCAs (slot 3 and slot 11), with P0 on each as the active connection and P1 on each as the standby connection.

- Application Domain:

- **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *second* CPU in the domain.

For the first Medium Domain in a Half Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 11 (standby).

- **Oracle Solaris Cluster private network:** Connections through P0 (active) on the IB HCA associated with the *first* CPU in the domain and P1 (standby) on the IB HCA associated with the *second* CPU in the domain.

For first Medium Domain in a Half Rack, these connections are through P0 on the IB HCA installed in slot 3 (active) and P1 on the IB HCA installed in slot 11 (standby).

- **Small Domains:** The HCAs and the ports on those HCAs that are used to connect to the IB network will vary, depending on the type of domain:

- Database Domain:

- **Storage private network:** Connections through P1 (active) and P0 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P1 (active) and P0 (standby) on that IB HCA.

- **Exadata private network:** Connections through P0 (active) and P1 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P0 (active) and P1 (standby) on that IB HCA.

- Application Domain:

- **Storage private network:** Connections through P1 (active) and P0 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P1 (active) and P0 (standby) on that IB HCA.

- **Oracle Solaris Cluster private network:** Connections through P0 (active) and P1 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P0 (active) and P1 (standby) on that IB HCA.

## Understanding the Network Configuration for the Full Rack

The following sections provide more detailed information on how the networks would be configured at that initial installation:

- [“Host Management Network Before Zone Creation \(Full Rack\)” on page 47](#)
- [“Client Access Network Before Zone Creation \(Full Rack\)” on page 48](#)
- [“IB Network Before Zone Creation \(Full Rack\)” on page 51](#)

### Host Management Network Before Zone Creation (Full Rack)

There are four host management ports (NET0–NET3) on each SPARC T5-8 server. The method that you use to connect to the host management network through these four host management ports varies, depending on the type of domain that you have set up on each SPARC T5-8 server:

- **Giant Domain:** Two out of the four 1GbE host management ports are part of one IPMP group for this domain:
  - NET0
  - NET3

- **Large Domains:** Two 1GbE host management ports are part of one IPMP group for each Large Domain. Following are the 1GbE host management ports associated with each Large Domain, depending on how many domains are on the SPARC T5-8 server in the Full Rack:
  - First Large Domain: NET0-1
  - Second Large Domain, if applicable: NET2-3
- **Medium Domains:** The number and type of 1GbE host management ports that are assigned to each Medium Domain varies, depending on the CPU that the Medium Domain is associated with:
  - CPU0/CPU1: NET0 to 1
  - CPU2/CPU3: NET1, NET3 (through virtual network devices)
  - CPU4/CPU5: NET0, NET2 (through virtual network devices)
  - CPU6/CPU7: NET2 to 3

See “[Full Rack Domain Configurations](#)” on page 14 for more information on which CPU would be associated with certain Small Domains.
- **Small Domains:** The number and type of 1GbE host management ports that are assigned to each Small Domain varies, depending on the CPU that the Small Domain is associated with:
  - CPU0: NET0-1
  - CPU1: NET1, NET3 (through virtual network devices)
  - CPU2: NET0, NET2 (through virtual network devices)
  - CPU3: NET1, NET3 (through virtual network devices)
  - CPU4: NET0, NET2 (through virtual network devices)
  - CPU5: NET1, NET3 (through virtual network devices)
  - CPU6: NET0, NET2 (through virtual network devices)
  - CPU7: NET2-3

See “[Full Rack Domain Configurations](#)” on page 14 for more information on which CPU would be associated with certain Small Domains.

## Client Access Network Before Zone Creation (Full Rack)

This required 10GbE network connects the SPARC T5-8 servers to your existing client network and is used for client access to the servers. 10GbE NICs installed in the PCIe slots are used for connection to this network. The number of 10GbE NICs and the PCIe slots that they are installed in varies depending on the type of configuration that is set up on the SPARC T5-8 server.

In each case, IPMP grouping provides redundancy to the 10GbE client access network, where two physical ports are part of one IPMP group and a single data address accesses these two physical ports. For the situations where the two physical ports are on separate dual-port 10GbE

NICs, that data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Giant Domain:** All eight PCI root complex pairs, and therefore eight 10GbE NICs, are associated with the logical domain on the server in this configuration. However, only two of the eight 10GbE NICs are used with this domain. One port is used on each dual-ported 10GbE NIC. The two ports on the two separate 10GbE NICs is part of one IPMP group. One port from the dual-ported 10GbE NICs is connected to the 10GbE network in this case, with the remaining unused ports and 10GbE NICs unconnected.

The following 10GbE NICs and ports are used for connection to the client access network for this configuration:

- PCIe slot 1, port 0 (active)
- PCIe slot 14, port 1 (standby)

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Large Domains:** Four PCI root complex pairs, and therefore four 10GbE NICs, are associated with the Large Domain on the SPARC T5-8 server in the Full Rack. One port is used on each dual-ported 10GbE NIC. The two ports on the two separate 10GbE NICs are part of one IPMP group.

The following 10GbE NICs and ports are used for connection to the client access network for this configuration:

- First Large Domain:
  - PCIe slot 1, port 0 (active)
  - PCIe slot 10, port 1 (standby)
- Second Large Domain, if applicable:
  - PCIe slot 5, port 0 (active)
  - PCIe slot 14, port 1 (standby)

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10GbE NICs fail.

- **Medium Domains:** Two PCI root complex pairs, and therefore two 10GbE NICs, are associated with the Medium Domain on the SPARC T5-8 server in the Full Rack. One port is used on each dual-ported 10GbE NIC. The two ports on the two separate 10GbE NICs are part of one IPMP group.

The following 10GbE NICs and ports are used for connection to the client access network for the Small Domains, depending on the CPU that the Small Domain is associated with:

- CPU0/CPU1:
  - PCIe slot 1, port 0 (active)
  - PCIe slot 9, port 1 (standby)

- CPU2/CPU3:
  - PCIe slot 2, port 0 (active)
  - PCIe slot 10, port 1 (standby)
- CPU4/CPU5:
  - PCIe slot 5, port 0 (active)
  - PCIe slot 13, port 1 (standby)
- CPU6/CPU7:
  - PCIe slot 6, port 0 (active)
  - PCIe slot 14, port 1 (standby)

See [“Full Rack Domain Configurations” on page 14](#) for more information on which CPU are associated with certain Small Domains.

A single data address is used to access these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if the connection to one of the two ports on the 10GbE NIC fails.

- **Small Domains:** One PCI root complex pair, and therefore one 10GbE NIC, is associated with the Small Domain on the SPARC T5-8 server in the Full Rack. Both ports are used on each dual-ported 10GbE NIC, and both ports on each 10GbE NIC would be part of one IPMP group. Both ports from each dual-ported 10GbE NIC is connected to the 10GbE network for the Small Domains.

The following 10GbE NICs and ports are used for connection to the client access network for the Small Domains, depending on the CPU that the Small Domain is associated with:

- CPU0:
  - PCIe slot 1, port 0 (active)
  - PCIe slot 1, port 1 (standby)
- CPU1:
  - PCIe slot 9, port 0 (active)
  - PCIe slot 9, port 1 (standby)
- CPU2:
  - PCIe slot 2, port 0 (active)
  - PCIe slot 2, port 1 (standby)
- CPU3:
  - PCIe slot 10, port 0 (active)
  - PCIe slot 10, port 1 (standby)
- CPU4:
  - PCIe slot 5, port 0 (active)

- PCIe slot 5, port 1 (standby)
- CPU5:
  - PCIe slot 13, port 0 (active)
  - PCIe slot 13, port 1 (standby)
- CPU6:
  - PCIe slot 6, port 0 (active)
  - PCIe slot 6, port 1 (standby)
- CPU7:
  - PCIe slot 14, port 0 (active)
  - PCIe slot 14, port 1 (standby)

A single data address accesses these two physical ports. That data address allows traffic to continue flowing to the ports in the IPMP group, even if the connection to one of the two ports on the 10GbE NIC fails.

## IB Network Before Zone Creation (Full Rack)

Connection to the IB network is achieved through the IB host channel adapters (HCAs) installed in express module slots in each SPARC T5-8 server. There are four IB HCAs in each SPARC T5-8 server.

The IB HCAs, and the ports on each HCA, that are assigned to each domain vary, depending on the type of domains on each SPARC T5-8 server:

- **Giant Domain:** The HCAs and the ports on those HCAs that are used to connect to the IB network vary, depending on the type of domain:
  - Database Domain:
    - **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the *first* processor module (PM0) in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the *last* processor module (PM3) in the domain.
    - For a Giant Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 16 (standby).
    - **Exadata private network:** Connections through P0 (active) and P1 (standby) on all IB HCAs associated with the domain.
    - For a Giant Domain in a Full Rack, connections are made through all eight IB HCAs, with P0 on each as the active connection and P1 on each as the standby connection.
  - Application Domain:

- **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the *first* processor module (PM0) in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the *last* processor module (PM3) in the domain.

For a Giant Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 16 (standby).

- **Oracle Solaris Cluster private network:** Connections through P0 (active) on the IB HCA associated with the *first* CPU in the *second* processor module (PM1) in the domain and P1 (standby) on the IB HCA associated with the *first* CPU in the *third* processor module (PM2) in the domain.

For a Giant Domain in a Full Rack, these connections are through P0 on the IB HCA installed in slot 4 (active) and P1 on the IB HCA installed in slot 7 (standby).

- **Large Domains:** The HCAs and the ports on those HCAs that are used to connect to the IB network will vary, depending on the type of domain:

- Database Domain:

- **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the domain.

For example, for the first Large Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 12 (standby).

- **Exadata private network:** Connections through P0 (active) and P1 (standby) on all IB HCAs associated with the domain.

For a Large Domain in a Full Rack, connections are made through all eight IB HCAs, with P0 on each as the active connection and P1 on each as the standby connection.

- Application Domain:

- **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *last* CPU in the domain.

For example, for the first Large Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 12 (standby).

- **Oracle Solaris Cluster private network:** Connections through P0 (active) on the IB HCA associated with the *second* CPU in the domain and P1 (standby) on the IB HCA associated with the *third* CPU in the domain.

For example, for the first Large Domain in a Full Rack, these connections are through P0 on the IB HCA installed in slot 11 (active) and P1 on the IB HCA installed in slot 4 (standby).

- **Medium Domains:** The HCAs and the ports on those HCAs that are used to connect to the IB network will vary, depending on the type of domain:
  - Database Domain:
    - **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *second* CPU in the domain.  
For example, for the first Medium Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 11 (standby).
    - **Exadata private network:** Connections through P0 (active) and P1 (standby) on all IB HCAs associated with the domain.  
For example, for the first Medium Domain in a Full Rack, connections are made through both IB HCAs (slot 3 and slot 11), with P0 on each as the active connection and P1 on each as the standby connection.
  - Application Domain:
    - **Storage private network:** Connections through P1 (active) on the IB HCA associated with the *first* CPU in the domain and P0 (standby) on the IB HCA associated with the *second* CPU in the domain.  
For example, for the first Medium Domain in a Full Rack, these connections are through P1 on the IB HCA installed in slot 3 (active) and P0 on the IB HCA installed in slot 11 (standby).
    - **Oracle Solaris Cluster private network:** Connections through P0 (active) on the IB HCA associated with the *first* CPU in the domain and P1 (standby) on the IB HCA associated with the *second* CPU in the domain.  
For example, for first Medium Domain in a Full Rack, these connections are through P0 on the IB HCA installed in slot 3 (active) and P1 on the IB HCA installed in slot 11 (standby).
- **Small Domains:** The HCAs and the ports on those HCAs that are used to connect to the IB network will vary, depending on the type of domain:
  - Database Domain:
    - **Storage private network:** Connections through P1 (active) and P0 (standby) on the IB HCA associated with the CPU associated with each Small Domain.  
For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P1 (active) and P0 (standby) on that IB HCA.
    - **Exadata private network:** Connections through P0 (active) and P1 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P0 (active) and P1 (standby) on that IB HCA.

- Application Domain:

- **Storage private network:** Connections through P1 (active) and P0 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

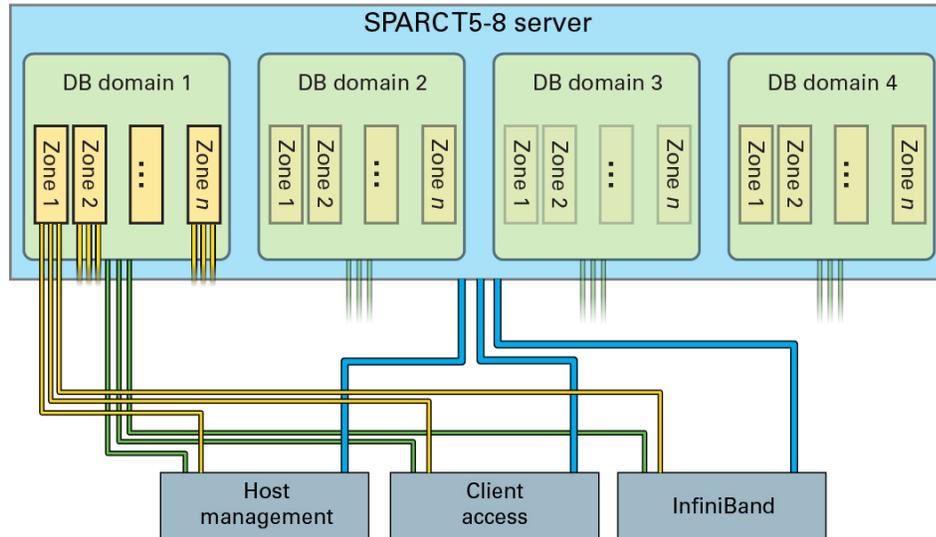
- For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P1 (active) and P0 (standby) on that IB HCA.

- **Oracle Solaris Cluster private network:** Connections through P0 (active) and P1 (standby) on the IB HCA associated with the CPU associated with each Small Domain.

- For example, for a Small Domain that is associated with CPU0, these connections are through the IB HCA installed in slot 3, through ports P0 (active) and P1 (standby) on that IB HCA.

## Understanding the Networking Setup After Zone Creation

After you create zones on the Database Domains, the existing networking setup that was configured before the zone creation remains in place with no modifications (see [“Understanding the Networking Setup Before Zone Creation” on page 41](#) for more information on the networking setup at the initial installation of your Oracle SuperCluster T5-8). When you create zones on the Database Domains, additional virtual networking interfaces are created to accommodate those zones, allowing the zones access to the host management, client access and IB networks. The following figure shows a graphical representation of how the zones would connect to these networks.



These topics describe how the networks are set up for the zones that you are creating:

- [“Host Management Network After Zone Creation” on page 55](#)
- [“Client Access Network After Zone Creation” on page 56](#)
- [“IB Network After Zone Creation” on page 56](#)

## Host Management Network After Zone Creation

In most cases, when you configure zones on the Database Domains, part of the zone creation process also automatically creates virtual network interface cards (VNICs), which are used as access to the host management network for the zones. You can create VNICs on top of physical interfaces, so in the situations where access to the host management network for the domains is achieved through actual physical interfaces, zones on those domains can use VNICs to connect to the host management network.

However, some domains do not connect to the host management network through actual physical interfaces, but rather through VNETs. Due to limitations with Oracle Solaris, you can not create VNICs on top of virtual network interfaces (VNETs). When creating zones on those domains that connect to the host management network through VNETs, you must manually create VNETs instead, alongside the domain VNETs, to use for the zones to access the host management network.

Whether or not VNICs are created automatically for you as part of the zone creation process, or if you have to manually create VNETs to use as access to the host management network for the

zones you want to create on the Database Domains, depends on the types of domains that you have on each SPARC T5-8 server and the zones that you want to create on certain domains.

- **Half Rack:** You must manually create VNETs for the Small Domains that are associated with the following CPUs:

- CPU1
- CPU6

See [“Half Rack Domain Configurations” on page 13](#) for more information on which Small Domains would have these CPUs associated with them.

- **Full Rack:** You must manually create VNETs for the following domains:

- Medium Domains associated with the following CPUs:
  - CPU2/CPU3
  - CPU4/CPU5
- Small Domains associated with the following CPUs:
  - CPU1
  - CPU2
  - CPU3
  - CPU4
  - CPU5
  - CPU6

See [“Full Rack Domain Configurations” on page 14](#) for more information on which Medium Domains and Small Domains would have these CPUs associated with them.

