

**Oracle® SuperCluster M8 and
SuperCluster M7 Zones With Oracle
Database on Database Domains
Configuration Guide**

ORACLE®

Part No: E58636-05
June 2020

Part No: E58636-05

Copyright © 2016, 2020, Oracle and/or its affiliates.

License Restrictions Warranty/Consequential Damages Disclaimer

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.

Warranty Disclaimer

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.

Restricted Rights Notice

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 embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

Hazardous Applications Notice

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.

Trademark Notice

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

Third-Party Content, Products, and Services Disclaimer

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.

Pre-General Availability Draft Label and Publication Date

Pre-General Availability: 2020-01-15

Pre-General Availability Draft Documentation Notice

If this document is in public or private pre-General Availability status:

This documentation is in pre-General Availability status and is intended for demonstration and preliminary use only. It may not be specific to the hardware on which you are using the software. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this documentation and will not be responsible for any loss, costs, or damages incurred due to the use of this documentation.

Oracle Confidential Label

ORACLE CONFIDENTIAL. For authorized use only. Do not distribute to third parties.

Revenue Recognition Notice

If this document is in private pre-General Availability status:

The information contained in this document is for informational sharing purposes only and should be considered in your capacity as a customer advisory board member or pursuant to your pre-General Availability trial agreement only. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle Master Agreement, Oracle License and Services Agreement, Oracle PartnerNetwork Agreement, Oracle distribution agreement, or other license agreement which has been executed by you and Oracle and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

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

Référence: E58636-05

Copyright © 2016, 2020, Oracle et/ou ses affiliés.

Restrictions de licence/Avis d'exclusion de responsabilité en cas de dommage indirect et/ou consécutif

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.

Exonération de garantie

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.

Avis sur la limitation des droits

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 embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

Avis sur les applications dangereuses

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 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.

Marques

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 Inside 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, Epyc, et le logo AMD sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Avis d'exclusion de responsabilité concernant les services, produits et contenu tiers

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.

Date de publication et mention de la version préliminaire de Disponibilité Générale ("Pre-GA")

Version préliminaire de Disponibilité Générale ("Pre-GA") : 15.01.2020

Avis sur la version préliminaire de Disponibilité Générale ("Pre-GA") de la documentation

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère public ou privé :

Cette documentation est fournie dans la Version préliminaire de Disponibilité Générale ("Pre-GA") et uniquement à des fins de démonstration et d'usage à titre préliminaire de la version finale. Celle-ci n'est pas toujours spécifique du matériel informatique sur lequel vous utilisez ce logiciel. Oracle Corporation et ses affiliés déclinent expressément toute responsabilité ou garantie expresse quant au contenu de cette documentation. Oracle Corporation et ses affiliés ne sauraient en aucun cas être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'utilisation de cette documentation.

Mention sur les informations confidentielles Oracle

INFORMATIONS CONFIDENTIELLES ORACLE. Destinées uniquement à un usage autorisé. Ne pas distribuer à des tiers.

Avis sur la reconnaissance du revenu

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère privé :

Les informations contenues dans ce document sont fournies à titre informatif uniquement et doivent être prises en compte en votre qualité de membre du customer advisory board ou conformément à votre contrat d'essai de Version préliminaire de Disponibilité Générale ("Pre-GA") uniquement. Ce document ne constitue en aucun cas un engagement à fournir des composants, du code ou des fonctionnalités et ne doit pas être retenu comme base d'une quelconque décision d'achat. Le développement, la commercialisation et la mise à disposition des fonctions ou fonctionnalités décrites restent à la seule discrétion d'Oracle.

Ce document contient des informations qui sont la propriété exclusive d'Oracle, qu'il s'agisse de la version électronique ou imprimée. Votre accès à ce contenu confidentiel et son utilisation sont soumis aux termes de vos contrats, Contrat-Cadre Oracle (OMA), Contrat de Licence et de Services Oracle (OLSA), Contrat Réseau Partenaires Oracle (OPN), contrat de distribution Oracle ou de tout autre contrat de licence en vigueur que vous avez signé et que vous vous engagez à respecter. Ce document et son contenu ne peuvent en aucun cas être communiqués, copiés, reproduits ou distribués à une personne extérieure à Oracle sans le consentement écrit d'Oracle. Ce document ne fait pas partie de votre contrat de licence. Par ailleurs, il ne peut être intégré à aucun accord contractuel avec Oracle ou ses filiales ou ses affiliés.

Accessibilité de la documentation

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

Accès aux services de support Oracle

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

Contents

Using This Documentation	11
Product Documentation Library	11
Feedback	11
Planning to Set Up Zones on Database Domains	13
Important Cautions	13
Naming Conventions	14
PDomains Overview	14
Logical Domain Configuration Overview	15
Determining the Cores Available for Domains and Zones	15
Cores Available for Domains and Zones	15
Example Configuration	16
Memory Available for Domains	18
Guidelines for Planning Zones	18
Guidelines for Planning Zone Clusters	19
Guidelines for Planning the Storage Server Disk Group Layout	20
Planning the Exadata Storage Server Disk Group Layout	21
Understanding an Example Scenario	24
Recording Your Existing and Planned Configuration Information	26
Recording Your Zone Configuration Information	26
Recording Your Cluster Configuration Information	32
Recording Your Exadata Storage Server Disk Group Layout	33
Preparing to Configure Zones on Database Domains	35
▼ Prepare to Configure Zones	35
Existing Network Worksheet	37
▼ Update the Base Software	38
▼ Determine the Repository Location	39

▼ Install or Update Packages From the Remote Repository	40
▼ Install or Update Files From the Local Repository	42
▼ Verify Configuration Tool Installation	46
Creating Configuration Files Using OEDA	49
Example Configuration	49
▼ Verify Exadata Storage Server Disk Space for Additional Zones	50
Creating Configuration Files Using OEDA	51
▼ Locate the Necessary Files to Set Up Zones	52
▼ Set Up the OEDA	54
▼ Start the OEDA	54
▼ Import the Most Recent OEDA Configuration File	57
▼ Review Existing Configuration Information	58
▼ Review the Information in the Identify Compute Node Operating System Page	59
▼ Review the Information in the Management and Private Networks Page	61
▼ Complete the Define Clusters Page	62
▼ Set Zone Default Configurations	70
▼ Complete the Cluster Page	73
▼ Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page	77
▼ Verify Remaining Configuration Information	79
▼ Generate the Configuration Files	79
Creating the Template Zone on Each Database Domain	83
▼ Create a Template Zone on a Database Domain	83
▼ Delete a Template Zone From a Database Domain	85
Running the <code>install.sh</code> Script	87
▼ Set Up Public Key Authentication for ZFS Storage Controllers	87
▼ Set Up the <code>install.sh</code> Script	88
▼ Run the <code>install.sh</code> Script	90
NTP Issue	92
Creating Additional Links on the IB Storage Network for Zones	93
▼ Create Additional Links on the IB Storage Network for Zones	93

Setting Up VLAN Tagging for Zones	97
▼ Set Up VLAN Tagging for Zones	97
Index	101

Using This Documentation

- **Overview** – Describes how to set up zones and Real Application Clusters (RACs) on the Database Domains in Oracle SuperCluster M8 and Oracle SuperCluster M7
- **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 http://docs.oracle.com/cd/E58626_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 M8 or Oracle SuperCluster M7 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 *Oracle SuperCluster Systems Zones on Application Domains Configuration Guide*, located here:

http://docs.oracle.com/cd/E58626_01/index.html

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

- “Important Cautions” on page 13
- “Naming Conventions” on page 14
- “PDomains Overview” on page 14
- “Logical Domain Configuration Overview” on page 15
- “Determining the Cores Available for Domains and Zones” on page 15
- “Memory Available for Domains” on page 18
- “Guidelines for Planning Zones” on page 18
- “Guidelines for Planning Zone Clusters” on page 19
- “Guidelines for Planning the Storage Server Disk Group Layout” on page 20
- “Recording Your Existing and Planned Configuration Information” on page 26

Important Cautions



Caution - Ensure that you back up all existing databases before running the `install.sh` script.



Caution - You must use the latest version of the OEDA 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 the `install.sh` script 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 the `install.sh` script 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:

- *scm7number* for the Oracle SuperCluster M8 or Oracle SuperCluster M7
- *rnumber* for the cluster number
- *cnumber* for the compute server number
- *dbnumber* for the Database Domain number
- *znumber* for the zone number

PDomains Overview

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

Each compute server is split into two partitions (two PDomains), where the bottom four CMIOU slots are part of the first partition (PDomain 0), and the top four CMIOU slots are part of the second partition (PDomain 1). Each compute server can have either two CMIOUs (quarter-populated), four CMIOUs (half-populated), or eight CMIOUs (fully-populated).

For detailed information about SuperCluster PDomains, refer to this section in the *Oracle SuperCluster M8 and SuperCluster M7 Overview Guide*: https://docs.oracle.com/cd/E58626_01/html/E58633/z4000d75165586.html

Logical Domain Configuration Overview

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

For detailed information about SuperCluster logical domain configurations, refer to this section in the *Oracle SuperCluster M8 and SuperCluster M7 Overview Guide*: https://docs.oracle.com/cd/E58626_01/html/E58633/cgffeihd.html#scrolltoc

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 15
- “Example Configuration” on page 16

Cores Available for Domains and Zones

Every CMIOU has 32 cores of CPU resources. The amount of CPU resources available for each domain varies, depending on the number of CMIOUs that are associated with that domain.

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). The remaining cores in the Database Domain are available for the nonglobal zones (the zones in the

Database Domain). The number of cores that are set aside for the global zone varies, depending on the number of CMIOUs that are associated with the domain:

- One CMIOU associated with an LDom: 2 cores are reserved for the global zone, and the remaining cores are available for the nonglobal zones.
- Two or more CMIOUs associated with an LDom: 4 cores are reserved for the global zone, and the remaining cores are available for the nonglobal zones.

When using the information above, 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 compute server. Carefully consider the expected workload on each zone that you create, so that you allot the appropriate number of cores to those zones.

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 49](#), where:

- Oracle SuperCluster M7 with a single compute node, with Exadata Storage Servers with High Capacity disks.
- The single compute node with two PDomains, or an R1 PDomain configuration.
- Four CMIOUs in each PDomain, or an R1-4 PDomain configuration.
- The U4-4 LDom configuration (4 LDom) will be set up on both PDomains in Oracle SuperCluster M7.
- The following LDom types will be assigned to each LDom in both PDomains:
 - LDom 1: Database Domain with 4 zones
 - LDom 2: Database Domain (without zones)
 - LDom 3: Solaris 11 Application Domain
 - LDom 4: Root Domain
- 5 RAC instances set up on the Oracle SuperCluster M7:
 - One RAC instance for each zone on the Database Domains with zones (4 RAC instances total for the zones)

- One RAC instance for the two Database Domains (without zones)

For the purposes of this exercise, assume the following:

- Four zones will be set up on LDom 1 (the first Database Domain) in both PDomains.
- No zones will be set up on LDom 2 (the second Database Domain) in both PDomains.

Based on the U4-4 LDom configuration (refer to https://docs.oracle.com/cd/E58626_01/html/E58633/cgffeihd.html#scrolltoc) and “Cores Available for Domains and Zones” on page 15, you can determine that 32 cores are allocated to both Database Domains on each PDomain.

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

- “LDom 1 (First Database Domain)” on page 17
- “LDom 2 (Second Database Domain)” on page 17

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 PDomain will contain four zones, and it has 32 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 8 cores ($32 / 4 = 8$).

However, based on the information provided in “Cores Available for Domains and Zones” on page 15, a certain number of cores will be set aside for the global zone. This domain has one CMIOU associated with it, which means that two cores will be set aside for the global zone, so you actually have 30 cores available for the four zones in this Database Domain. So, for the zones in this Database Domain, you could allocate 8 cores for two of the zones (16 cores) and 7 cores for the other two zones (14 cores).

LDom 2 (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 PDomains.

Memory Available for Domains

The amount of memory resources available for database zones depends on the type of SuperCluster you have. Each CMIOU has 16 DIMM slots, for:

- A total of 1 TB (64 GB DIMMs) of total available memory in Oracle SuperCluster M8
- A total of 512 GB (32 GB DIMMs) of total available memory in Oracle SuperCluster M7

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

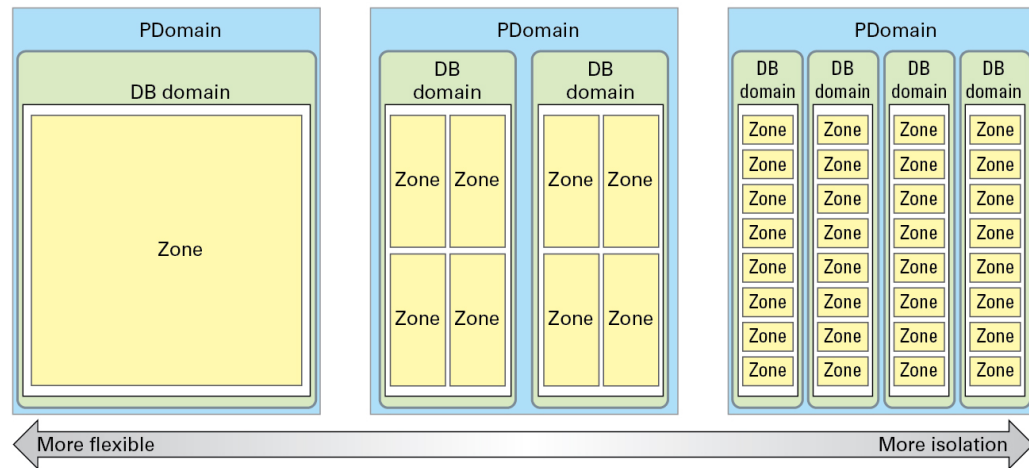
For example, assume you have Oracle SuperCluster M7, and you have a Database Domain that has one CMIOU associated with it. By default, 512 GB of memory would be available to that Database Domain. You could therefore have four equal-sized database zones within that Database Domain, where each database zone has 100 GB of memory assigned to it, for a total of 400 GB of memory for all database zones. The remaining 112 GB of memory in this Database Domain could then be saved for future database zones that you might want to create on this Database Domain.

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 zones and clusters on the Database Domains:

- When you first install the operating system instances on a domain, that domain is automatically designated as the global zone. Any 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.
- You can create clusters with zones of different sizes. However, each zone that is a member of a cluster should be the same size with regards to the number of cores being used by each zone in the cluster. For example, if you are creating a two-node cluster and the first zone in that cluster uses four cores, the other zone in that cluster should also use four cores.

- You can cluster zones together across multiple Database Domains that are of different sizes. For example, you can create a two-node cluster, where there is one zone on the Database Domains on the two PDomains, and those zones are clustered together, even if one of the Database Domains has one CMIOU associated with it and the other Database Domain has two CMIOUs associated with it. 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.
- While it is possible to create a cluster with zones in the same Database Domain or on the same PDomain, that configuration is not a highly-available configuration. Instead, create clusters where the zones are on separate Database Domains, and the cluster spans across PDomains, so that there is not a single point of failure.

Guidelines for Planning the Storage Server Disk Group Layout

An Oracle representative should have set up 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 Oracle SuperCluster M7 with three 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 Oracle SuperCluster M7 with 11 Exadata Storage Servers and an Oracle Exadata Storage Expansion with 18 Exadata Storage Servers connected to your Oracle SuperCluster M7. 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 21](#)
- [“Understanding an Example Scenario” on page 24](#)

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 (scm7r1 and scm7r2), then you might use these names for the three disk groups for each cluster:

- First cluster (scm7r1):
 - DATA_scm7r1
 - RECO_scm7r1

- DBFS_scm7r1
- Second cluster (scm7r2):
 - DATA_scm7r2
 - RECO_scm7r2
 - DBFS_scm7r2

Record your Exadata Storage Server disk group layout in the worksheet in [“Recording Your Exadata Storage Server Disk Group Layout”](#) on page 33.

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

- [“Determine the Available Existing Storage Capacity”](#) on page 22
- [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster”](#) on page 23
- [“Determine the Amount of Space for Each Disk Group in Each Cluster”](#) on page 24

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.

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

- X7-2L Exadata Storage Server:
 - 8 x 6.4 TB Extreme Flash drives
 - 12 x 10 TB High Capacity drives
- X6-2L Exadata Storage Server:
 - 8 x 3.2 TB Extreme Flash drives
 - 12 x 8 TB High Capacity drives
- X5-2L Exadata Storage Server:
 - 8 x 1.6 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 types of Exadata Storage Servers.

TABLE 1 X7-2L Single Storage Server Capacities

Capacity Type	8 x 6.4 TB Extreme Flash Drives	12 x 10 TB High Capacity Drives
Raw disk capacity	51.2 TB	120 TB
Usable mirrored capacity (ASM normal redundancy)	20 TB	50 TB
Usable triple mirrored capacity (ASM high redundancy)	17.2 TB	37.5 TB

TABLE 2 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

TABLE 3 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

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 PDomain, the size of each Database Domain, and so on.