## Client Access Network After Zone Creation

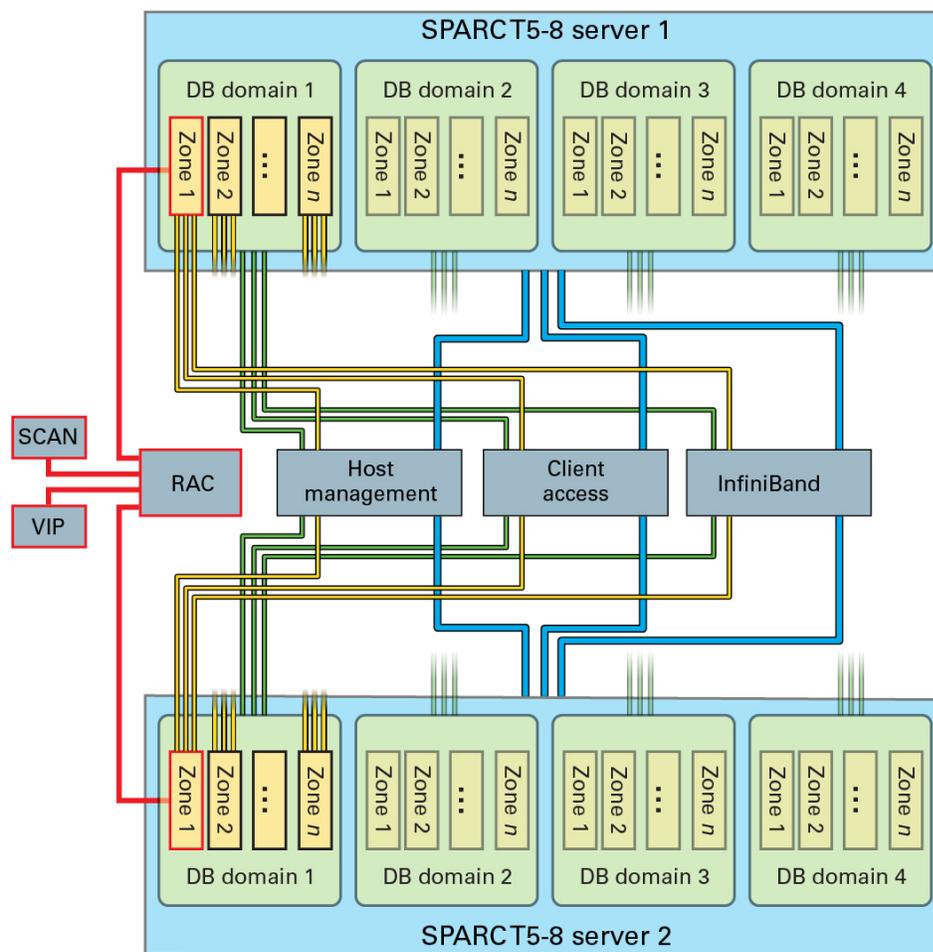
When you configure zones on the Database Domains, part of the zone creation process also automatically creates the VNICs that access the 10GbE client access network for the zones. You should not have to manually configure the VNICs that are used for the client access network for the zones.

## IB Network After Zone Creation

When you configure zones on the Database Domains, part of the zone creation process also automatically creates the VNICs that access the IB network for the zones. You should not have to manually configure the VNICs that are used for the IB network for the zones.

## Network Setup for Clusters

As part of the zone creation process, you can also cluster zones together, where they are members of a cluster. Access to the client access network for each cluster is achieved through the Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) addresses. The following figure shows a graphical representation of how clusters are accessed through the client access network.



In this respect, access to the clusters that have zones as members is no different than how clusters are accessed when clusters are used on two or more physical computers (each with an Oracle RDBMS instance) to concurrently access a single database.

# 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 59](#)
- [“Existing Network Worksheet” on page 61](#)
- [“Update the Base Software” on page 62](#)
- [“Determine the Repository Location” on page 63](#)
- [“Install or Update Packages From the Remote Repository” on page 64](#)
- [“Install or Update Files From the Local Repository” on page 65](#)
- [“Verify Configuration Tool Installation” on page 70](#)

## ▼ Prepare to Configure Zones

Complete these tasks before setting up zones and clusters on the Database Domains with zones in the Oracle SuperCluster T5-8.

### 1. **Back up the data on your Exadata Storage Servers.**

You might have to repartition the drives in the Exadata 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 21](#) for more information. Refer to the Exadata 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:
  - 1GbE administration (management) network
  - 10GbE client access network
  - Private IB network
  - Backup/Data Guard network (optional)

### 3. Plan the Exadata Storage Server layout.

You must make the following decisions on the Exadata Storage Server layout before beginning the processes in this document:

- The number of Exadata 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 T5-8
- IP addresses and host names for Oracle ILOM for the following components:
  - SPARC T5-8 servers with Database Domains that you are setting up zones on
  - Exadata Storage Servers
- IP addresses and host names for the 1GbE host management administration network for the following components:
  - Exadata Storage Servers
  - Switches (IB leaf and spine switches, and the Cisco switch)
  - PDUs
- IP addresses and host names for the IB network for the Exadata Storage Servers

In addition, as part of the initial installation when your Oracle SuperCluster T5-8 was first set up, the Oracle installer should have saved a configuration file 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.

You need this information and configuration files for certain procedures in this document.

**5. Record the information in the Existing Network Worksheet.**

See [“Existing Network Worksheet”](#) on page 61.

## Existing Network Worksheet

Existing Host Names and IP Addresses		
Component	Host Name	IP Address
Oracle ILOM for SPARC T5-8 Server 1		
Oracle ILOM for SPARC T5-8 Server 2		
Oracle ILOM for Exadata Storage Server 1		
Oracle ILOM for Exadata Storage Server 2		
Oracle ILOM for Exadata Storage Server 3		
Oracle ILOM for Exadata Storage Server 4		
Oracle ILOM for Exadata Storage Server 5		
Oracle ILOM for Exadata Storage Server 6		
Oracle ILOM for Exadata Storage Server 7		
Oracle ILOM for Exadata Storage Server 8		
Host management network for Exadata Storage Server 1		
Host management network for Exadata Storage Server 2		
Host management network for Exadata Storage Server 3		
Host management network for Exadata Storage Server 4		
Host management network for Exadata Storage Server 5		
Host management network for Exadata Storage Server 6		
Host management network for Exadata Storage Server 7		
Host management network for Exadata Storage Server 8		
Host management network for IB spine switch		
Host management network for IB leaf switch 1		
Host management network for IB leaf switch 2		
Host management network for PDU A		

Existing Host Names and IP Addresses		
Component	Host Name	IP Address
Host management network for PDU B		
IB network for Exadata Storage Server 1		
IB network for Exadata Storage Server 2		
IB network for Exadata Storage Server 3		
IB network for Exadata Storage Server 4		
IB network for Exadata Storage Server 5		
IB network for Exadata Storage Server 6		
IB network for Exadata Storage Server 7		
IB network for Exadata Storage Server 8		

## ▼ Update the Base Software

1. Log in to the first SPARC T5-8 server.
2. Determine if your system has the correct level of base software installed.

```
# pkg info entire
```

- If you see output similar to the following, then you have the correct level of software installed. Go to [“Determine the Repository Location” on page 63](#):

```
FMRI: pkg://solaris/entire@0.5.11,5.11-0.175.1.7.0.5.0:20130504T024145Z
```

where the values shown in **bold** above denote:

- Solaris 11 Update 1 OS
- SRU 7
- Build 5 or later
- If the SRU version is less than SRU 7.5, then you do not have the correct level of software installed. Follow these instructions to get the latest version of the software:
  - 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 T5-8.**  
Select any of the Oracle SuperCluster T5-8 options.
- e. **In the Release field, select the Oracle SuperCluster T5-8 1.5 release or later.**
- f. **Click the Search button.**
- g. **Select the QUARTERLY FULL STACK DOWNLOAD PATCH FOR EXADATA (*release-date*) (PATCH) option.**
- h. **Click the Read Me button to get installation instructions for the patch.**
- i. **Click the Download button to download the ZIP files for the patch.**

## ▼ Determine the Repository Location

Your Oracle SuperCluster T5-8 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 T5-8

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 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 64.](#)

```
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 T5-8. Go to [“Install or Update Files From the Local Repository” on page 65.](#)

```
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-examv 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 3a](#) 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 70](#).

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 70](#).

## ▼ Install or Update Files From the Local Repository

This procedure describes<sup>7</sup> how to install or update files from the local repository on the ZFS storage appliance.

---

**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, if 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 T5-8. Select any of the Oracle SuperCluster T5-8 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.



- a. **Locate the ISO image ZIP file that you downloaded in [Step 8](#) and copy that file to a temporary directory on the first SPARC T5-8 server.**

For example, copy the ZIP file to the `/var/tmp` directory on the first SPARC T5-8 server.

- b. **Unzip the ISO image ZIP file.**
- c. **Locate the ISO image in this directory.**
- 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 63.

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 will be 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 70.**

## ▼ **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:

- `advnet-wrapper.sh`
  - `advnet.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 Java OneCommand utility later on in this process.**

Go to [“Creating Configuration Files Using OEDA”](#) on page 73.



# Creating Configuration Files Using OEDA

---

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

- [“Example Configuration” on page 73](#)
- [“Verify Exadata Storage Server Disk Space for Additional Zones” on page 74](#)
- [“Determine Software Build Level” on page 75](#)
- [“Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier” on page 76](#)
- [“Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later” on page 120](#)

## Example Configuration

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

- The system is the Full Rack version of the Oracle SuperCluster T5-8, with Exadata Storage Servers with 1.2 TB high performance disks
- The F4-2 LDom configuration (four LDom) is set up on both SPARC T5-8 servers in the Oracle SuperCluster T5-8.
- The following LDom types are assigned to each LDom in both SPARC T5-8 servers:
  - **LDom 1:** Database Domain with four zones
  - **LDom 2:** Solaris 10 Application Domain
  - **LDom 3:** Solaris 11 Application Domain
  - **LDom 4:** Database Domain (without zones)
- Five clusters are set up on the Oracle SuperCluster T5-8:
  - Four clusters for the eight zones on the Database Domain with zones (four zones on each SPARC T5-8 server, where two zones, one from each SPARC T5-8 server, are members of a cluster)
  - One cluster for the two Database Domains (without zones)

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 compute servers. So for the zones on Database Domains with zones for the Oracle SuperCluster T5-8, these zones are members of the first cluster:

- First zone in the Database Domain with zones in the *first* SPARC T5-8 server
- First zone in the Database Domain with zones in the *second* SPARC T5-8 server

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 Exadata 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 Exadata Storage Servers:

- If you are not sharing Exadata Storage Servers across clusters and you have entire Exadata 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 Exadata Storage Servers to the new zones as part of this procedure.
- If you are sharing Exadata Storage Servers across clusters, you must determine how much free space is available on the Exadata Storage Servers that you are using for these new zones before proceeding.

---

**Note** - If you are sharing Exadata Storage Servers across clusters, you must follow the procedures in this section so that you will have the correct information before deciding on the size of the disk groups for the new clusters that you will be creating. Failure to do so could result in your existing disk groups being overwritten.

---

### 1. Determine the amount of free space available on the Exadata 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 Exadata Storage Servers.**
3. **Use the information on the amount of available space on each of the Exadata 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. **Determine the software build level on your system.**

See [“Determine Software Build Level” on page 75](#).

## ▼ Determine Software Build Level

The method you use to create configuration files using OEDA varies, depending on the software build level on your system. Follow these procedures to determine the software build level on your system:

1. **Type the following command:**

```
# svcprop -p configuration/build ssc/identification
```

- If the following error message appears:

```
svcprop: Pattern 'ssc/identification' doesn't match any entities
```

then you probably have software build 1.5.8 or later on your system. Go to [Step 2](#) to verify.

- If output similar to the following appears:

```
ssc-1.5.3
```

- If the output shows a software build level of **1.5.4 or earlier**, go to “[Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier](#)” on page 76.
- If the output shows a software build level of **1.5.8 or later**, go to “[Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later](#)” on page 120.

## 2. Type the following command:

```
# svcprop -p configuration/build svc:/system/oes/id:default
```

Output similar to the following appears:

```
ssc-1.5.10
```

- If the output shows a software build level of **1.5.4 or earlier**, go to “[Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier](#)” on page 76.
- If the output shows a software build level of **1.5.8 or later**, go to “[Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later](#)” on page 120.

## Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier



---

**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.

---

**Note** - The procedures in this section apply if you have a software build level of 1.5.4 or earlier on your system, as determined in the section [“Determine Software Build Level” on page 75](#). If your software build level is 1.5.8 or later, follow the procedures in [“Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later” on page 120](#).

---

Follow these procedures to use OEDA to create your configuration files for a system with a software build of 1.5.4 or earlier:

- [“Oracle Exadata Deployment Assistant Overview” on page 78](#)
- [“Set Up the OEDA” on page 79](#)
- [“Start the OEDA” on page 79](#)
- [“Complete the Customer Details Page” on page 81](#)
- [“Complete the Hardware Selection Page” on page 83](#)
- [“Complete the Define Customer Networks Page” on page 87](#)
- [“Complete the Administration Network Page” on page 89](#)
- [“Complete the Client Ethernet Network Page” on page 95](#)
- [“Complete the IB Network Page” on page 98](#)
- [“Complete the Identify Compute Node OS and Enable Capacity-on-Demand Configuration Page” on page 101](#)
- [“Review the Information in the Management and Private Networks Page” on page 102](#)
- [“Complete the Define Clusters Page” on page 104](#)
- [“Set Zone Default Configurations” on page 108](#)
- [“Complete the Cluster Page” on page 111](#)
- [“Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page” on page 115](#)
- [“Verify Remaining Configuration Information” on page 117](#)
- [“Generate the Configuration Files” on page 117](#)

## Oracle Exadata Deployment Assistant Overview

The Oracle Exadata Deployment Assistant is normally used when the system is initially installed and configured, and its purpose is to gather networking information for the Database Domains and for the physical components in the Oracle SuperCluster T5-8, such as the SPARC T5-8 servers, Exadata Storage Servers, switches and PDUs. However, in this section, you will be using the Oracle Exadata Deployment Assistant to gather networking information for the clusters and single-instance database zones that you want to set up.

Because you will be using the Oracle Exadata Deployment Assistant to gather networking information for clusters and single-instance database zones, rather than for physical components as the tool was originally intended for, it's useful to keep these things in mind as you perform the procedures in this section:

- The Database Machine Prefix would normally refer to the name of the Oracle SuperCluster T5-8 as a whole, and would be used as a prefix for the hostnames for the components inside the Oracle SuperCluster T5-8. However, for the procedures in this section, the Database Machine Prefix refers to the name for the cluster that you will be setting up, and will be used as a prefix for the zones that are a part of that cluster.
- The Oracle Database Deployment Assistant would normally assign hostnames to the compute nodes (the SPARC T5-8 servers) in several areas, such as the Administration Network. However, in some areas in this section, information that is being entered for compute nodes will actually apply to the zones that are part of this particular cluster, not to the compute nodes themselves. The instructions in this section will point out when information that needs to be entered for compute nodes actually applies to zones rather than compute nodes.
- In other areas in this process, information that is being entered for a compute node actually does apply to the compute node and not to a zone in the cluster, such as the ILOM host name and IP address for the compute node. The instructions in this section will point out when information that needs to be entered for compute nodes actually does apply for the compute nodes, and not for zones in this cluster.
- When the Oracle Database Deployment Assistant is first run on your Oracle SuperCluster T5-8 during the initial installation and configuration, part of the process would include assigning host names and IP addresses to the switches and PDUs. Because you are using this same tool in this situation to gather networking information to be used for the clusters and single-instance database zones, certain screens will have networking information for the switches and PDUs that will be based on the information you are entering for this particular cluster, and would therefore be different from the information that was originally entered for those components when your system was first set up. The networking information for the switches and PDUs are not cluster-dependant, so you will be asked to manually change the networking information for those components back to their original values several times in these procedures.

## ▼ Set Up the OEDA

### 1. Prepare to configure zones.

See [“Preparing to Configure Zones on Database Domains” on page 59](#).

### 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 70](#) 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 79](#).

## ▼ Start the OEDA

### 1. Start the OEDA.

- On Linux, MacOS, or the Oracle Solaris OS, type:

```
config.sh
```

- On Microsoft Windows, type:

```
config.cmd
```

## Start the OEDA

The OEDA Welcome window is displayed.



2. In the Welcome window, click Next to begin filling out the forms in the Oracle Exadata Deployment Assistant.

The Customer Details page is displayed. Go to [“Complete the Customer Details Page”](#) on page 81.

## ▼ Complete the Customer Details Page

1. Complete the fields in the Customer Details page.

The value that you enter in the Name Prefix will be the prefix that will be used for the prefix for the host names for the zones or clusters that you will be creating. This value should be eight characters or fewer. This prefix will be used for this set of zones that you are entering information for.