See [“Guidelines for Planning Zones” on page 18](#) 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 24](#)
- [“Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster \(Example Scenario\)” on page 25](#)
- [“Determine the Amount of Space for Each Disk Group in Each Cluster” on page 25](#)

Determine the Available Existing Storage Capacity (Example Scenario)

Consider the following example scenario:

- Oracle SuperCluster M7, with three X6-2L Exadata Storage Servers

- 12 high-capacity (8 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 22](#), you have **120 TB** of usable storage capacity available for the Exadata Storage Servers in your Oracle SuperCluster M7.

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 120 TB of usable disk capacity, and you want to create two clusters for the zones, you could use the existing 120 TB of space in one of several ways:

- Use the entire 120 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 60 TB.
 - Create two clusters of different sizes, where the total disk space used by the disk groups for the two clusters equals 120 TB. For example, one cluster with the total disk space used by the disk groups equaling 80 TB and the other cluster with the total disk space used by the disk groups equaling 40 TB.

Keep in mind that by using the entire 120 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 120 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 60 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 60 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 M7, 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 Zone Configuration Information” on page 26](#)
- [“Recording Your Cluster Configuration Information” on page 32](#)
- [“Recording Your Exadata Storage Server Disk Group Layout” on page 33](#)

Recording Your Zone Configuration Information

You can use the tables in this section to plan the zone configuration, and to gather the configuration information that you are eventually asked to provide.

Use the set of tables to record the following information for each zone:

- The number of zones that you want to deploy on each Database Domain (1 to 8 zones)
- The number of cores for each zone (1 minimum, the maximum depends on how many cores are available in the Database Domain minus 2 or 4 (for the Database Domain), and the number of zones planned for that domain. See [“Determining the Cores Available for Domains and Zones” on page 15](#)).

- 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

Enter the zone configuration information in the following worksheets:

- [“ Database Domain 1 Zone Configuration Information ” on page 27](#)
- [“ Database Domain 2 Zone Configuration Information ” on page 28](#)
- [“ Database Domain 3 Zone Configuration Information ” on page 29](#)
- [“ Database Domain 4 Zone Configuration Information ” on page 31](#)

Database Domain 1 Zone Configuration Information

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

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

Database Domain 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:	Host management:

Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
		Client access: IB: VIP:	Client access: IB: VIP:
4	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
5	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
6	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:
8	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

Database Domain 3 Zone Configuration Information

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

Recording Your Existing and Planned Configuration Information

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

Database Domain 4 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: VIP:	Host management: Client access: IB: VIP:
7	Cores:	Host management: Client access: IB: VIP:	Host management: Client access: IB: VIP:

Zone Number	Number of Cores Assigned to This Zone	Host Names	IP Addresses
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, scm7r1 or scm7r2)
- 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 PDomains for high availability. If you are creating a cluster, and a zone in a Database Domain is a member of that cluster on one PDomain, then other zones that are members of that cluster should be in a similar Database Domain on the second PDomain.

The worksheet has entries for eight clusters. But you can have more than eight clusters in your Oracle SuperCluster M8 or Oracle SuperCluster M7, 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 eight 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:

Cluster Number and Name	Zone Membership Information	SCAN Host Names and IP Addresses for Each Cluster
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:

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 eight clusters, but you could have more than eight clusters in your Oracle SuperCluster M7, 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 eight 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 21](#). For example, if the name of a particular cluster is scm7r1, then the names of the disk groups for that cluster should be something similar to the following:

- DATA_scm7r1
- RECO_scm7r1
- DBFS_scm7r1

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:	

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

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 35](#)
- [“Existing Network Worksheet” on page 37](#)
- [“Update the Base Software” on page 38](#)
- [“Determine the Repository Location” on page 39](#)
- [“Install or Update Packages From the Remote Repository” on page 40](#)
- [“Install or Update Files From the Local Repository” on page 42](#)
- [“Verify Configuration Tool Installation” on page 46](#)

▼ Prepare to Configure Zones

Complete these tasks before setting up zones and clusters on the Database Domains with zones in your Oracle SuperCluster M8 or Oracle SuperCluster M7.

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 20](#) 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 M8 or Oracle SuperCluster M7
- IP addresses and host names for Oracle ILOM for the following components:
 - Compute 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 M8 or Oracle SuperCluster M7 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 37](#).

Existing Network Worksheet

Component	Existing Host Names and IP Addresses	
	Host Name	IP Address
Oracle ILOM for compute server 1		
Oracle ILOM for compute 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		
Oracle ILOM for Exadata Storage Server 9		
Oracle ILOM for Exadata Storage Server 10		
Oracle ILOM for Exadata Storage Server 11		
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 Exadata Storage Server 9		

Component	Existing Host Names and IP Addresses	
	Host Name	IP Address
Host management network for Exadata Storage Server 10		
Host management network for Exadata Storage Server 11		
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		
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		
IB network for Exadata Storage Server 9		
IB network for Exadata Storage Server 10		
IB network for Exadata Storage Server 11		

▼ Update the Base Software

1. Log in to the first PDomain.
2. Determine if your system has the correct level of base software installed.

```
# pkg info entire
```

- If you see output similar to the following, where you have the build level shown below or later, then you have the correct level of software installed. Go to [“Determine the Repository Location” on page 39](#):

```
FMRI: pkg://solaris/entire@0.5.11,5.11-0.175.3.1.0.5.0:20130504T024145Z
```

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

- Solaris 11 Update 3 OS

- SRU 1
- Build 5 or later
- **If you do not have the correct level of software installed (if you have a build earlier than Oracle Solaris 11.3, SRU 1.5), 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 M8 or Oracle SuperCluster M7.**
Select any of the Oracle SuperCluster M8 or Oracle SuperCluster M7 options.
 - e. **In the Release field, select the SuperCluster V2.0 release or later.**
 - f. **Click the Search button.**
 - g. **Select the QUARTERLY FULL STACK DOWNLOAD PATCH FOR SUPERCLUSTER (*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 M8 or Oracle SuperCluster M7 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 M8 or Oracle SuperCluster M7

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 40](#).

```
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 M8 or Oracle SuperCluster M7. Go to [“Install or Update Files From the Local Repository” on page 42](#).

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

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

▼ Install or Update Packages From the Remote Repository

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

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

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 46](#).

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
```

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 46](#).

▼ Install or Update Files From the Local Repository

This procedure describes how to install or update files from the local repository on the ZFS storage appliance.

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/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 M8 or Oracle SuperCluster M7.**
Select any of the Oracle SuperCluster M8 or Oracle SuperCluster M7 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.**
 - h. **In the Patch Search Results page, select the option for the most recent release (1.1 or later) and click the Read Me button.**
 - i. **Compare the version number of each package shown in the Read Me file to the version numbers displayed when you entered the `pkg info` command earlier in these procedures.**
 - If the version numbers match for each package, then you have the latest versions. Go to [Step 11](#).
 - If the versions of the packages on your local repository are older than the package versions listed in the Read Me file, or if there is no package version information in the Read Me file, then you must update the packages on the local repository. Go to [Step 8](#).
7. **On a system that has access to the Internet, locate the repository ISO image from the remote Oracle site.**
- 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 M7.**
Select any of the Oracle SuperCluster M7 options.
 - e. **In the Release field, select all releases.**
 - f. **Choose Description Contains in the pull-down menus, then enter `exa-family` in the search field.**
 - g. **Click the Search button.**

8. **In the Patch Search Results page, select the option for the most recent release (1.5.9 or later) and click the Download button.**

Do not unzip the file at this time.

9. **Update the local repository on the ZFS storage appliance.**
 - a. **Locate the ISO image ZIP file that you downloaded in [Step 8](#) and copy that file to a temporary directory on the first PDomain.**

For example, copy the ZIP file to the `/var/tmp` directory on the first PDomain.

- 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 39](#).