In addition, the values that you enter in the Customer Name and the Application fields in this page will be used as the name of the directory that will contain the configuration files at the end of this process.

---

**Note** - When creating multiple clusters, you will repeat these procedures for each cluster that you want to create. When you get to the Customer Details page for every cluster that you want to create, you must use the same values for the Customer Name and the Application fields in the Customer Details page for every pass through the Oracle Exadata Deployment Assistant so that all the configuration files are stored in a single directory.

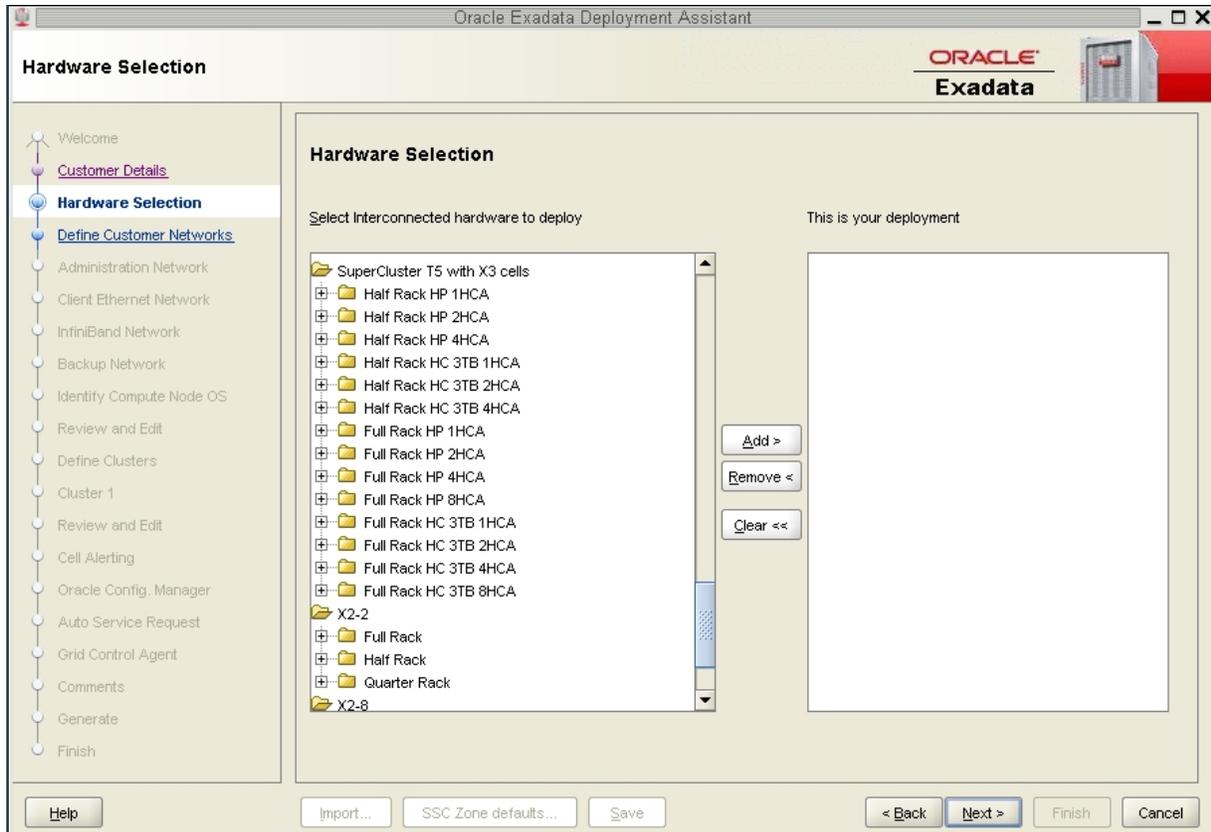
---

**2. When you have completed the Customer Details page, click Next.**

The Hardware Selection page is displayed. Go to [“Complete the Hardware Selection Page” on page 83.](#)

## ▼ Complete the Hardware Selection Page

### 1. Review the Hardware Selection page.



### 2. Determine what version of the Exadata Storage Servers that you have in your Oracle SuperCluster T5-8:

- If you have the older, X3 version of the Exadata Storage Servers, scroll down the list until you see the SuperCluster T5 with X3 cells folder.
- If you have the newer, X4 version of the Exadata Storage Servers, scroll down the list until you see the SuperCluster T5 with X4 cells folder.

### 3. Open the proper folder in the Hardware Selection page for your configuration.

You must determine the following things before you can make a choice in this page:

- The type of Oracle SuperCluster T5-8 that you have (Half Rack or Full Rack).
- What domain configuration you have set up on the SPARC T5-8 servers in your Oracle SuperCluster T5-8 (H1-1 configuration for a Half Rack, F1-1 configuration for a Full Rack, and so on). See [“Understanding Domain Configurations” on page 12](#) for more information.
- Which Database Domain you will be creating zones on in that domain configuration.
- The number of IB HCAs that are associated with that Database Domain. See [“Understanding Domain Configurations” on page 12](#) for more information.
- Whether you have high performance disks (HP) or high capacity (HC) disk in the Exadata Storage Servers.

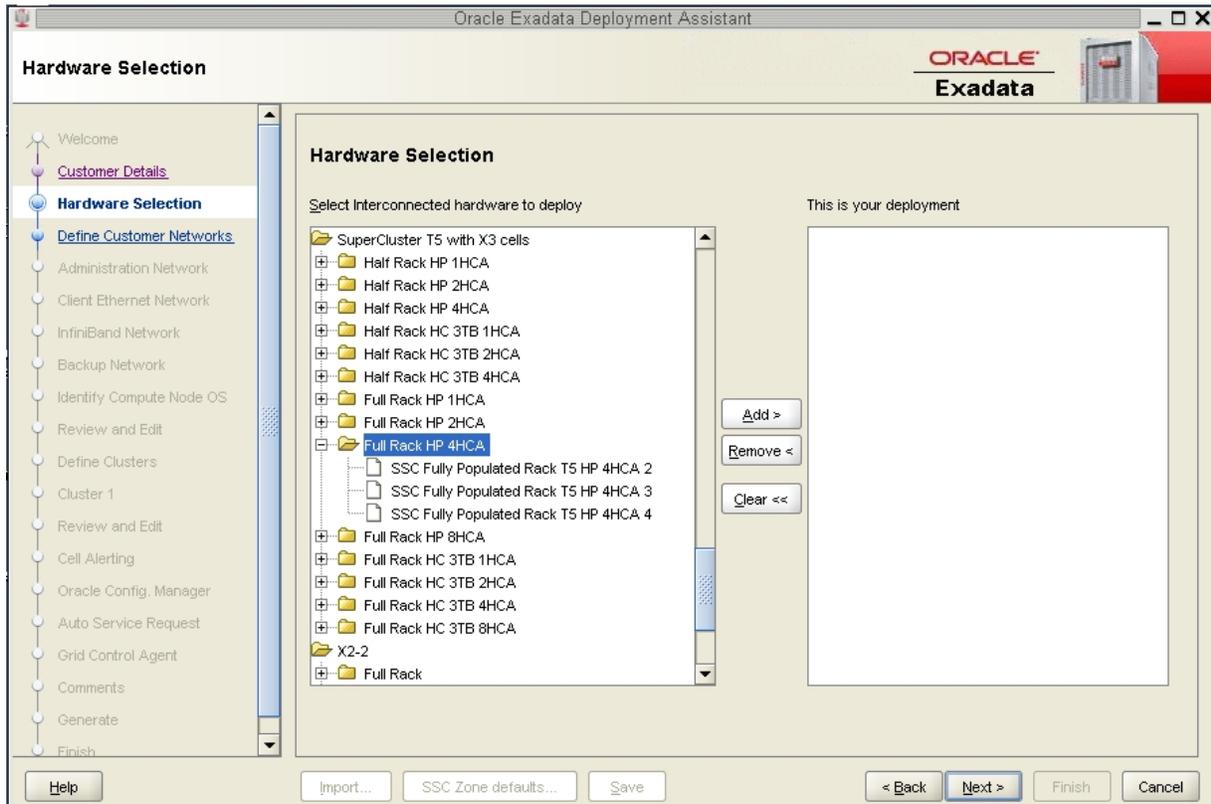
For example, if you are entering information for the following cluster setup at this point:

- A Full Rack version of the Oracle SuperCluster T5-8
- Both SPARC T5-8 servers have the F4-2 domain configuration set up on the servers
- You are setting up zones on the first Database Domain in the F4-2 configuration (the Large Domain)
- There are four IB HCAs associated with this Large Domain
- You have high performance (HP) disks in the Exadata Storage Servers

then you would open the following folder in the Hardware Selection page:

Full Rack HP 4HCA

The following figure shows the selections that would be available for that folder.



#### 4. Make the appropriate selection within the folder that you opened.

The selections provided in this folder reflect the total number of similar Database Domains that could house zones in your system.

For example, within the Full Rack HP 4HCA folder, you would see the following options:

- SSC Fully Populated Rack T5 HP 4HCA 2
- SSC Fully Populated Rack T5 HP 4HCA 3
- SSC Fully Populated Rack T5 HP 4HCA 4

The number at the end of each option denotes the number of Large Domains (domains with 4 IB HCAs) that could be part of this cluster. Therefore:

- If you are creating a cluster that spans across *two* Large Domains, you would select SSC Fully Populated Rack T5 HP 4HCA 2
- If you are creating a cluster that spans across *three* Large Domains, you would select SSC Fully Populated Rack T5 HP 4HCA 3
- If you are creating a cluster that spans across *four* Large Domains, you would select SSC Fully Populated Rack T5 HP 4HCA 4

These are the only choices available in this folder because, for a Full Rack, there is one domain configuration option with two Large Domains on each SPARC T5-8 server, or four Large Domains total in your system (F2-1 on both SPARC T5-8 servers), and several domain configuration options with one Large Domain on each SPARC T5-8 server, or two Large Domains total in your system (F3-1, F4-2, and F5-2 on both SPARC T5-8 servers).

For this example configuration, where you would have the F4-2 domain configuration on both SPARC T5-8 servers, you would have two Large Domains total in your system, so the SSC Fully Populated Rack T5 HP 4HCA 2 option would apply in this example configuration.

5. **When you have selected the proper configuration for your system, click the Add button to add that selection to the This is your deployment window, then click Next.**

The Define Customer Networks page is displayed. Go to [“Complete the Define Customer Networks Page” on page 87](#).

## ▼ Complete the Define Customer Networks Page

1. Complete the fields in the Define Customer Networks page.

**Define Customer Networks**

Oracle Exadata Deployment Assistant

ORACLE  
Exadata

Define Customer Networks

Exadata requires a minimum of 2 separate customer subnets. This page allows you to describe those subnets, for completeness it also includes subnet 3 which is the Private infiniband network. Some customers have more than 2 subnets. In those cases Exadata can configure one of those additional subnets for 'backup', 'replication', 'dr' or for an 'independent client' network in multi cluster environments. This is included here as subnet 4 however this subnet is NOT mandatory for deployment

**Subnet 1**

Name : **Admin**  Bonded  Non Bonded

Subnet Mask : 255.255.255.0

Gateway :

Network Format :  1/10 Gbit Copper Base-T  10 Gbit Optical

**Subnet 2**

Name : **Client**  Bonded  Non Bonded

Subnet Mask : 255.255.255.0

Gateway :

Network Format :  1/10 Gbit Copper Base-T  10 Gbit Optical

**Subnet 3**

Name : **Private**  Bonded  Non Bonded

Subnet Mask : 255.255.252.0

Network Format :  Infiniband

**Subnet 4**

Available Network : Backup  Bonded  Non Bonded

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

- a. In the Subnet 1 (Admin) area, provide the following information:

Field	Entry
Subnet Mask	Choose the subnet mask for this cluster.
Gateway	Type the gateway for this cluster.
Network Format	Leave the 1/10 Gbit Copper Base-T option selected.

- b. In the Subnet 2 (Client) area, provide the following information:

Field	Entry
Bonded or Non Bonded selection	Leave the Bonded option selected.
Subnet Mask	Choose the subnet mask for this cluster.
Gateway	Type the gateway for this cluster.
Network Format	Select the 10 Gbit Optical option.

**c. In the Subnet 3 (Private) area, provide the following information:**

Field	Entry
Subnet Mask	Choose the subnet mask for this cluster.

**2. Click the Next button.**

The Administration Network page is displayed. Go to [“Complete the Administration Network Page” on page 89.](#)

## ▼ Complete the Administration Network Page

1. Complete the fields in the Administration Network page.

Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.
Is the default gateway for database servers	Select this option if you want the administration network to be the default gateway for this cluster. Note that Oracle recommends the client access network for the default gateway.
Defines the Hostname for the database servers	Leave this option selected.

2. In the Sample first host names field, click the Modify... button.

The Admin Network Format Masks page is displayed.

**Admin Network Format Masks**

**Compute Node** \_\_\_\_\_ **Compute ILOM** \_\_\_\_\_

The %% sign in the 'Compute' names will be replaced with the Start ID.

Name :  Name :

Starting ID :  Starting ID :

Sample Name :  Sample Name :

**Cell Node** \_\_\_\_\_ **Cell ILOM** \_\_\_\_\_

The %% sign in the 'Cell' names will be replaced with the Start ID.

Name :  Name :

Starting ID :  Starting ID :

Sample Name :  Sample Name :

**Switches and PDUs** \_\_\_\_\_

The % sign in the 'Switch' and 'PDU' names will be replaced with the Rack ID.

Cisco switch :  InfiniBand Spine :

KVM switch :  InfiniBand Leaf :

PDU-A :  InfiniBand Leaf :

PDU-B :

3. **Modify the host names in this page for the zones in this cluster.**
  - a. **In the Compute Node area, enter a unique name for this set of zones in the Name field in the Compute Node area.**

The Compute Node area will be populated with a name that was automatically generated, based on the Database Machine Name prefix that you entered earlier. However, this automatically-generated name does not provide necessary zone-specific information, so you should change the name in this field to reflect information on the zones in this cluster.

For example, if you entered `osc01r1` for the Database Machine Name prefix in the Customer Details page, you should see the following text in the Name field in the Compute Node area:

```
osc01r1dbadmin%%
```

The `%%` part of the name allows for a number to be automatically generated at the end of this process. You can enter information before or after the `%%` text.

The name that you enter in this field will be used as the host name for the 1GbE host management for zones that are part of this cluster, so you should enter information in this field that distinguishes this set of zones from other zones in other clusters.

The format you use for this field is up to you. The following format is an example of a format that you could use for this field:

```
cluster-namez#db#cn%
```

where:

- *cluster-name* is the Database Machine Prefix that you entered in the Customer Details page.
- *z#* is used for the zone numbers for this cluster. The zones numbers would be the same as the number for the cluster that they reside in. For example, if this is the first cluster that you are configuring (r1), then all the zones in this cluster would be z1 zones.
- *db#* is used for the Database Domain numbers on each SPARC T5-8 server that will house the zones that will be members of this cluster.
- *cn%* is used for the SPARC T5-8 servers (compute nodes) that house the Database Domains that will contain the zones that you are creating. Note that the `%` part of the name allows for a number to be automatically generated at the end of this process.

For example, assume that you want to create a set of zones for a cluster with these characteristics:

- This is the first cluster and set of zones that you will be setting up on your system
- The zones for this cluster will all reside on the second Database Domain on each of the SPARC T5-8 servers
- This will be a two-node cluster (two zones in this cluster), and the zones will reside on the Database Domains in SPARC T5-8 servers 1 and 2 in the Oracle SuperCluster T5-8

Then you could use this as the host name for this cluster and set of zones:

osc01r1z1db2cn%

where:

- osc01r1 is the Database Machine Prefix that you chose earlier
- z1 indicates that these zones are for the first cluster that you will be configuring
- db2 indicates that this set of zones that will be members of this cluster all reside on the second Database Domain on each SPARC T5-8 server
- The cn before the % allows for the compute node numbers to be automatically generated at the end of the process

---

**Note** - There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T5-8 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T5-8 server, and the second zone resides on the second Database Domain on the second SPARC T5-8 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

---

**b. In the Compute ILOM area, change the information in the Name field that was automatically populated.**

The Name field in the Compute ILOM area will automatically populate, depending on the information that you entered in the Name field in the Compute Node area. For example, if you entered osc01db2cn%-z1 in the Name field in the Compute Node area, then the following information will automatically populate in the Name field in the Compute ILOM area:

osc01db2cn%-z1-ilom

However, the Oracle ILOM information applies to the SPARC T5-8 server as a whole, not to a cluster with zones that reside on that server. You must change this field so that the format matches the host name for Oracle ILOM that has already been assigned to your SPARC T5-8 servers, where the %% in the field denotes a numerical wild card that will be automatically assigned.

---

**Note** - For some Oracle SuperCluster T5-8s, the Oracle ILOM host names assigned to the SPARC T5-8 servers use a single-digit number, such as osc01cn1-ilom, instead of a double-digit number, such as osc01cn01-ilom. Each percentage sign (%) wild card generates a single digit, so if the Oracle ILOM host names assigned to your SPARC T5-8 servers use a single-digit number, remove one of the two percentage signs so that host names that are automatically generated use only a single digit.

---

---

For example, if the host names for Oracle ILOM for your SPARC T5-8 servers are *SuperCluster-namecnT4-4-number-ilom* (for example, osc01cn1-ilom), then you would delete the necessary text in this field until you have this:

```
osc01cn%-ilom
```

**c. In the Cell Node area, change the information in the Name field that was automatically populated.**

The Name field in the Cell Node area will automatically populate, depending on the Database Machine Prefix that you chose earlier. For example, if you used osc01r1 as the Database Machine Prefix, then the following information will automatically populate in the Name field in the Cell Node area:

```
osc01r1celadmin%%
```

However, this field is meant for the Exadata Storage Servers, so zone naming information should not be used in this area. You must change this field so that the format matches the host name for the 1GbE host management that has already been assigned to your Exadata Storage Servers, where the %% in the field denotes a numerical wild card that will be automatically assigned.

---

**Note** - For some Oracle SuperCluster T5-8s, the 1GbE host management host names assigned to the Exadata Storage Servers use a single-digit number, such as osc01cel**1**, instead of a double-digit number, such as osc01cel**01**. Each percentage sign (%) wild card generates a single digit, so if the 1GbE host management host names assigned to your Exadata Storage Servers use a single-digit number, remove one of the two percentage signs so that host names that are automatically generated use only a single digit.

---

For example, if the host names for the 1GbE host management for your Exadata Storage Servers are *SuperCluster-namecelESS-number* (for example, osc01cel1), then you would delete the necessary text in this field until you have this:

```
osc01cel%
```

Note that the Name field in the Cell ILOM field automatically populates with the changes you make in the Name field in the Cell Node area. Verify that the format in the Cell ILOM field matches the format for the host names for Oracle ILOM for the Exadata Storage Servers.

---

**Note** - Each field has to end with at least once percentage (%) sign, where each percentage sign is used to generate a single numerical digit at the end of the host name. If the host name for a component does not end with a numeral, you will still have to leave a percentage sign at the end of the host name, and you will edit the host names for these components to remove the numerical digit later on in these procedures.

---

**d. Change all information for the switches and PDUs.**

In the Switches and PDUs area, manually change all of the information in the beginning of each field, before the sw, so that the information matches the original name that was assigned to the Oracle SuperCluster T5-8 at the initial deployment. Doing so will create generated names that will match the host names for the 1GbE host management that were originally assigned to each switch and PDU.

For example, if the original name of the Oracle SuperCluster T5-8 was osc01, then you would enter the following information for the switches and PDUs:

- Cisco switch: osc01sw-admin%
- KVM switch: osc01sw-kvm%
- PDU-A: osc01sw-pdua%
- PDU-B: osc01sw-pdub%
- IB switch: osc01sw-ibs%
- IB Leaf: osc01sw-iba%
- IB Leaf: osc01sw-ibb%

The following table provides information on how the fields in this screen should be populated once you have made all the necessary changes.

Component/Cluster	Correct Entry
Compute Node Name and Starting ID	Should have a value that applies to the zones in this particular cluster
Compute ILOM Name and Starting ID	Should have a value that applies to the SPARC T5-8 servers
Cell Node Name and Starting ID	Should have a value that applies to the Exadata Storage Servers
Cell ILOM Name and Starting ID	Should have a value that applies to the Exadata Storage Servers
Switches and PDUs Names	Should have values that apply to the switches and PDUs

**e. Click the Save button to exit the Admin Network Format Masks page.**

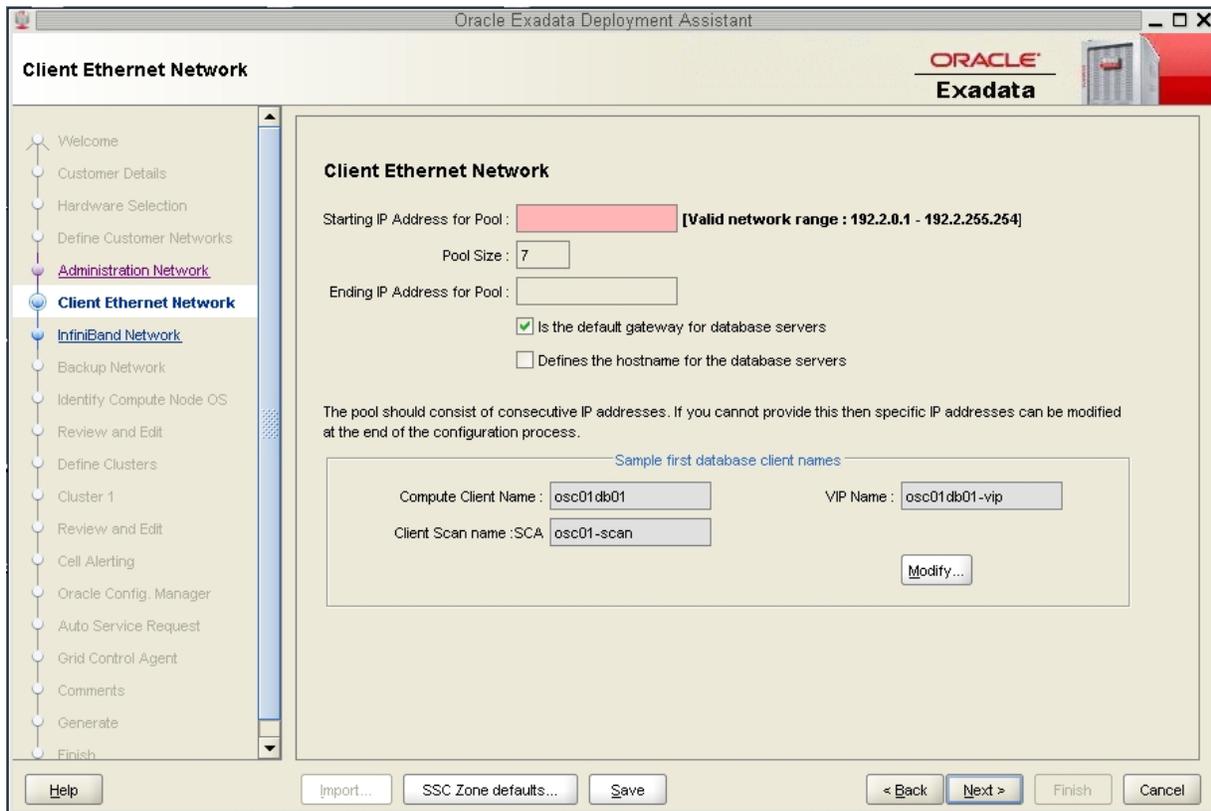
The main Administration Network page is displayed, showing the results of the changes you made in the Admin Network Format Masks page.

**4. In the main Administration Network page, click Next.**

The Client Ethernet Network page is displayed. Go to [“Complete the Client Ethernet Network Page”](#) on page 95.

## ▼ Complete the Client Ethernet Network Page

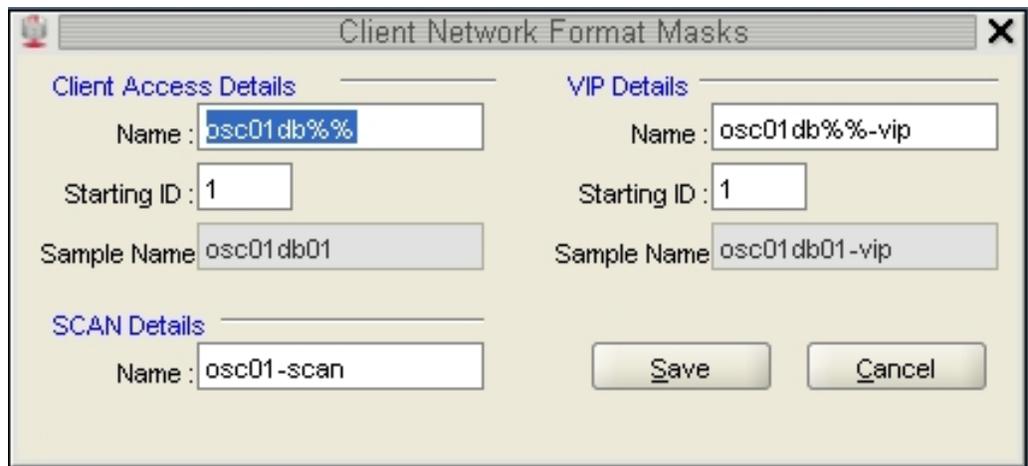
**1. Complete the fields in the Client Ethernet Network page.**



Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.
Is the default gateway for database servers	Select this option if you want the client Ethernet network to be the default gateway for the cluster.
Defines the Hostname for the database servers	Select this option.

2. **In the Sample first database client names field, click the Modify... button.**

The Client Network Format Masks page is displayed.



3. **Modify the host names in this page for the zones in this cluster.**

- a. **In the Client Access Details area, enter a unique name for this set of zones in the Name field.**

The Client Access Details area will be populated with a name that was automatically generated, based on the Database Machine Name prefix that you entered earlier. However, this automatically-generated name does not provide necessary zone-specific information, so you should change the name in this field to reflect information on this cluster.

For example, if you entered osc01r1 for the Database Machine prefix in the Customer Details page, you should see this entry in the Name field in the Client Access Details area:

osc01r1client%%

The name that you enter in this field will be used as the host name for the client access network for zones that are part of this cluster. You should use the same naming convention for this field as you did for the administration network page. The following is an example format for this field:

*cluster-namez#db#cn%-z#-client*

For example:

*osc01r1z1db2cn%-client*

---

**Note** - There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T5-8 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T5-8 server, and the second zone resides on the second Database Domain on the second SPARC T5-8 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

---

Note that the Name field in the VIP Details area automatically populates with the changes that you made in the Client Access Name field in the Client Access Details area. Because this name applies to each zone within this particular cluster, this automatically-generated information should be correct as-is.

Do not modify the information in the SCAN Details section. This will generate a single SCAN name and three SCAN IP addresses for the cluster or the single-instance database, which is correct.

**b. Click the Save button to exit the Client Network Format Masks page.**

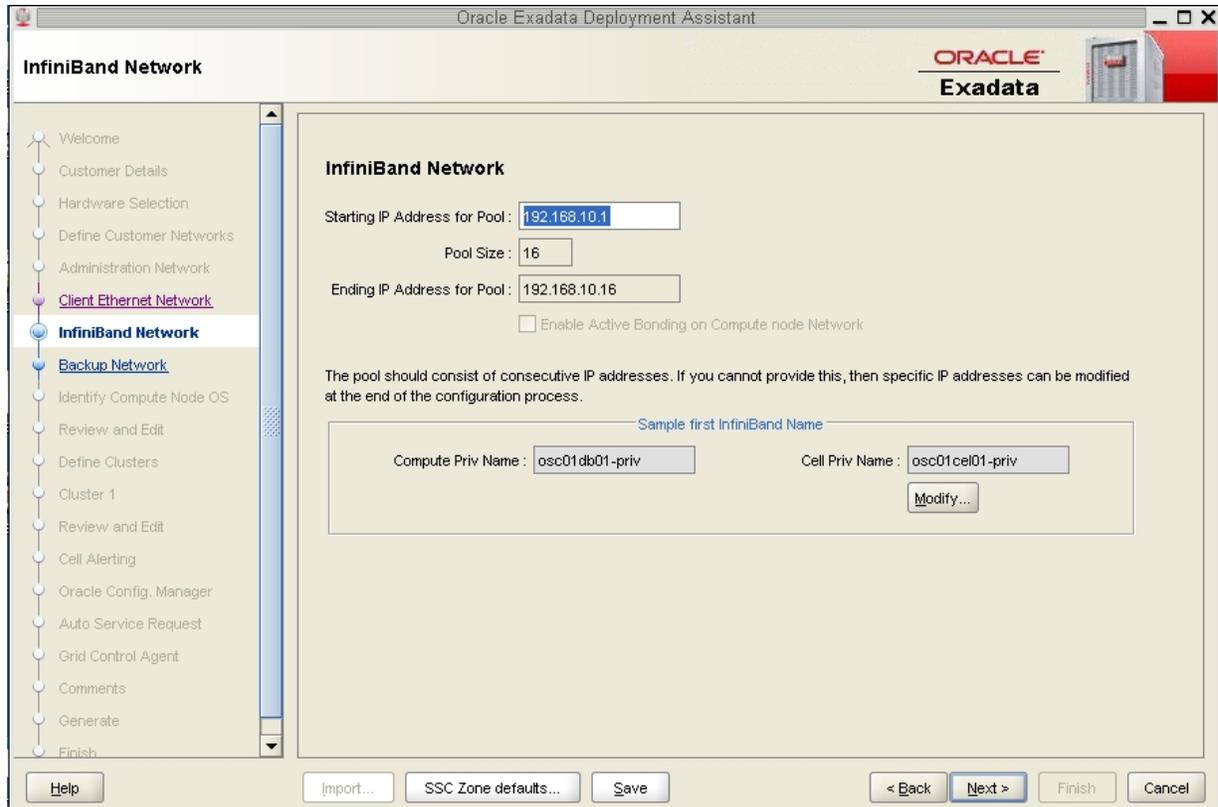
The main Client Access Network page is displayed, showing the results of the changes you made in the Client Network Format Masks page.

**4. In the main Client Ethernet Network page, click Next.**

The IB Network page is displayed. Go to [“Complete the IB Network Page”](#) on page 98.

## ▼ Complete the IB Network Page

1. Complete the fields in the IB Network page.



Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.

2. In the Sample first IB Name field, click the Modify... button.

The Private (IB) Network Format Masks page is displayed.

**3. Modify the host names in this page for the zones in this cluster.**

- a. In the Compute Private Details area, modify the information in the Name field that was automatically populated.**

The Compute Private Details area will be populated with a name that was automatically generated, based on the Database Machine Name prefix that you entered earlier. However, this automatically-generated name does not provide necessary zone-specific information, so you should change the name in this field to reflect information on this cluster.

For example, if you entered `osc01r1` for the Database Machine prefix in the Customer Details page, you should see this entry in the Name field in the Compute Private Details area:

```
osc01r1db%%-priv
```

The name that you enter in this field will be used as the host name for the IB network for zones that are part of this cluster. You should use the same naming convention for this field as you did for the administration network page. The following is an example format for this field:

```
cluster-namez#db#cn%-priv
```

For example:

```
osc01r1z1db2cn%-priv
```

---

**Note** - There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T5-8 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T5-8 server, and the second zone resides on the second Database Domain on the second SPARC T5-8 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

---

**b. In the Cell Private Details area, modify the information in the Name field that was automatically populated.**

The Name field in the Cell Private Details area will automatically populate, depending on the Database Machine Prefix that you chose earlier. For example, if you used osc01r1 as the Database Machine Prefix, then the following information will automatically populate in the Name field in the Cell Private Details area:

```
osc01r1cel%-priv
```

However, this field is meant for the Exadata Storage Servers, so zone naming information should not be used in this area. You must change this field so that the format matches the host name for the IB network that has already been assigned to your Exadata Storage Servers, where the % in the field denotes a numerical wild card that will be automatically assigned.

For example, if the host names for the IB network for your Exadata Storage Servers are *supercluster-namecelESS-number-priv* (for example, osc01cel01-priv), then you would delete the necessary text in this field until you have this:

```
osc01cel%-priv
```

The following table provides information on how the fields in this screen should be populated once you have made all the necessary changes.

Component/Cluster	Correct Entry
Compute Private Details Name and Starting ID	Should have a value that applies to the zones in this particular cluster
Cell Private Details Name and Starting ID	Should have a value that applies to the Exadata Storage Servers

**c. Click the Save button to exit the Private (IB) Network Format Masks page.**

The main IB Network page is displayed, showing the results of the changes you made in the Private (IB) Network Format Masks page.

**4. In the main IB Network page, click Next.**

The Backup Network page is displayed.

**5. Complete the fields in the optional Backup Network page, if necessary.**

When you have completed this page, click Next. The Identify Compute Node OS and Enable Capacity-on-Demand Configuration page is displayed.