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
```

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 46](#).

▼ 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 (MOS) note 888828.1, which provides information on the latest version of the patch with the Oracle OneCommand utility.**

Locate the Oracle Exadata Deployment Assistant (OEDA) section in the note.
 - b. **Select the most recent patch from the table.**

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 22694905 would be a more recent version of the patch than Patch 22491458.
 - c. **Select the appropriate option from the Platform or Language field for the operating system on the machine that you are using for OEDA (`config.sh`) and the `install.sh` script.**

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:

- `applyconfig_ssc.sh`
 - `dcli`
 - `gen_cellaffinity_ssc.sh`
 - `ib_set_node_desc_ssc.sh`
 - `iscsi-lun.sh`
 - `iscsi-zpool.sh`
 - `oidcli`
 - `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 OEDA (`config.sh`) to create the configuration files to be used by the `install.sh` script later on in this process.**

Go to [“Creating Configuration Files Using OEDA”](#) on page 49.

Creating Configuration Files Using OEDA

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

- [“Example Configuration” on page 49](#)
- [“Verify Exadata Storage Server Disk Space for Additional Zones” on page 50](#)
- [“Creating Configuration Files Using OEDA” on page 51](#)

Example Configuration

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

- Oracle SuperCluster M7 with a single compute node, with Exadata Storage Servers with High Capacity disks.
- The single compute node with two PDomains, or an R1 PDomain configuration.
- Four CMIOUs in each PDomain, or an R1-4 PDomain configuration.
- The U4-4 LDom configuration (4 LDom) will be set up on both PDomains in Oracle SuperCluster M7.
- The following LDom types will be assigned to each LDom in both PDomains:
 - LDom 1: Database Domain with 4 zones
 - LDom 2: Database Domain (without zones)
 - LDom 3: Solaris 11 Application Domain
 - LDom 4: Root Domain
- 5 RAC instances set up on the Oracle SuperCluster M7:
 - One RAC instance for each zone on the Database Domains with zones (4 RAC instances total for the zones)
 - One RAC instance for the two Database Domains (without zones)

For each RAC instance, the example configuration assumes that you have zones within the Database Domain with zones, or Database Domains (without zones), that reside on different PDomains. So for the zones on Database Domains with zones for the Oracle SuperCluster M7, these zones are members of the first RAC instance:

- First zone in the Database Domain with zones in the *first* PDomain
- First zone in the Database Domain with zones in the *second* PDomain

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)

Creating Configuration Files Using OEDA



Caution - Ensure that you backup all existing databases before running the `install.sh` script.



Caution - You must use the latest version of the OEDA 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 the `install.sh` script 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 the `install.sh` script also destroy cell disks and griddisks with the Create Cell Disks step.

Follow these procedures to use OEDA to create your configuration files:

- [“Locate the Necessary Files to Set Up Zones” on page 52](#)
- [“Set Up the OEDA” on page 54](#)
- [“Start the OEDA” on page 54](#)
- [“Import the Most Recent OEDA Configuration File” on page 57](#)
- [“Review Existing Configuration Information” on page 58](#)
- [“Review the Information in the Identify Compute Node Operating System Page” on page 59](#)
- [“Review the Information in the Management and Private Networks Page” on page 61](#)
- [“Complete the Define Clusters Page” on page 62](#)
- [“Set Zone Default Configurations” on page 70](#)
- [“Complete the Cluster Page” on page 73](#)
- [“Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page” on page 77](#)
- [“Verify Remaining Configuration Information” on page 79](#)
- [“Generate the Configuration Files” on page 79](#)

▼ 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 49](#), and you provided the following information during your initial installation:

- Customer name: Example
- Name prefix: scm7

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:
 - *scm7zdbadm01-zoneConfig.xml*: Configuration file for the first zone
 - *scm7zdbadm02-zoneConfig.xml*: Configuration file for the second zone
 - *scm7zdbadm03-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:
 - *example-scm7-cluster-clu1.xml*: Configuration file for the first cluster (the first cluster with zones)
 - *example-scm7-cluster-clu2.xml*: Configuration file for the second cluster (the second cluster with zones)
 - *example-scm7-cluster-clu3.xml*: Configuration file for the third cluster (the third cluster with zones)
 - *example-scm7-cluster-clu4.xml*: Configuration file for the fourth cluster (the fourth cluster with zones)
 - *example-scm7-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:
 - `example-scm7.xml`

You will be using one of these configuration files when you use OEDA (`config.sh`) or the `install.sh` script 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 35](#).

2. **Ensure that the machine running the OEDA has Oracle JRE 1.7 or later.**

3. **Locate the OEDA ZIP files (`0cmd`) for your operating system.**

The OEDA files are part of the `0cmd` ZIP file that you downloaded earlier in this process. See [“Verify Configuration Tool Installation” on page 46](#) for more information.

There are different OEDA ZIP files for these operating systems:

- Linux
- MacOS
- Solaris (SPARC)
- Solaris (Intel)
- Windows

4. **Unzip the `0cmd` 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 `0cmd` ZIP file for your operating system.**

6. **Start the OEDA.**

Go to [“Start the OEDA” on page 54](#).

▼ Start the OEDA

1. **Start the OEDA.**

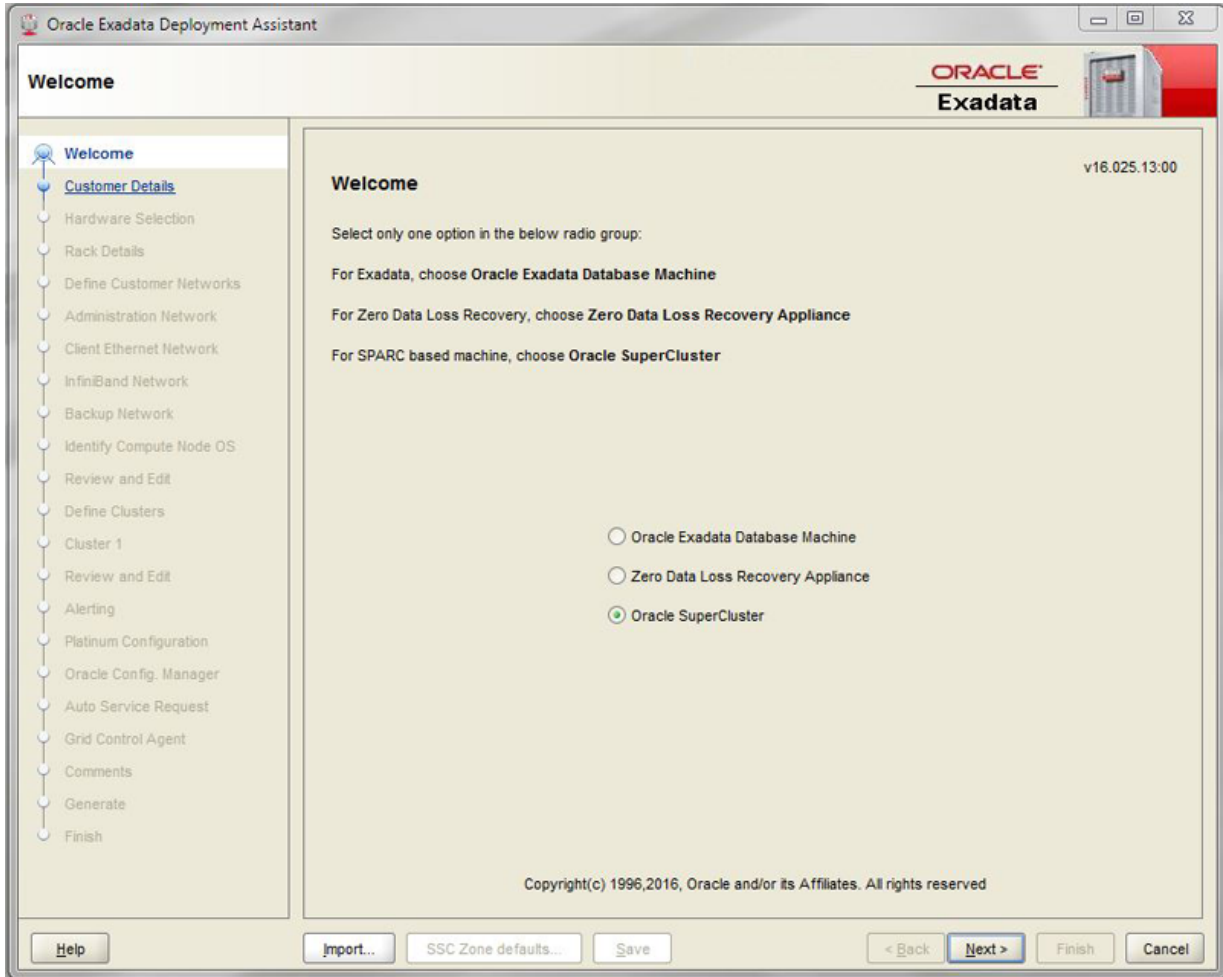
- On Linux, MacOS, or the Oracle Solaris OS, type:

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 57.

▼ 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 52](#) 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:** Example

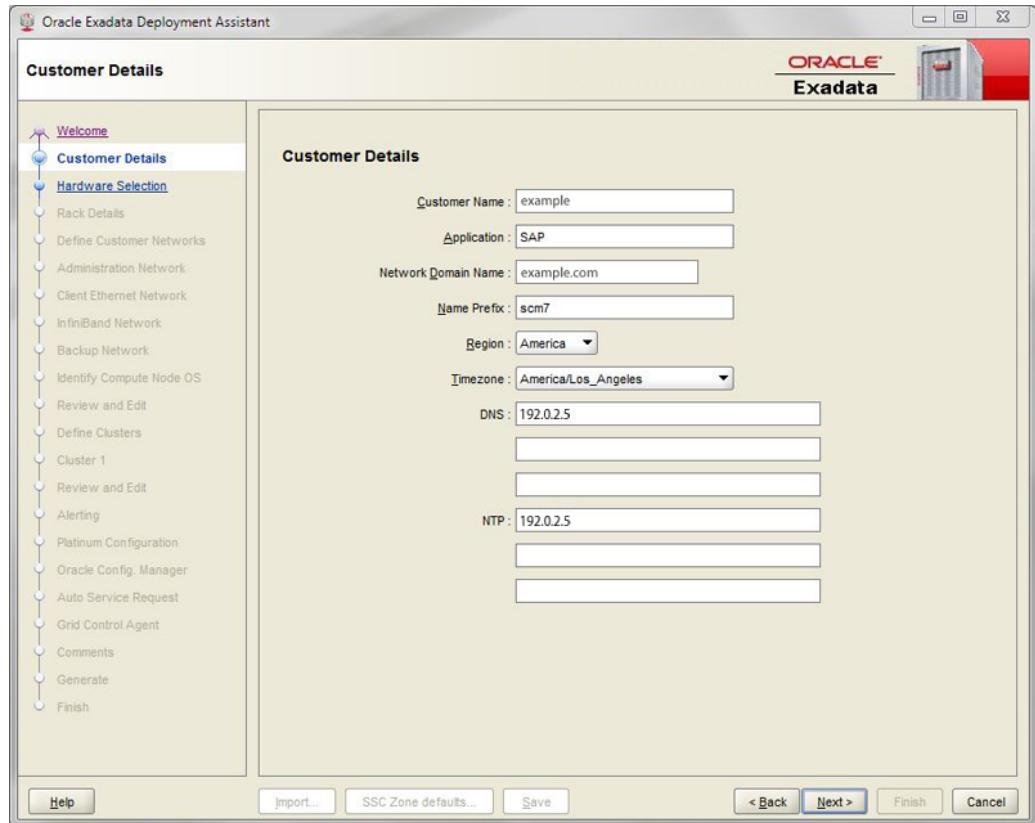
- **Name prefix:** scm7

Then the configuration file for this cluster would be named `example.scm7.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 58.

▼ 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 M7. 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 59.

▼ 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.