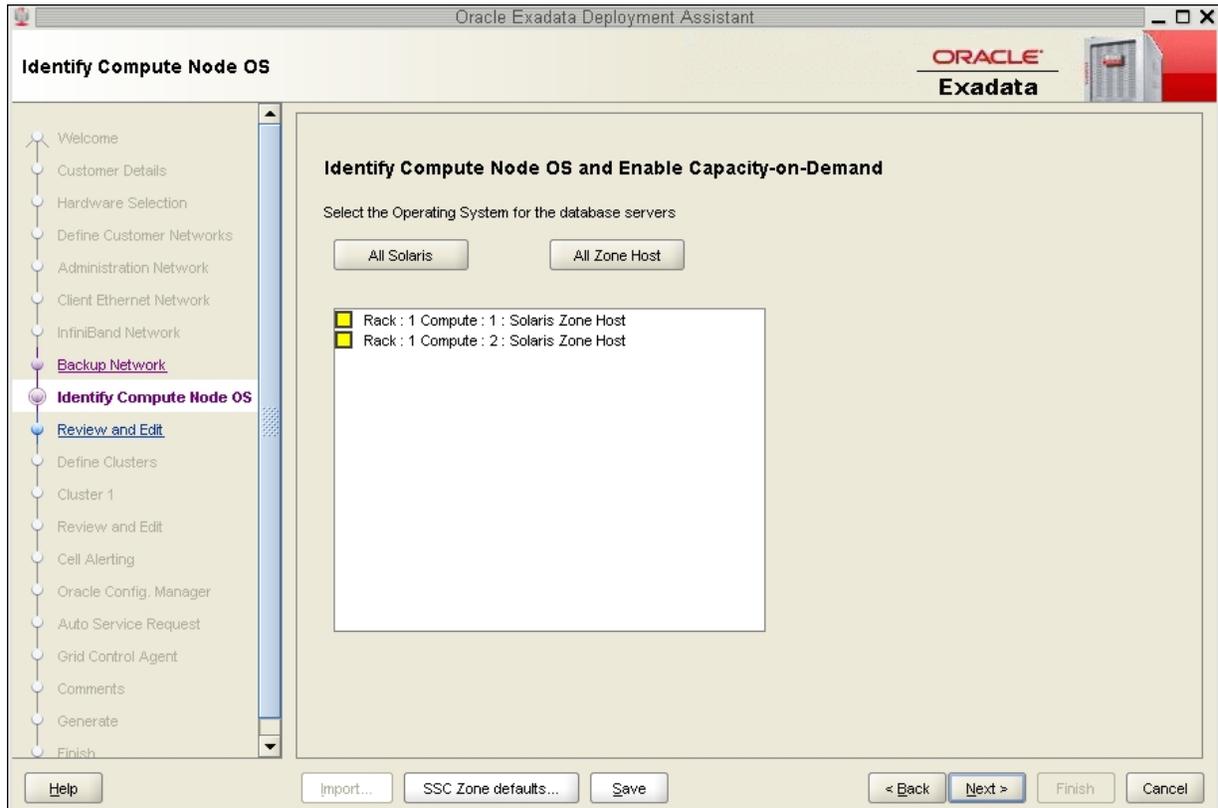
**6. Complete the fields in the Identify Compute Node OS and Enable Capacity-on-Demand Configuration page.**

Go to [“Complete the Identify Compute Node OS and Enable Capacity-on-Demand Configuration Page”](#) on page 101.

**▼ Complete the Identify Compute Node OS and Enable Capacity-on-Demand Configuration Page**

1. Click All Zone Host to select the Oracle Solaris OS for each zone in this cluster.

The selections change to Solaris Zone Host for all zones in the cluster.



**2. Click Next.**

The Management and Private Networks page is displayed. Go to [“Review the Information in the Management and Private Networks Page”](#) on page 102.

▼ **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)
- Exadata Storage Servers in the Oracle SuperCluster T5-8
- Switches and PDUs in the Oracle SuperCluster T5-8

The screenshot shows the Oracle Exadata Deployment Assistant (OEDA) interface. The title bar reads 'Oracle Exadata Deployment Assistant'. The main window is titled 'Review and Edit' and features the Oracle Exadata logo in the top right corner. On the left, a vertical navigation pane lists various steps: Welcome, Customer Details, Hardware Selection, Define Customer Networks, Administration Network, Client Ethernet Network, InfiniBand Network, Backup Network, Identify Compute Node OS (highlighted in purple), Review and Edit (highlighted in blue), Define Clusters, Cluster 1, Review and Edit, Cell Alerting, Oracle Config. Manager, Auto Service Request, Grid Control Agent, Comments, Generate, and Finish. The main content area is titled 'Management and Private Networks' and contains the following text: '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.' Below this text is a 'Re-Generate Data' button. The page displays configuration details for four components, each with a 'Rack 1 - Rack Location' label and three input fields for 'Admin Name', 'Admin IP', 'IloM Name', 'IloM IP', 'Priv Name', and 'Priv IP':

Component	Rack 1 - Rack Location	Admin Name	Admin IP	IloM Name	IloM IP	Priv Name	Priv IP
Exadata SuperCluster T5 Compute Node 1 - Zone Host	10	osc01dbadm01.acme.com	192.1.1.2	osc01dbadm01-iloM.acme.com	192.1.1.12	osc01db01-priv1.acme.com	192.168.10.1
Exadata SuperCluster T5 Compute Node 2 - Zone Host	27	osc01dbadm02.acme.com	192.1.1.3	osc01dbadm02-iloM.acme.com	192.1.1.13	osc01db02-priv1.acme.com	192.168.10.5
Exadata Cell Node HP 1	2	osc01celadm01.acme.com	192.1.1.4	osc01celadm01-iloM.acme.com	192.1.1.14	osc01cel01-priv.acme.com	192.168.10.9
Exadata Cell Node HP 2	4	osc01celadm02.acme.com	192.1.1.5	osc01celadm02-iloM.acme.com	192.1.1.15	osc01cel02-priv.acme.com	192.168.10.10

At the bottom of the window, there are several buttons: 'Help', 'Import...', 'SSC Zone defaults...', 'Save', '< Back', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted in blue.

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

The Define Clusters page is displayed.

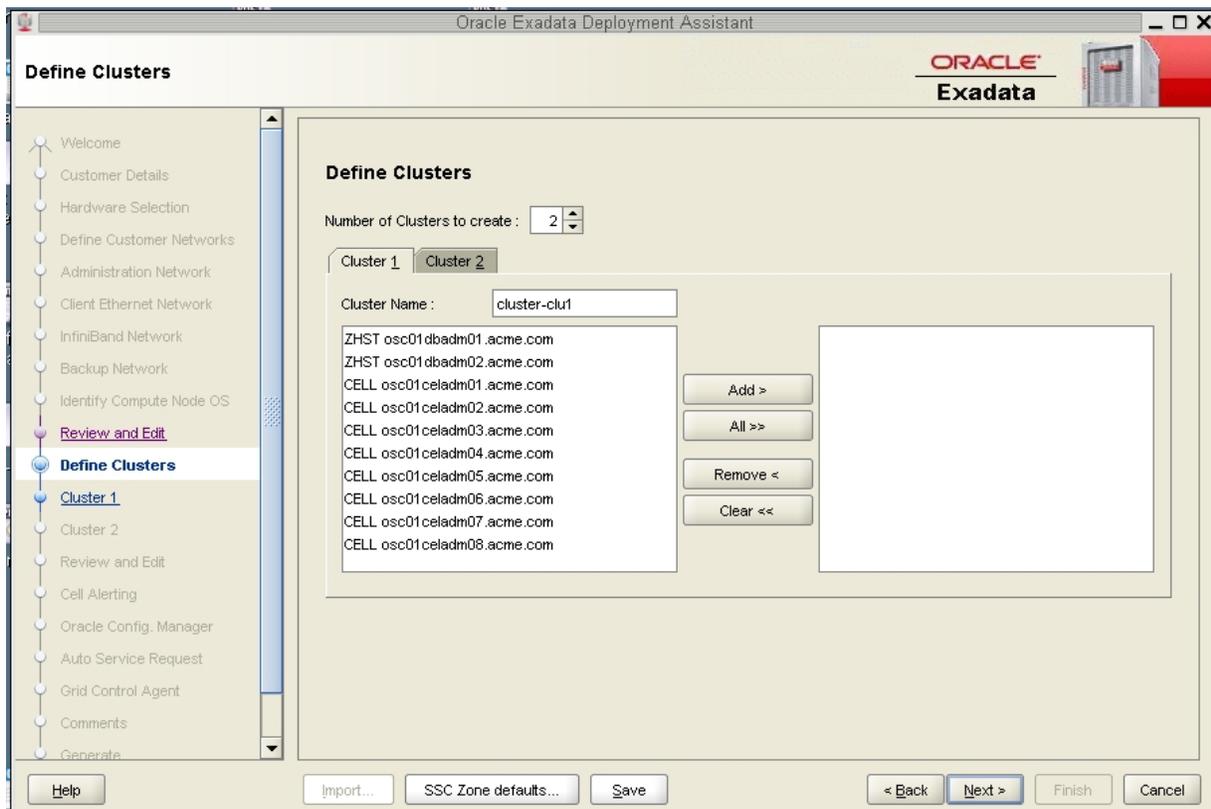
3. Complete the Define Clusters page.

Go to [“Complete the Define Clusters Page”](#) on page 104.

## ▼ Complete the Define Clusters Page

1. **Determine the number of clusters that you will create and select the appropriate number in the Number of Clusters to create area.**
2. **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 1. Note that cluster was automatically named `cluster-clu1` 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-clu1`). 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-clu1`).

**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 dictates that you should assign all Exadata Storage Servers to each cluster rather than split the Exadata Storage Servers between the clusters.

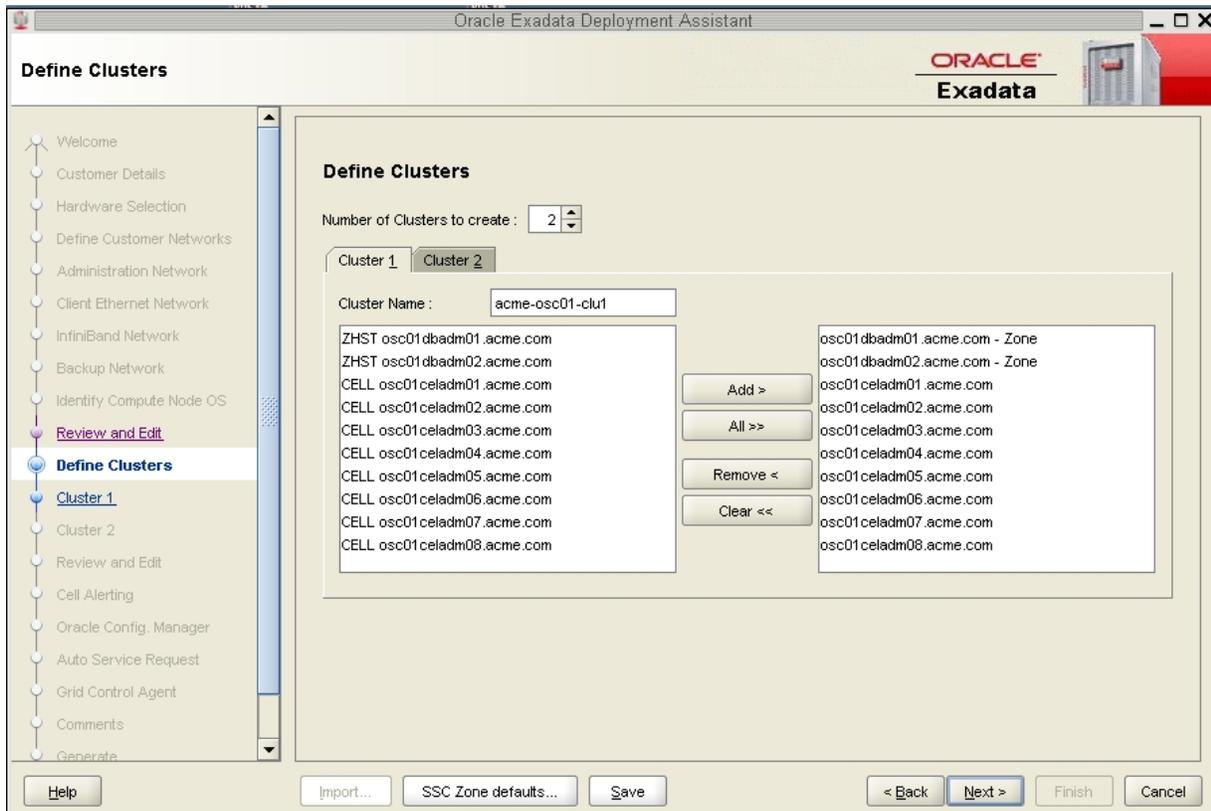
---

---

**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 Exadata Storage Server Disk Space for Additional Zones” on page 74](#) for more information.

---

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



3. Repeat [Step 2](#) to all of the new clusters that you will be creating to assign a unique name and select the Exadata Storage Servers for each cluster.
4. 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.
5. Set the zone default configurations for each cluster that you want to create.  
Go to [“Set Zone Default Configurations”](#) on page 108.

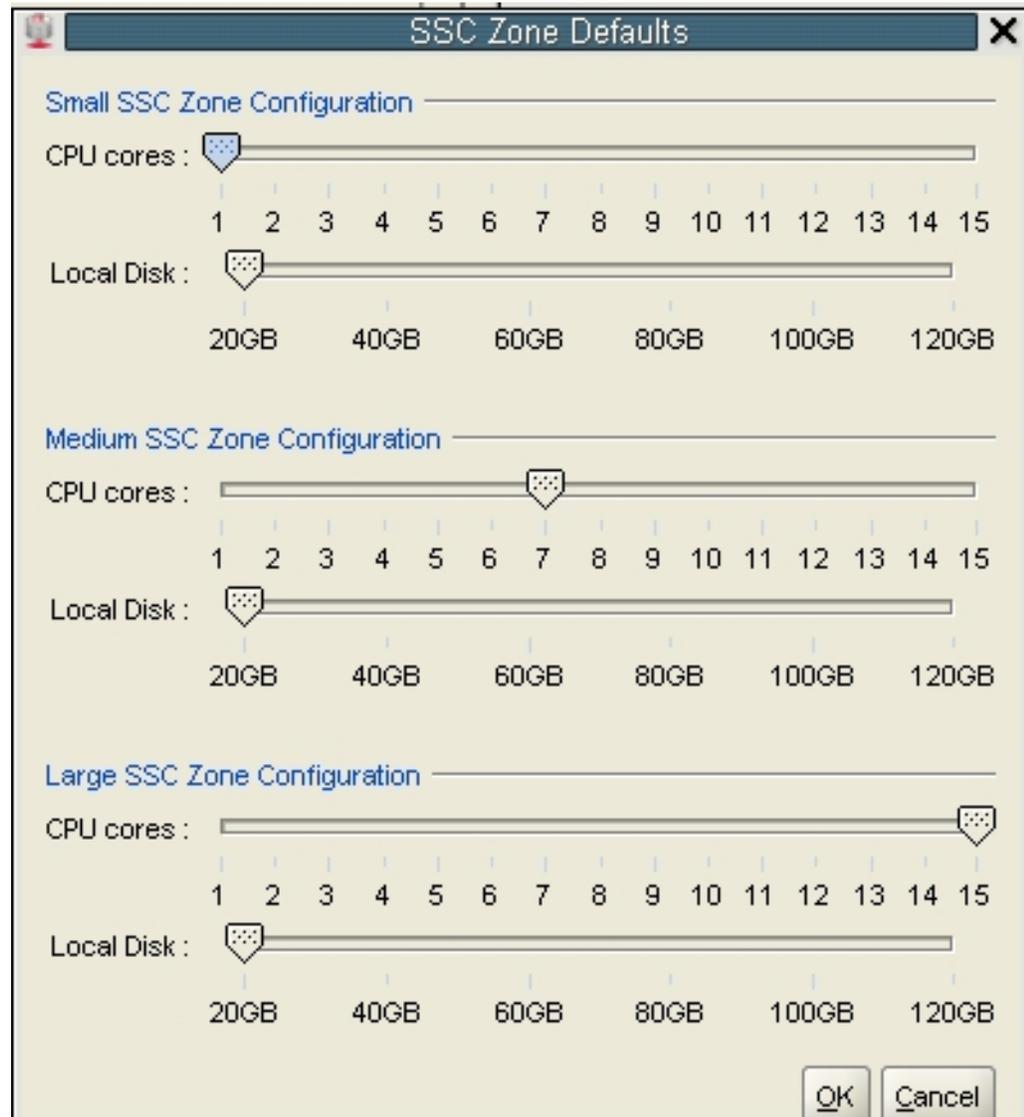
## ▼ Set Zone Default Configurations

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:
  - 15 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 15 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 Domains and Zones” on page 16](#) 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 111](#). The values you assign to each of these types of zone configurations

are completely user-definable, so you can 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:

- 10 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 111](#).