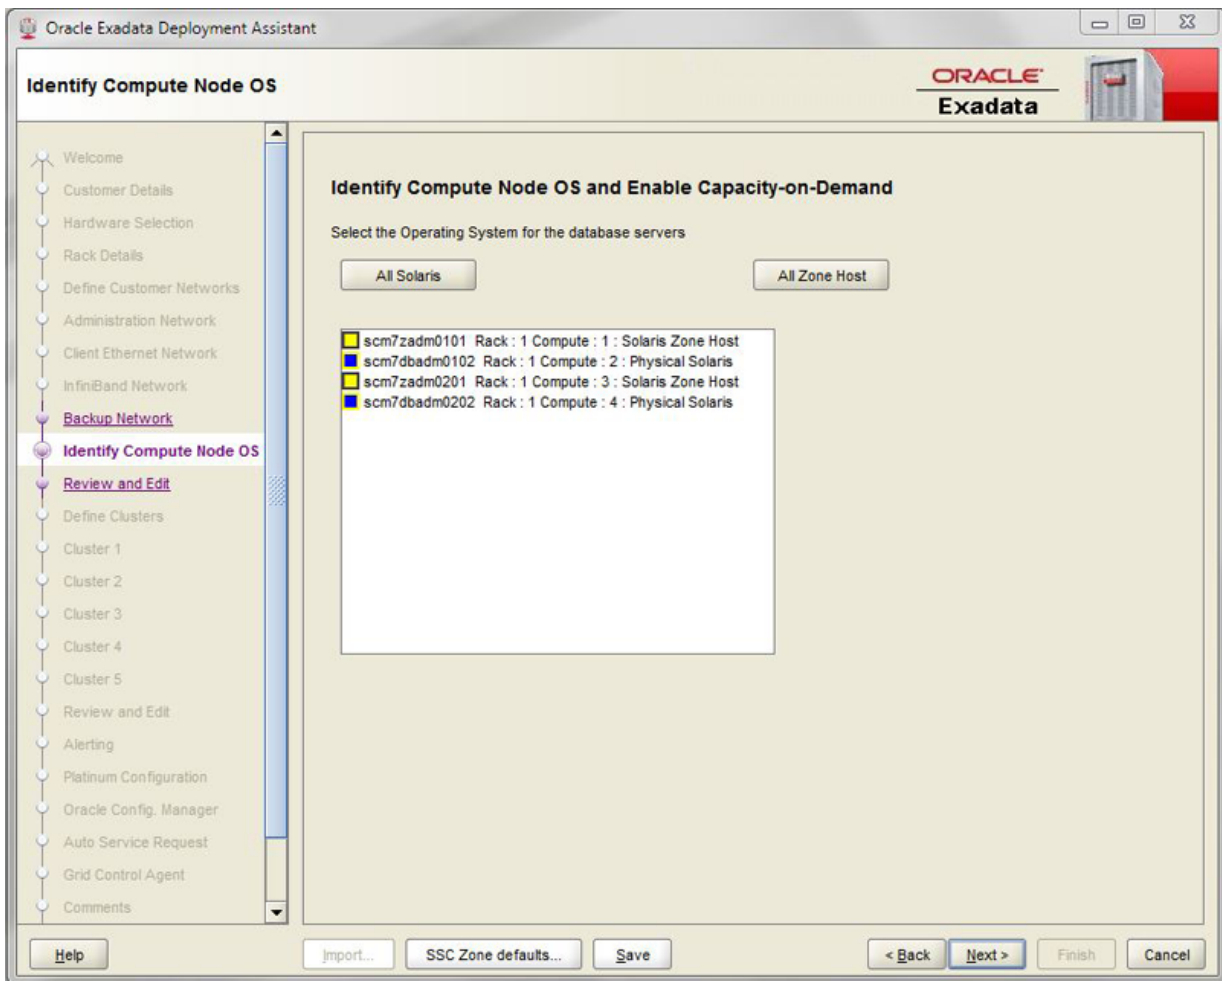
The Identify Compute Node Operating System page should display the total number of Database Domains with zones and Database Domains (without zones) on your Oracle SuperCluster M7 in your existing configuration. This page does not display information on any Applications Domains in your existing configuration.

For example, assume you have these LDoms as part of your existing configuration, as described in [“Example Configuration”](#) on page 49:

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

Then you should see the following screen:



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

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

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

2. **Click Next.**

The Review and Edit page is displayed.

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

Go to [“Review the Information in the Management and Private Networks Page”](#) on page 61.

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

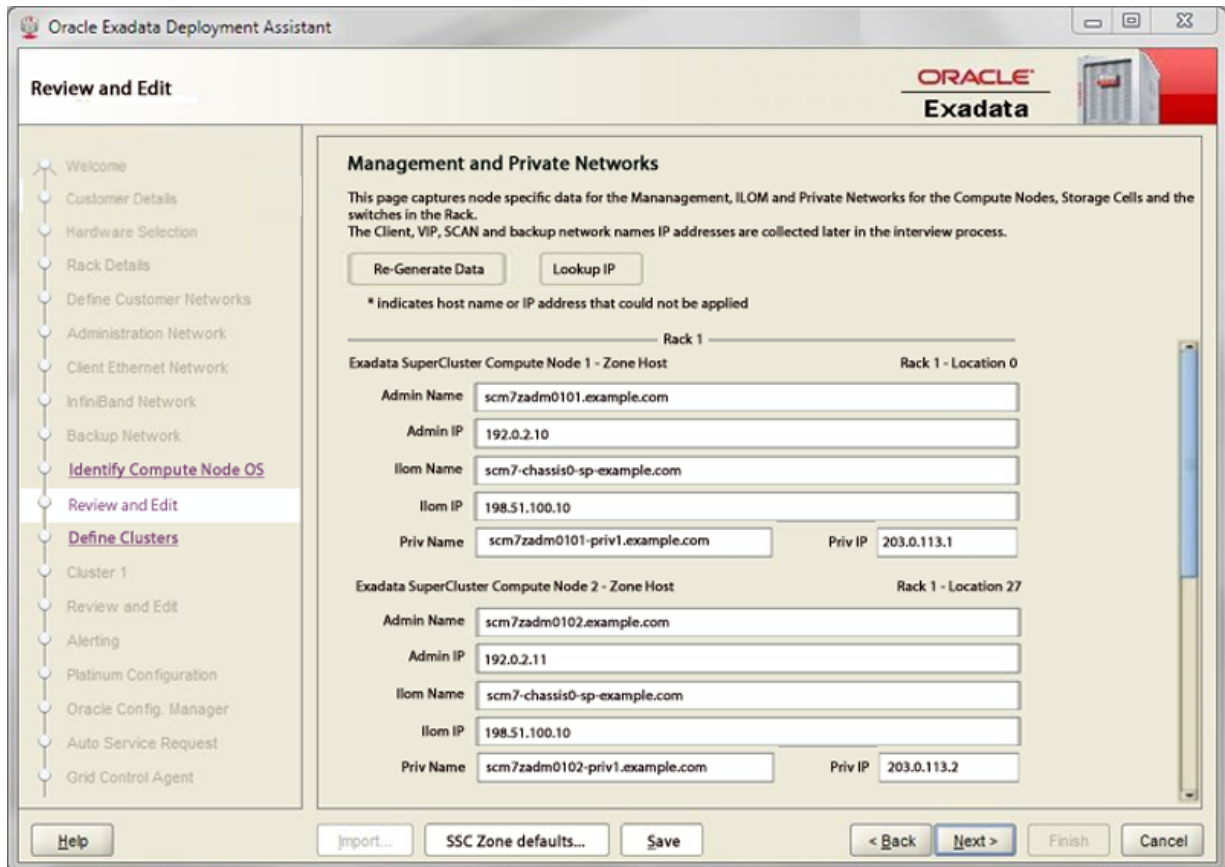
1. **Review the information in the Management and Private Networks page.**

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 your Oracle SuperCluster M8 or Oracle SuperCluster M7

- Switches and PDUs in your Oracle SuperCluster M8 or Oracle SuperCluster M7



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 62.

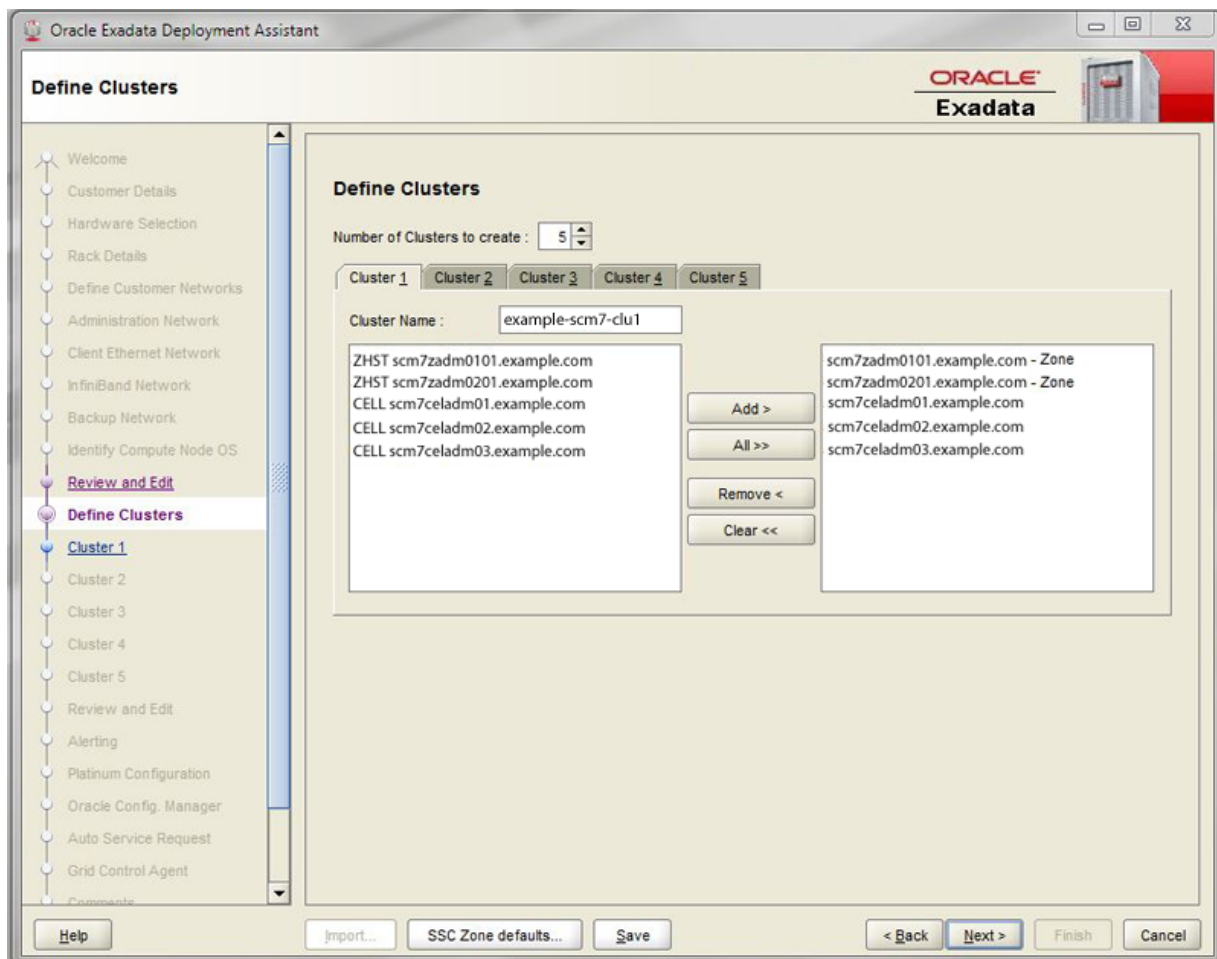
▼ Complete the Define Clusters Page

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 49:

- 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:

- scm7zadm0101.example.com
- scm7zadm0201.example.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.

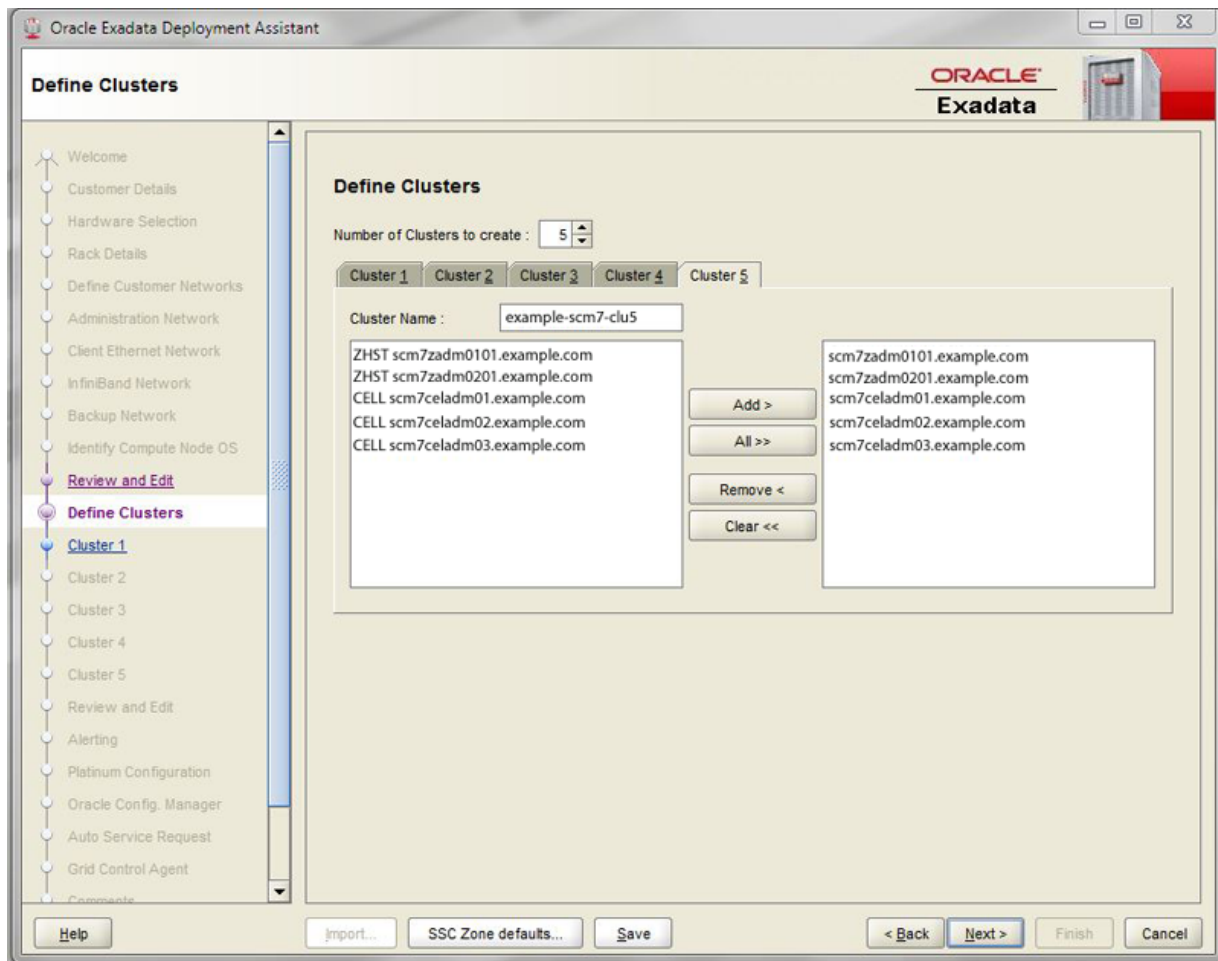
Because Cluster 5 is for the two Database Domains (without zones), that screen would show those two Database Domains and not the Database Domains with zones that are shown in Clusters 1 to 4.

In this example configuration, the screen for Cluster 5 shows the following information:

- scm7dbadm0102.example.com
- scm7dbadm0202.example.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 `example-scm7-clu5`, as shown in the following figure.



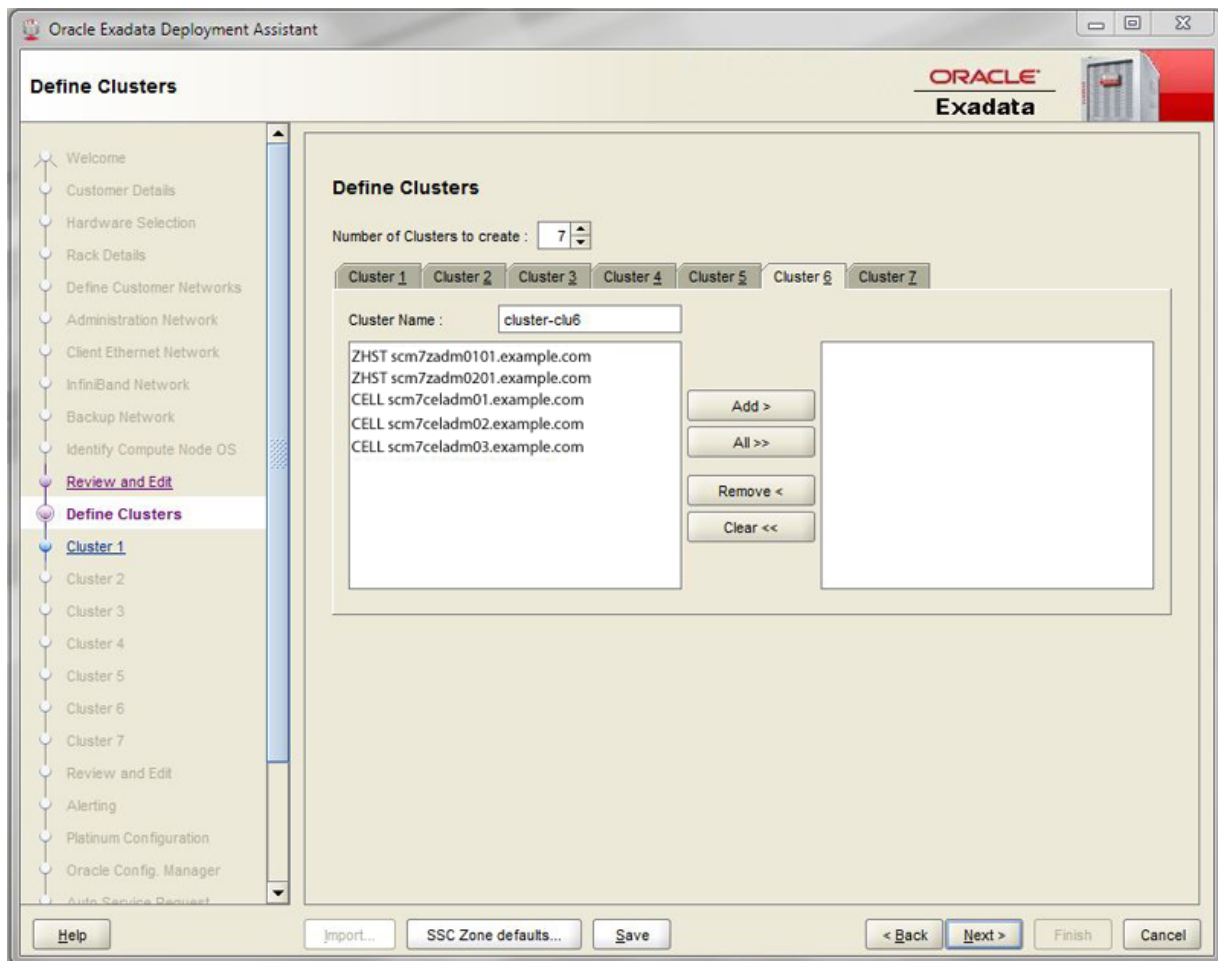