## ▼ 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.

**Cluster 1**

Cluster name:  Zones Cluster

Virtual Guest size:  Cores:  Local Disk:

Prefix:

DNS:

NTP:

Domain Name:

Region:  TimeZone:

**Users and Groups**

Role Separated

User name:  ID:  base:

DBA Group name:  ID:

OINSTALL Group name:  ID:

**Software Locations**

Inventory Location:

Grid Infrastructure Home:

Database Home Location:

Software Languages:

**Disk Group Details**

Diskgroup Layout:  Legacy 80%:20%  Legacy 40%:60%  Custom

DBFS DiskGroup:   Size:

DATA DiskGroup:   Size:

RECO DiskGroup:   Size:

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

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 Domains and Zones” on page 16](#) 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 108](#).

---

**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 will not update to the values that you assigned in the section [“Set Zone Default Configurations” on page 108](#) 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 1:
  - DBFS Diskgroup: DBFS\_DGC1
  - DATA Diskgroup: DATAC1
  - RECO Diskgroup: RECO1
- 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 Exadata Storage Server Disk Space for Additional Zones” on page 74](#) 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 will be 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. So 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 is 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 115.

## ▼ **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 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 117.](#)

## ▼ 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.**

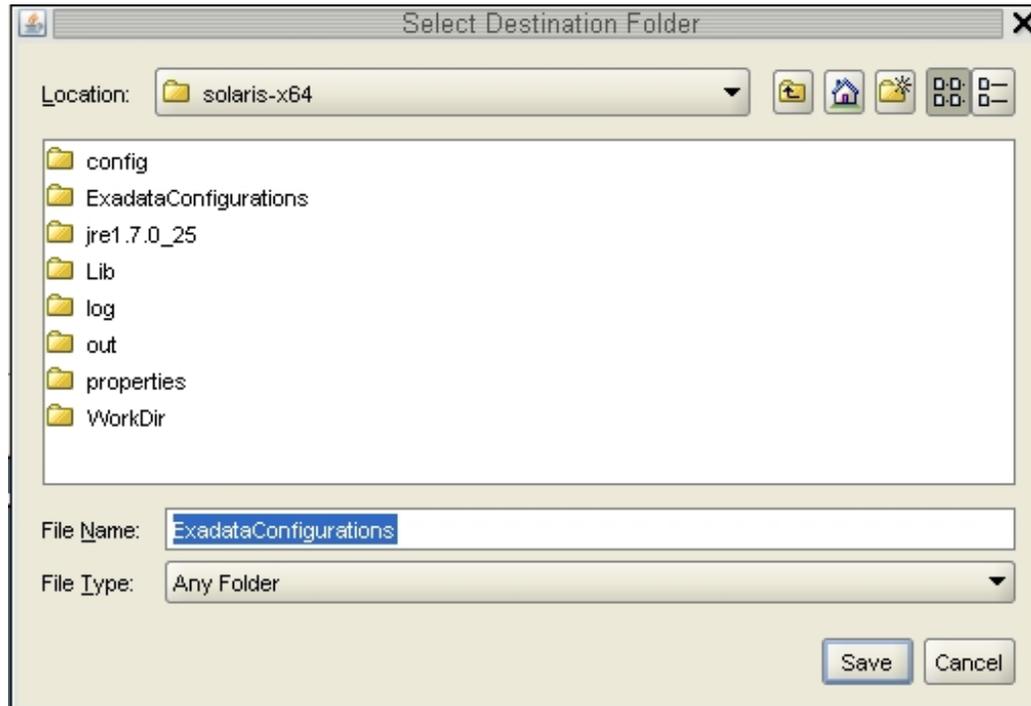
**3. Generate the configuration files.**

Go to [“Generate the Configuration Files” on page 117.](#)

## ▼ 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:

- 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 73](#), 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 zones that were set up on Database Domains with zones, such as:
  - *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 zone
- Individual configuration files for the clusters that were set up, such as:
  - *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)
- A single master configuration file that contains all the configuration information for your system, such as:
  - *Acme-osc01.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. Determine whether you want to create more clusters at this time.

- If you want to create more clusters at this time, now that you have gone through one complete pass at entering information with the new OEDA, you can now use the more automated procedures provided in [“Creating Configuration Files Using OEDA: Build Level](#)

[1.5.8 or Later](#)” on page 120, even though your system is technically at build level 1.5.4 or earlier.

- If you do not want to create more clusters at this time, go to [“Creating the Template Zone on Each Database Domain”](#) on page 149.

## Creating Configuration Files Using OEDA: Build Level 1.5.8 or Later



---

**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.

---

---

**Note** - The procedures in this section apply if you have a software build level of 1.5.8 or later on your system, as determined in the section [“Determine Software Build Level”](#) on page 75. If your software build level is 1.5.4 or earlier, follow the procedures in [“Creating Configuration Files Using OEDA: Build Level 1.5.4 or Earlier”](#) on page 76.

---

Follow these procedures to use OEDA to create your configuration files for a system with a software build of 1.5.8 or later:

- [“Locate the Necessary Files to Set Up Zones”](#) on page 121
- [“Set Up the OEDA”](#) on page 122
- [“Start the OEDA”](#) on page 123
- [“Import the Most Recent OEDA Configuration File”](#) on page 124
- [“Review Existing Configuration Information”](#) on page 126
- [“Review the Information in the Identify Compute Node Operating System Page”](#) on page 127

- [“Review the Information in the Management and Private Networks Page” on page 129](#)
- [“Complete the Define Clusters Page” on page 131](#)
- [“Set Zone Default Configurations” on page 137](#)
- [“Complete the Cluster Page” on page 140](#)
- [“Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page” on page 144](#)
- [“Verify Remaining Configuration Information” on page 146](#)
- [“Generate the Configuration Files” on page 146](#)

## ▼ 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 73](#), 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 zone

And 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

### 1. Prepare to configure zones.

See [“Preparing to Configure Zones on Database Domains” on page 59](#).

### 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 70](#) 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 123](#).

## ▼ 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 124.

**▼ 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 121](#) 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:

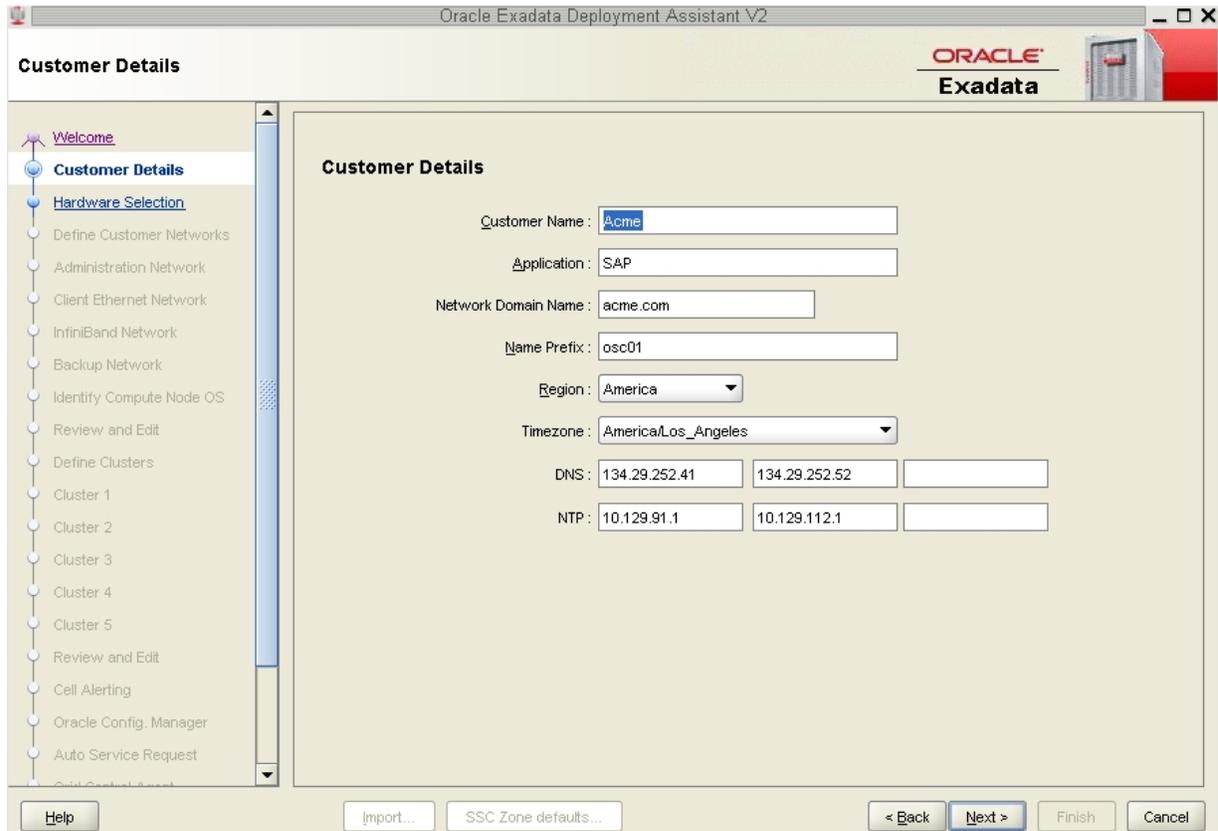
- **Customer name:** Acme
- **Name prefix:** osc01

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 126.

**▼ 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 T5-8. 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
- IB 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 127.](#)

## ▼ 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.

---

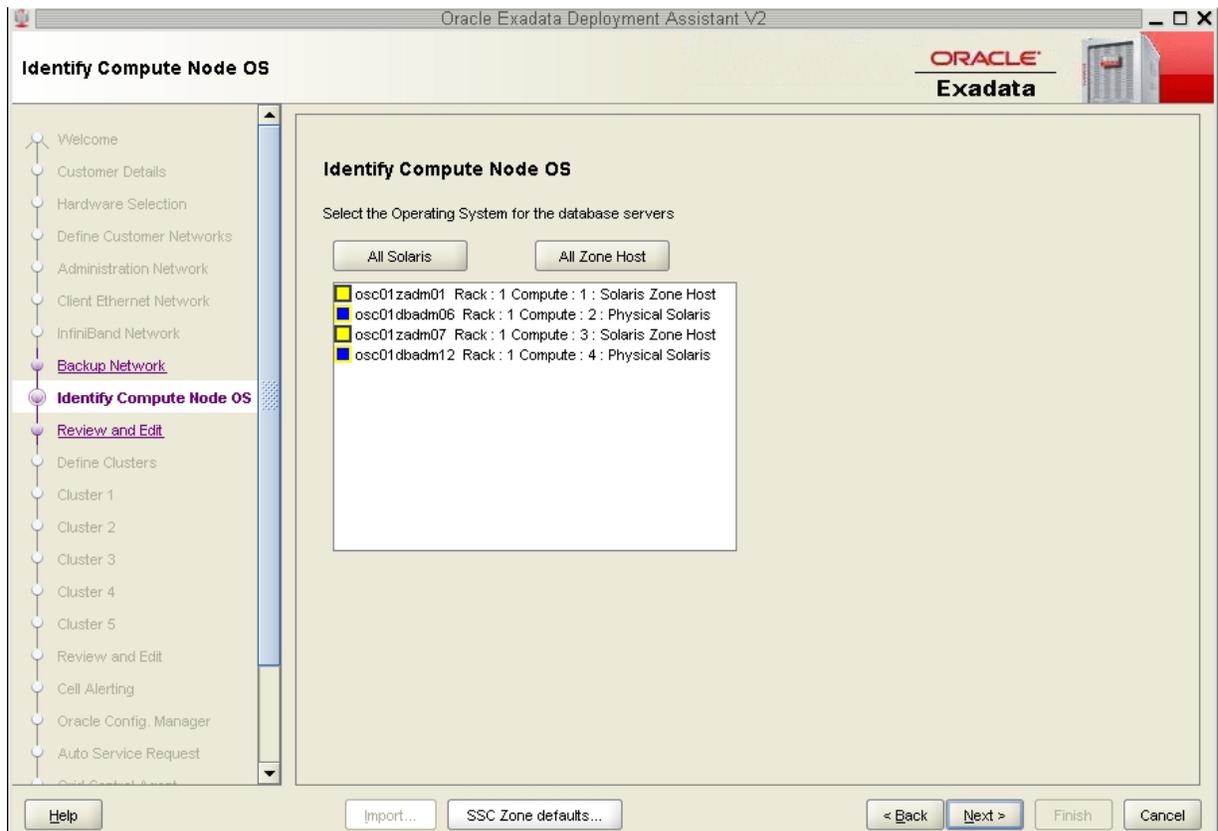
The Identify Compute Node Operating System page should display the total number of Database Domains with zones and Database Domains (without zones) on your Oracle SuperCluster T5-8 in your existing configuration. This page does not display information on any Applications Domains in your existing configuration.

For example, assume you have these LDoms as part of your existing configuration, as described in [“Example Configuration” on page 73:](#)

- One Database Domain with zones on each compute node (2 total)

- One Database Domain (without zones) on each compute node (2 total)
- One Oracle Solaris 10 Application Domain on each compute node (2 total, not displayed in this page)
- One Oracle Solaris 11 Application Domain on each compute node (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 129.

## ▼ **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 T5-8

- Switches and PDUs in the Oracle SuperCluster T5-8

**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

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

- Click Next after reviewing the information in this page.

The Define Clusters page is displayed.

- Complete the Define Clusters page.

Go to [“Complete the Define Clusters Page”](#) on page 131.

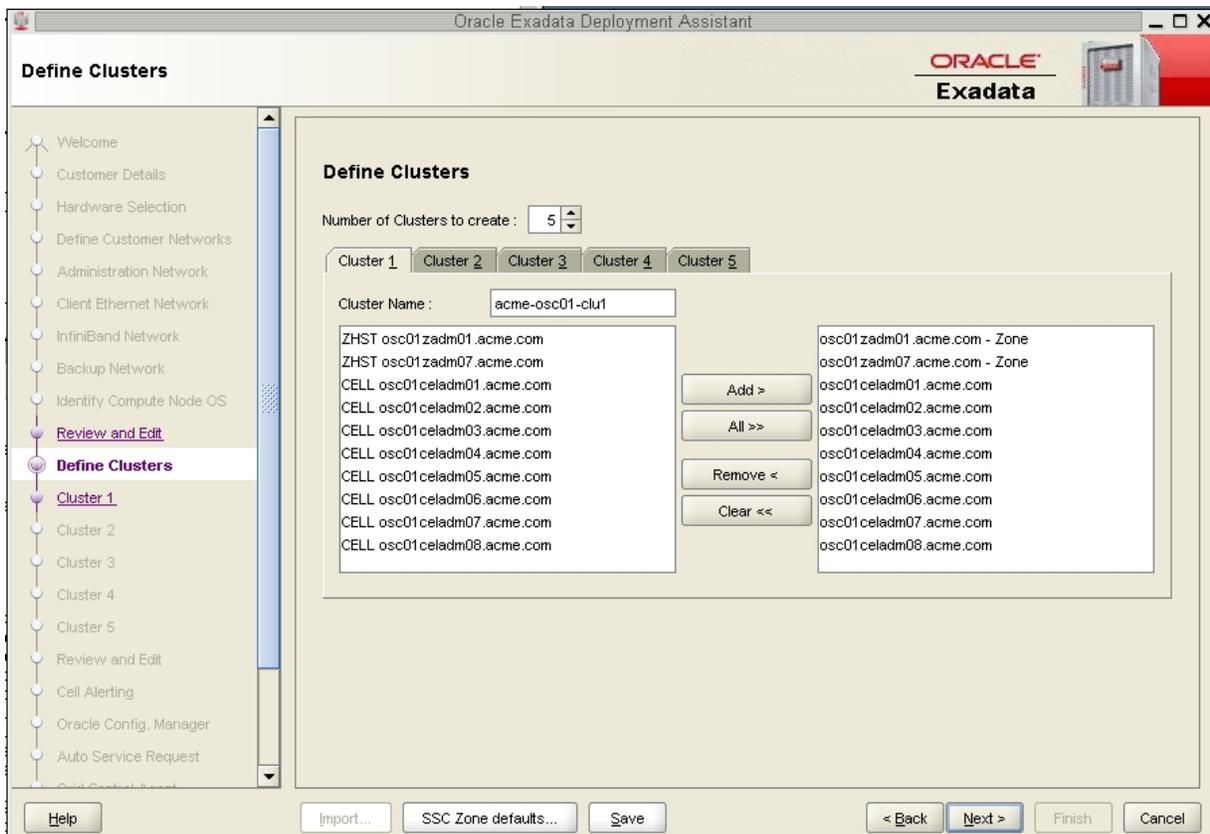
## ▼ Complete the Define Clusters Page

### 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 73](#):

- 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:

- Clusters 1 to 4 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)
- Cluster 5 should display the cluster that was set up for the Database Domains (without zones)

For Clusters 1 to 4, 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 1 to 4 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.

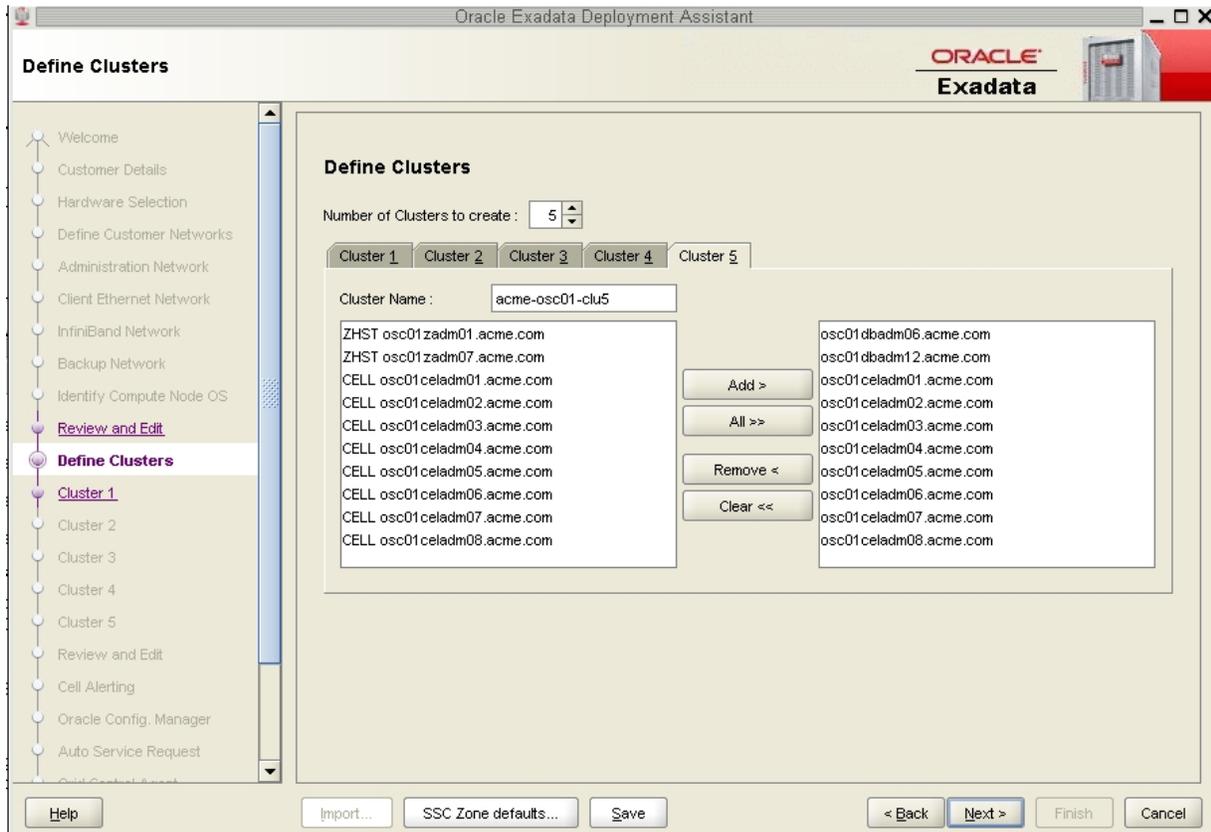
Because Cluster 5 is for the two Database Domains (without zones), that screen would show those two Database Domains and not the Database Domains with zones that are shown in Clusters 1 to 4.

In this example configuration, the screen for Cluster 5 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 73](#), 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.**

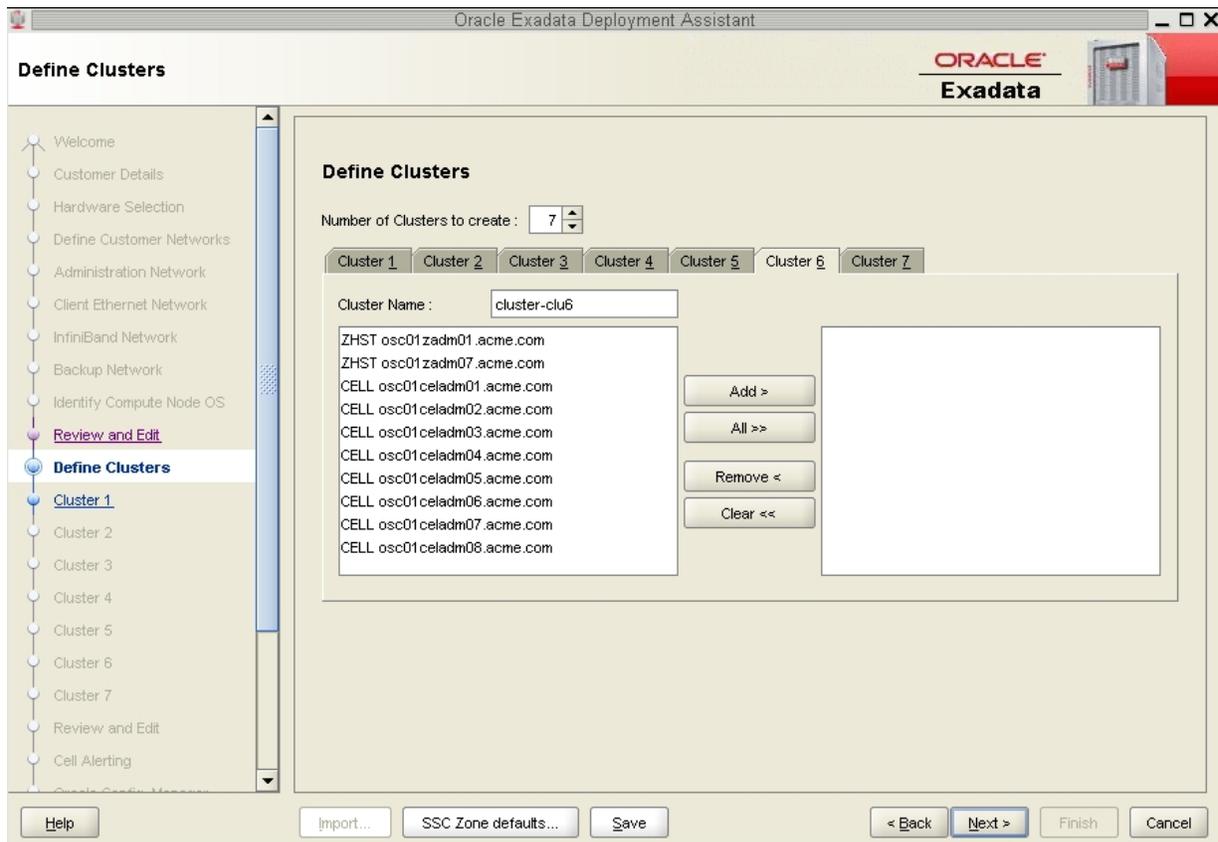
For example, if you are adding two new clusters, you would increase the number of clusters to seven total clusters:

- 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 the 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.

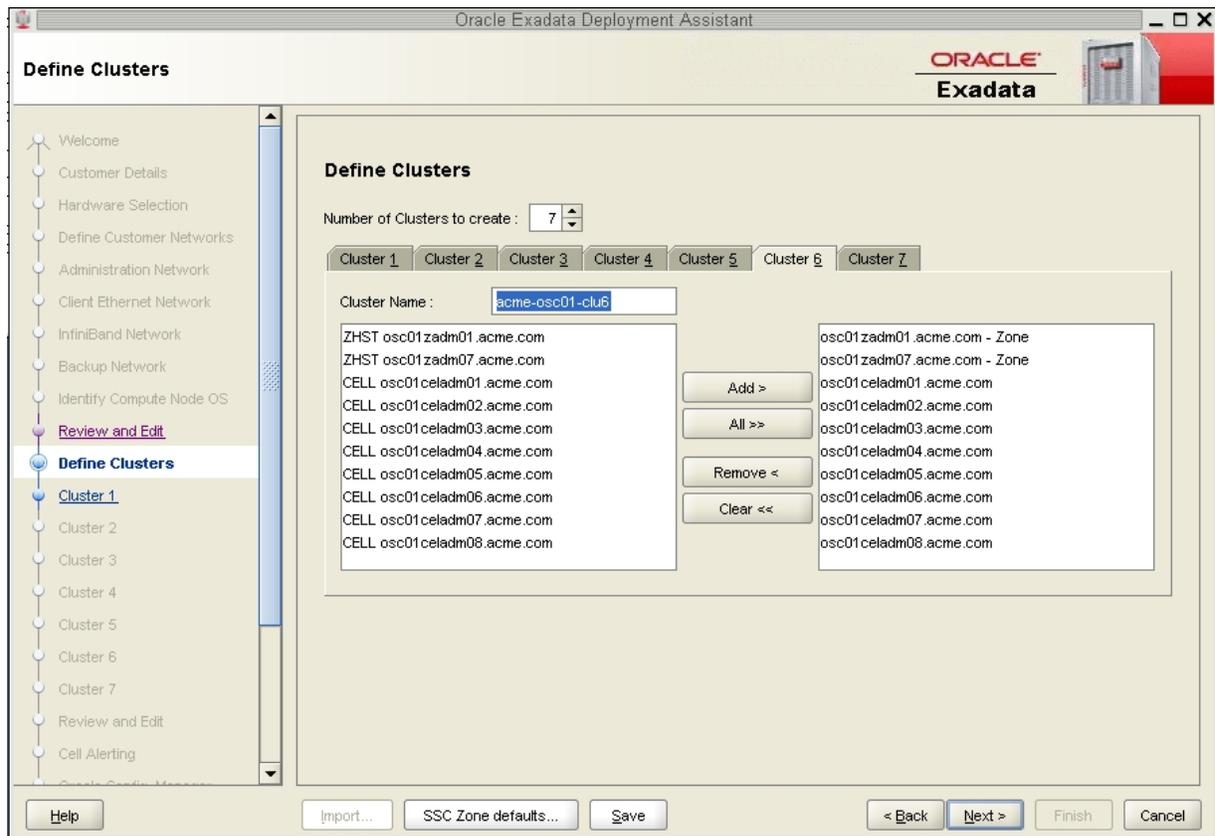
---

---

**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 Exadata Storage Server Disk Space for Additional Zones” on page 74](#) 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 137.

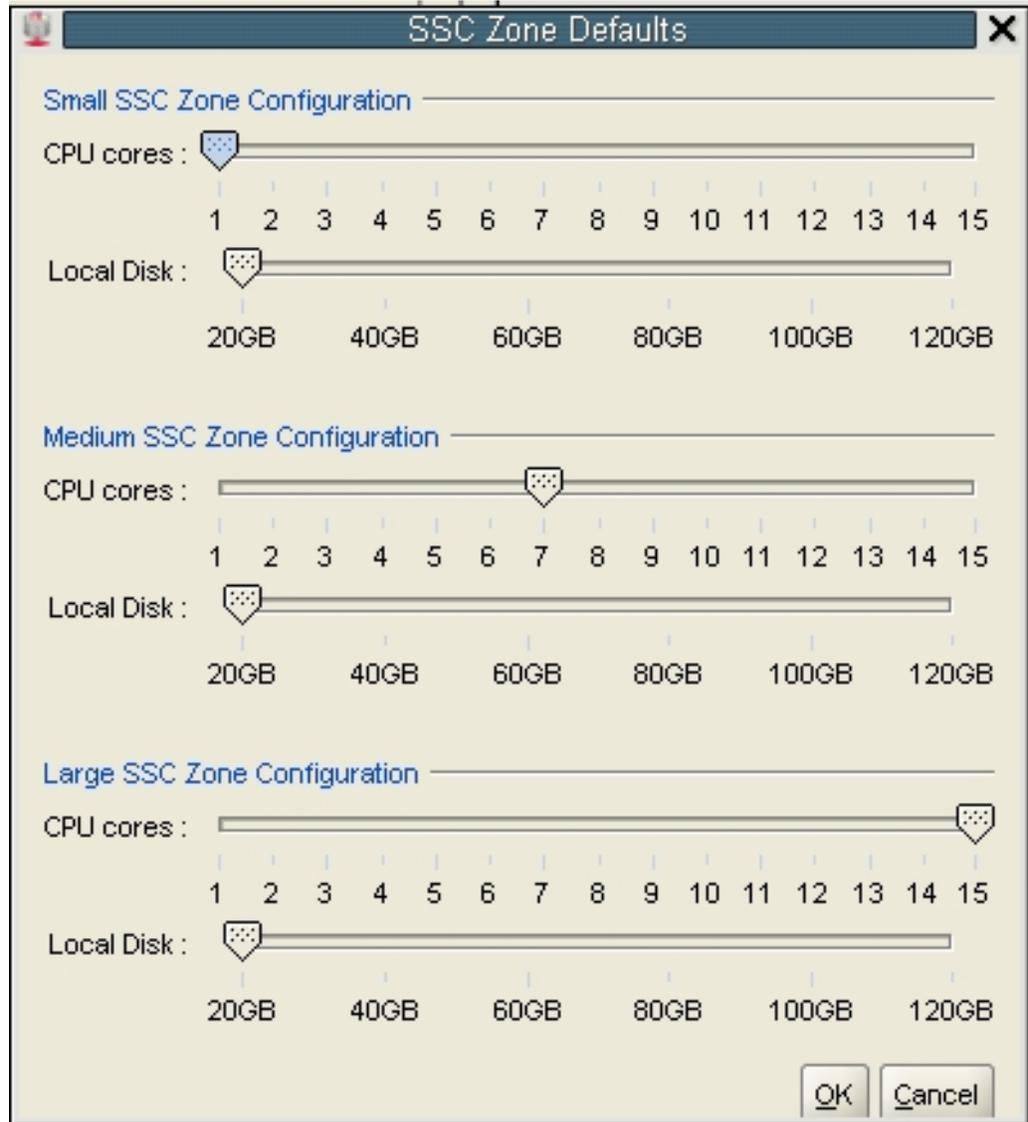
## ▼ Set Zone Default Configurations

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:
  - 15 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 15 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 Domains and Zones” on page 16](#) 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 140](#). 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:

- 10 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 140](#).

## ▼ 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 73](#), review the cluster configuration page for Cluster 6, which is the first new cluster that you are creating. Skip over the cluster configuration 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.

The screenshot shows the Oracle Exadata Deployment Assistant (OEDA) window for configuring Cluster 6. The interface includes a navigation pane on the left with steps from Welcome to Finish. The main area is titled 'Cluster 6' and contains the following configuration fields:

- Cluster name:** me-osc01-clu6 (Zones Cluster)
- Virtual Guest size:** Small (Cores: 4, Local Disk: 20GB)
- Prefix:** osc01
- DNS:** 192.1.1.1
- NTP:** 192.2.1.1
- Domain Name:** acme.com
- Region:** America (TimeZone: Los\_Angeles)
- Users and Groups:**
  - Role Separated
  - User name:** oracle (ID: 1001, base: /u01/app/oracle)
  - DBA Group name:** dba (ID: 1002)
  - OINSTALL Group name:** oinstall (ID: 1001)
- Software Locations:**
  - Inventory Location:** /u01/app/orainventory
  - Grid Infrastructure Home:** 11.2.0.4 BP8 (Path: /u01/app/11.2.0.4/grid)
  - Database Home Location:** 11.2.0.4 BP8 (Path: /u01/app/oracle/product/11.2.0.4/dbhome\_1)
  - Software Languages:** en
- Disk Group Details:**
  - Diskgroup Layout:** Legacy 80%:20% (selected), Legacy 40%:60%, Custom
  - DBFS DiskGroup:** DBFS\_DG (NORMAL, Size: default)
  - DATA DiskGroup:** DATA06 (HIGH, Size: 80%)
  - RECO DiskGroup:** RECO06 (HIGH, Size: 20%)

Buttons at the bottom include Help, Import..., SSC Zone defaults..., Save, < Back, Next >, Finish, and Cancel.

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 Domains and Zones” on page 16](#) 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 137](#).

---

**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 137](#), 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: RECO6
- 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 Exadata Storage Server Disk Space for Additional Zones” on page 74](#) 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.

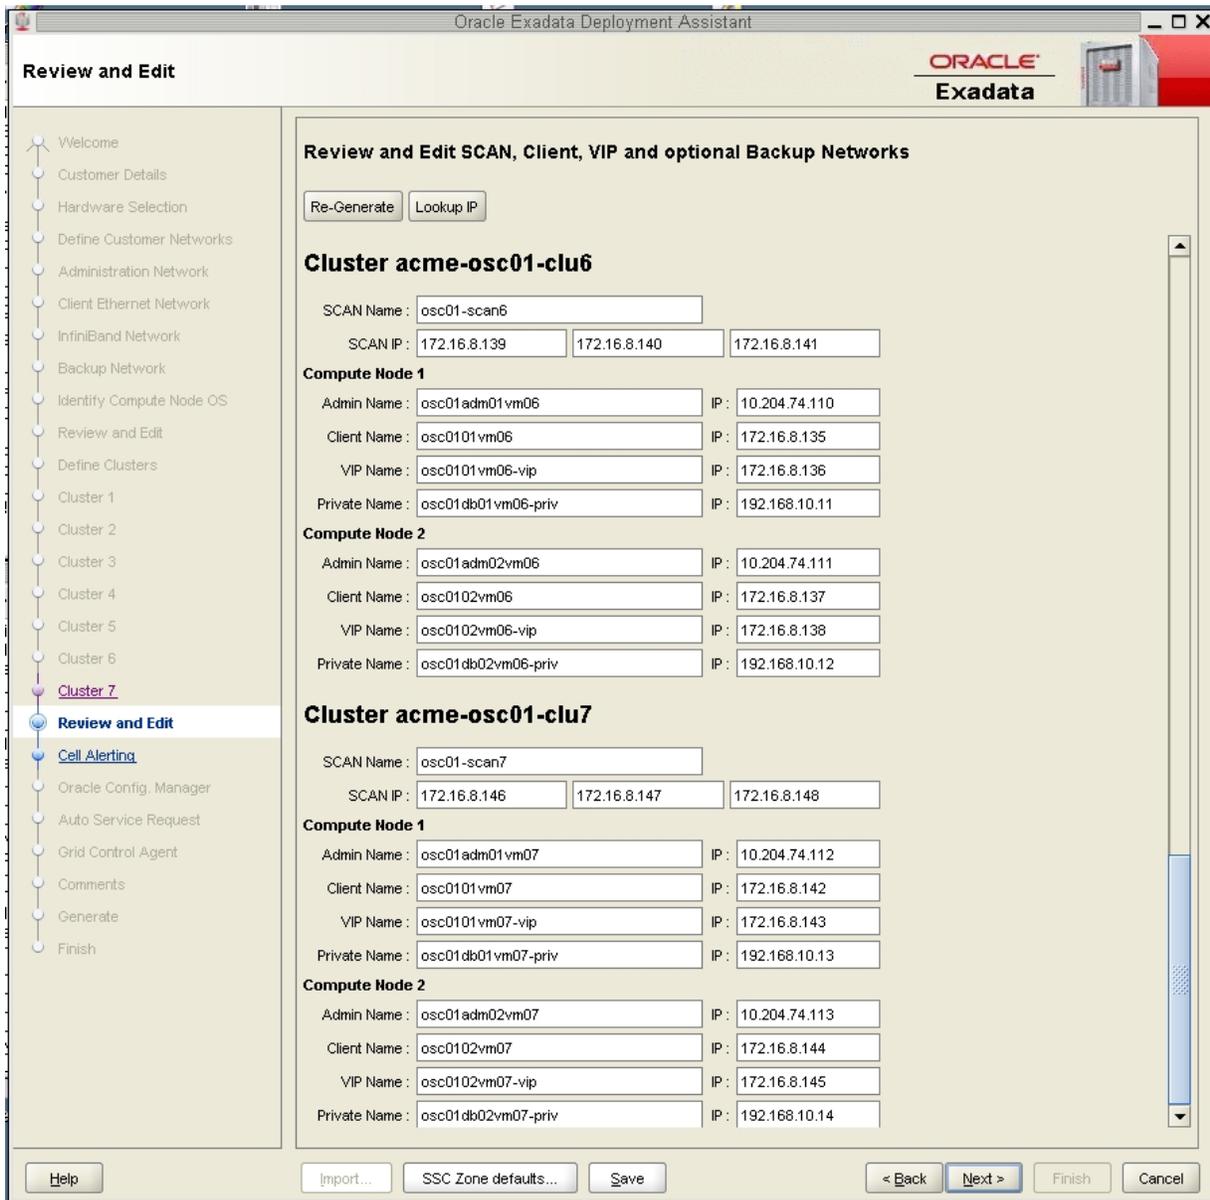
For example, if the first available starting IP address for the client network for the first new cluster is 172.16.8.150, and the Pool Size field next to the Start IP field for that network shows 7, then IP addresses 172.16.8.150 to 172.16.8.156 are assigned to the client access network for this cluster. 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 144.

## ▼ **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 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 146.

## ▼ 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.**

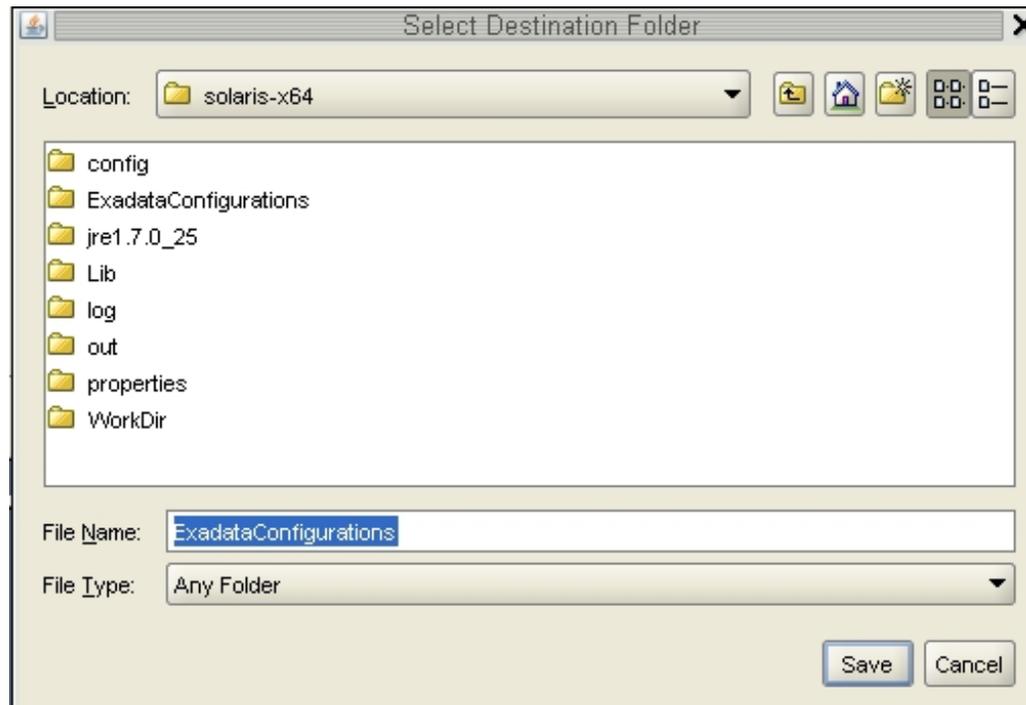
**3. Generate the configuration files.**

Go to [“Generate the Configuration Files”](#) on page 146.

## ▼ 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 149.

# 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 149](#)
- [“Delete a Template Zone From a Database Domain” on page 151](#)

## ▼ Create a Template Zone on a Database Domain

1. **Log in to the first Database Domain in your Oracle SuperCluster T5-8 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 153](#).

- 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 70](#) 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/
osc01dbadmin01-01-20121129093659.log and traces in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
osc01dbadmin01-01-20121129093659.trc and o/p in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
osc01dbadmin01-01-20121129093659.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.

## ▼ 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 70](#) 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.

---

**Note** - You can also destroy zones on Database Domains if necessary, using the command `ssc_exavm -destroy -virtualGuest zone-name`, where *zone-name* is the name of the zone. However, use caution when using this command to destroy zones on Database Domains.

---



## Creating VNETs

---

The following procedures describe how to create or delete VNETs.

- [“Determining if Additional VNETs Are Needed for a Database Domain” on page 153](#)
- [“Locate the Control Domain” on page 155](#)
- [“Create VNETs for Zones” on page 156](#)
- [“Delete VNETs” on page 158](#)

### 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 \(CPU-to-Database Domain Mapping\)” on page 153](#)
- [“Determine if Additional VNETs Are Needed \(Software Commands\)” on page 154](#)

#### ▼ Determine if Additional VNETs Are Needed (CPU-to-Database Domain Mapping)

You must create additional VNETs if you are setting up zones on the following Database Domains on your SPARC T5-8 server:

- **Half Rack:** For the Small Domains that are associated with the following CPUs:
  - CPU1
  - CPU6

See [“Half Rack Domain Configurations” on page 13](#) for more information on which Small Domains would have these CPUs associated with them.

- **Full Rack:** For the following domains:
  - Medium Domains associated with the following CPUs:
    - CPU2/CPU3
    - CPU4/CPU5
  - Small Domains associated with the following CPUs:
    - CPU1
    - CPU2
    - CPU3
    - CPU4
    - CPU5
    - CPU6

See [“Full Rack Domain Configurations” on page 14](#) for more information on which Medium Domains and Small Domains would have these CPUs associated with them.

- **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 161](#).
  - 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 156](#).

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

Another way to determine whether or not you will 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:**

```
# impstat -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 INTERFACES 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 161](#).
- If you see output similar to the following, then you must create additional VNETs for the zones that you will be 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 156](#).

## ▼ Locate the Control Domain

The tool for creating VNETs must be run on the control domain. The following procedure provides instructions for locating the control domain on your system.

**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 59](#) for the location of the `addvnet-wrapper.sh` script.

2. **On the first SPARC T5-8 server, list the Database Domains with zones 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 155](#) for instructions on locating the control domain.

4. **From the control domain, add VNETs to the guest domain on this SPARC T5-8 server:**

```
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 SPARC T5-8 server.

For example, if you have the following setup:

- The name of the guest domain is `osc01dbadmin02-01`.
- You will be setting up 4 zones on the Database Domains on this SPARC T5-8 server.

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 `mnet`, and ending with either a `0` or a `1` (for example, `z0mnet1`).
- Viewed from the 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 `zonbnetxx` (such as `zonbnet12`, `zonbnet13`, 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 151](#), the net pool is freed and renamed back to `zonenetxx` (such as `zonenet12`, `zonenet13`, and so on) so that any future addition of zones using zone tools can pick those VNETs.

##### 5. Check the guest domains to ensure that 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=mngnet1|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=z0mngnet0|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=z0mngnet1|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=z1mngnet0|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=z1mngnet1|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 the second SPARC T5-8 server.
7. Run the Java OneCommand.  
Go to [“Running the Java OneCommand” on page 161](#).

## ▼ 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 59](#) 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=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=mgnnet1|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 161](#)
- [“Set Up the Java OneCommand” on page 162](#)
- [“Run the Java OneCommand” on page 163](#)
- [“NTP Issue” on page 166](#)

### ▼ 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 SPARC T5-8 server 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 59.

**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 59.

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

See “[Verify Configuration Tool Installation](#)” on page 70 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 163.

## ▼ 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 files for the new clusters that were generated at the end of the OEDA process when you added zones to the Database Domains.**

The configuration files for the new clusters that you created should be named *customer-name-name-prefix-cluster-clucluster-number.xml*.

For example, using the example configuration, the two new cluster files that were created at the end of the OEDA process as described in [“Creating Configuration Files Using OEDA” on page 73](#) would be:

- Acme-osc01-cluster-clu6.xml
- Acme-osc01-cluster-clu7.xml

**2. Using the first new cluster configuration file, display the list of functions that will be performed by the Java OneCommand:**

```
./install.sh -cf /path-to-config-file/customer-name.name-prefix.xml -l
```

For example:

```
./install.sh -cf /path-to-config-file/Acme-osc01-cluster-clu6.xml -l
```

Information similar to the following is displayed:

1. Validate Configuration File
2. Setup Required Files
3. Create Users
4. Setup Cell Connectivity
5. Verify IB and Calibrate Cells
6. Create Cell Disks
7. Create Grid Disks

8. Install Cluster Software
9. Initialize Cluster Software
10. Install Database Software
11. Relink Database with RDS
12. Create ASM Diskgroups
13. Create Databases
14. Apply Security Fixes
15. Create Installation Summary
16. 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 IB and calibrate cells (step 5 in the example output above)
- Create cell disks (step 6 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 cluster 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-cluster-clu6.xml -s 1
```

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

- Verify IB 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-cluster-clu6.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 151](#) 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. Repeat [Step 2](#) through [Step 4](#) for each new cluster configuration file that you generated at the end of the OEDA process.
6. 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 167.

## NTP Issue

When your Oracle SuperCluster T5-8 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 will be 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) will only run 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 will cause ctssd to try to go into active mode, which will result 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 167](#)

### ▼ 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:

```
# impstat -g
GROUP      GROUPNAME  STATE  FDT  INTERFACES
stor_imp0  stor_imp0  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 173](#)

### ▼ 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)
bondeth0  bondeth0   ok      --
net4 (net7)
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 6b](#), 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
net4          Ethernet
up          10000 full     ixgbe2
net7          Ethernet
up          10000 full     ixgbe5

```

**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 6b](#)
- *net\_device#2* is the second client network device that you saw in the output in [Step 6b](#)

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, 153

## C

client access network  
  after zone creation, 56  
  before zone creation, 43, 48

## D

determining  
  amount of space for each disk group, 27  
  available storage capacity, 23  
  maximum amount of available space for RACs, 26  
  number of RACs, 26  
disk groups  
  considerations when planning layout, 21  
  determining amount of disk space, 27  
  planning layout, 22  
  recording planned layout, 38

## E

Exadata Storage Servers  
  considerations when planning disk group  
  layout, 21  
  example layout scenario, 27  
  planning disk group layout, 22  
  recording planned disk group layout, 38

rules and restrictions for disk group layout, 21

## H

host management network  
  after zone creation, 55  
  before zone creation, 42, 47

## I

InfiniBand network  
  after zone creation, 56  
  before zone creation, 45, 51  
iSCSI LUNs  
  locating iscsi-lun.sh script, 71  
iscsi-lun.sh script  
  locating, 71

## N

network setup  
  after zone creation, 54  
  before zone creation, 41  
  RACs, 57  
networks  
  client access  
    after zone creation, 56  
    before zone creation, 43, 48  
  host management  
    after zone creation, 55  
    before zone creation, 42, 47  
  InfiniBand

- after zone creation, 56
  - before zone creation, 45, 51
- O**
- Oracle Exadata Deployment Assistant
    - creating the configuration files, 79, 122, 162
    - downloading, 70
    - understanding, 78
  - Oracle OneCommand utility
    - downloading, 70
- P**
- preparing to set up zones, 59
  - process for setting up zones on Database Domains
    - tasks, 59, 73
- R**
- RACs
    - determining maximum allowable disk space, 26
    - determining number of RACs, 26
    - network setup, 57
    - planning number of zones and RACs, 17
    - recording planned configuration, 36
    - rules and recommendations, 20
  - recording
    - existing configuration information, 29
    - planned Exadata Storage Server disk group layout, 38
    - planned RAC configuration, 36
    - planned zone configuration, 32
  - repository
    - determining location of, 63
    - installing or updating files from local repository, 65
    - installing or updating files from remote repository, 64
  - restrictions
    - Exadata Storage Server disk group layout, 21
    - setting up zones and RACs, 20
- rmvnet.sh script
    - deleting VNETs, 158
    - locating, 71
  - rules
    - Exadata Storage Server disk group layout, 21
    - setting up zones and RACs, 20
- S**
- ssc\_exavm tool
    - creating template zone, 150
    - deleting template zone, 151
- T**
- template zone
    - creating, 150
    - deleting, 151
    - understanding, 149
- V**
- VNETs
    - creating, 153
    - deleting, 158
    - locating rmvnet.sh script, 71
- Z**
- zonemanifest.sh script, locating, 71
  - zones
    - network setup after zone creation, 54
    - network setup before zone creation, 41
    - number allowed per domain, 19
    - planning number of zones and RACs, 17
    - recording planned zone configuration, 32
    - rules and recommendations, 19