3. Determine the number of zones and clusters that you want to add to the Database Domains with zones.

Using the example configuration outlined in [“Example Configuration” on page 49](#), 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 `example-scm7-clu6`).

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

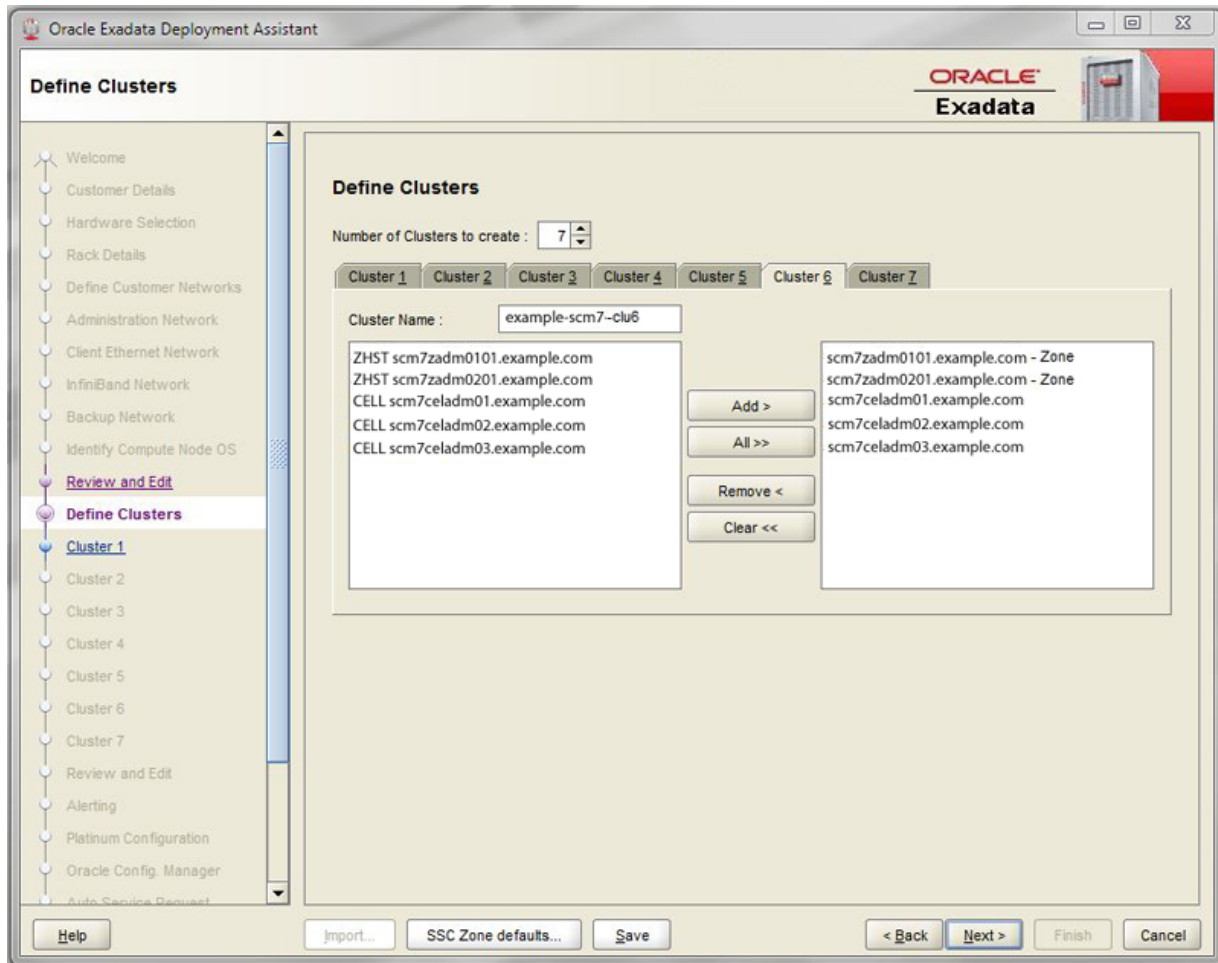
The Database Domains with zones are shown as ZHST in the left pane in the Define Clusters page. The Database Domains with zones that you selected appear in the right pane.

- 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 50](#) 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 70](#).

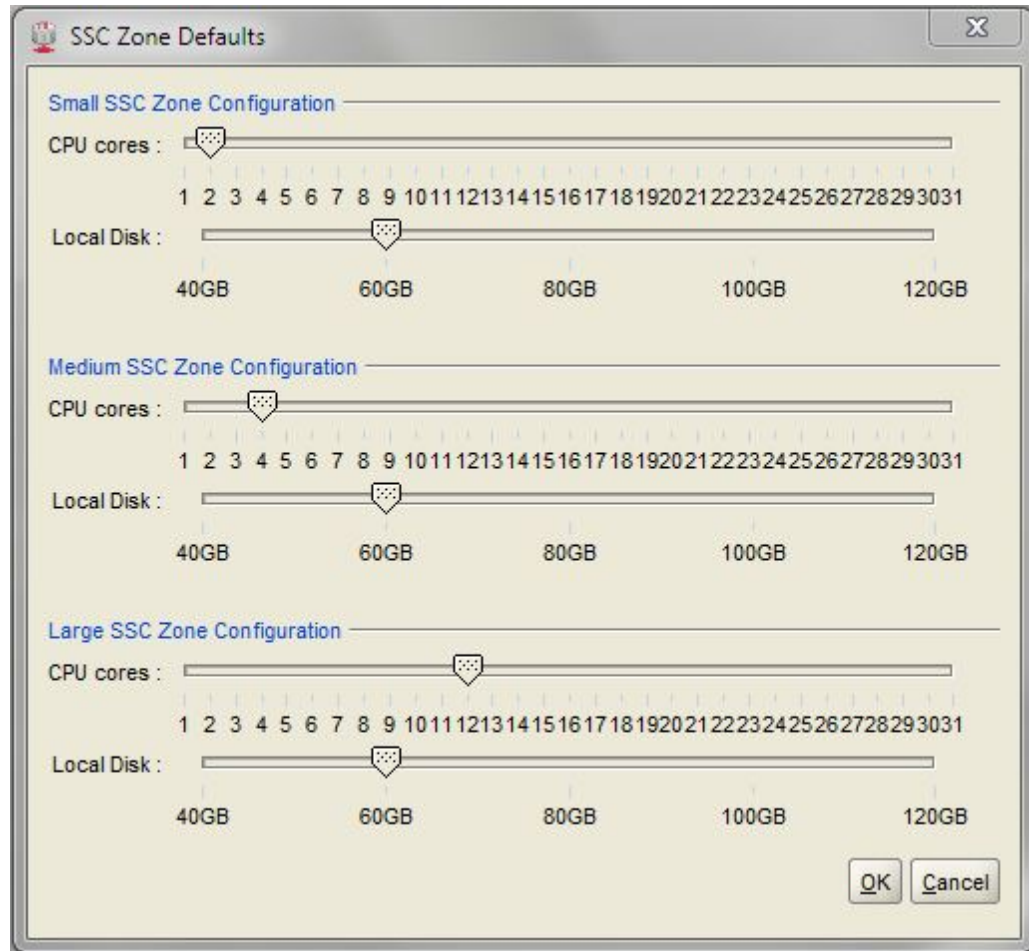
▼ Set Zone Default Configurations

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:
 - 2 cores
 - 60 GB of disk space

- Medium zone configuration:
 - 4 cores
 - 60 GB of disk space
- Large zone configuration:
 - 12 cores
 - 60 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 31 cores
 - **Amount of disk space:** Between 40 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 15](#) 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 73](#). 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 73](#).

▼ Complete the Cluster Page



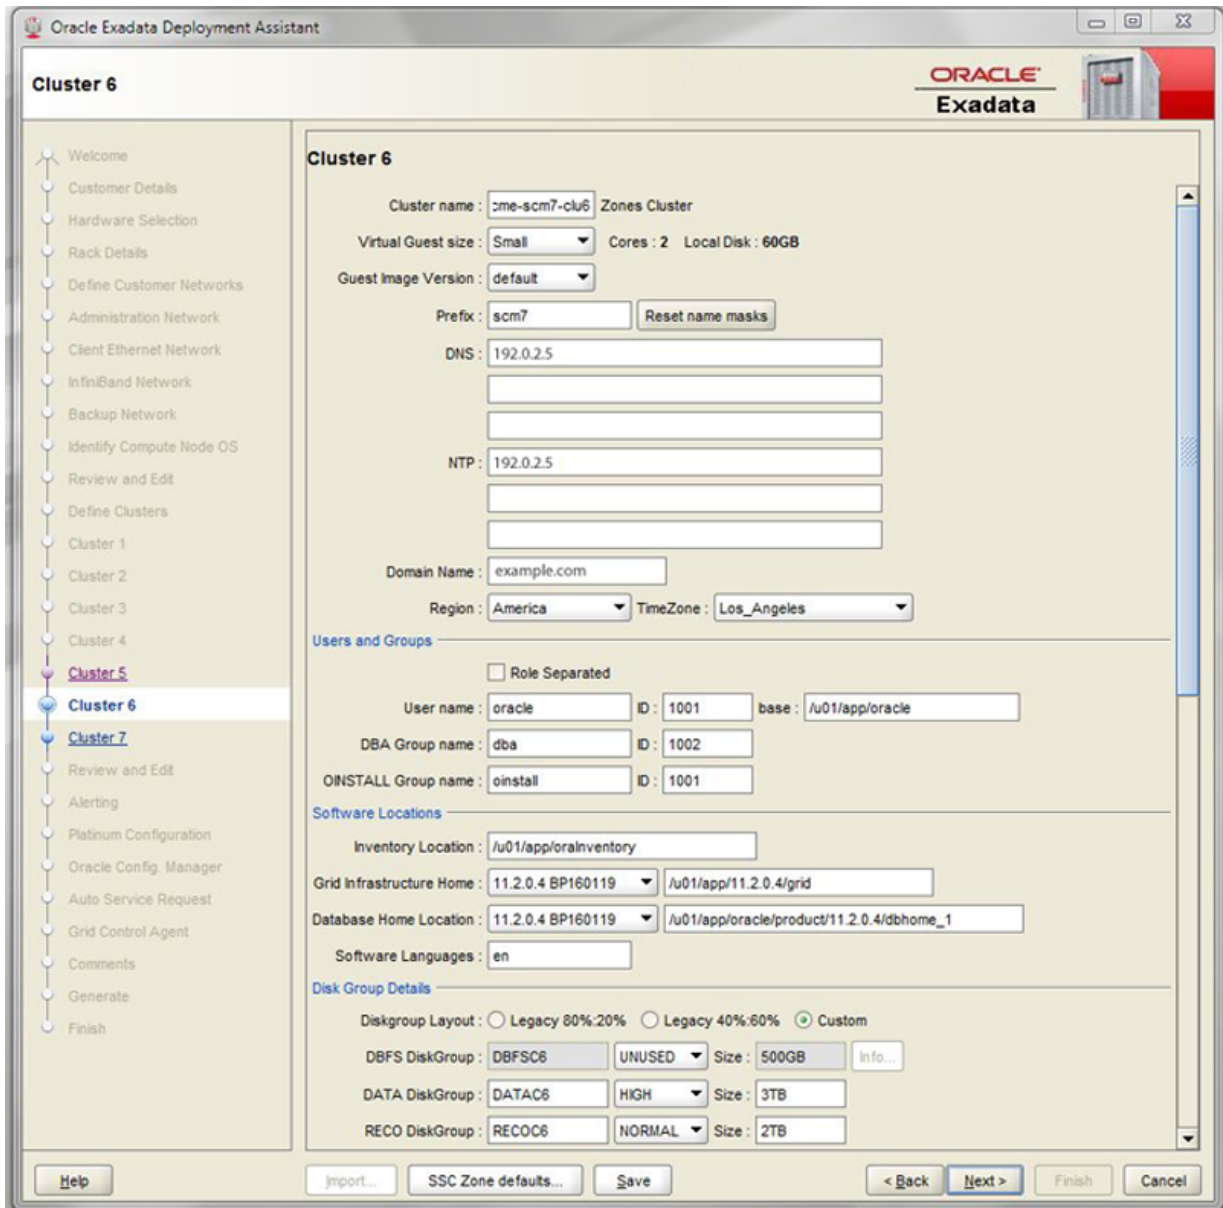
Caution - Take extreme care when specifying the DATA and RECO diskgroup names. The names must be unique for the cluster you are defining. Check existing griddisks on the storage cells and ensure they are not already in use. If an existing griddisk is specified for a new cluster, the griddisk will be 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 49](#), review the cluster configuration page for Cluster 6, which is the first new cluster that you are creating. Click the Next button to 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.



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 15](#) 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 70](#).

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 70](#), 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: DBFSC6
 - 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 50](#) 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.

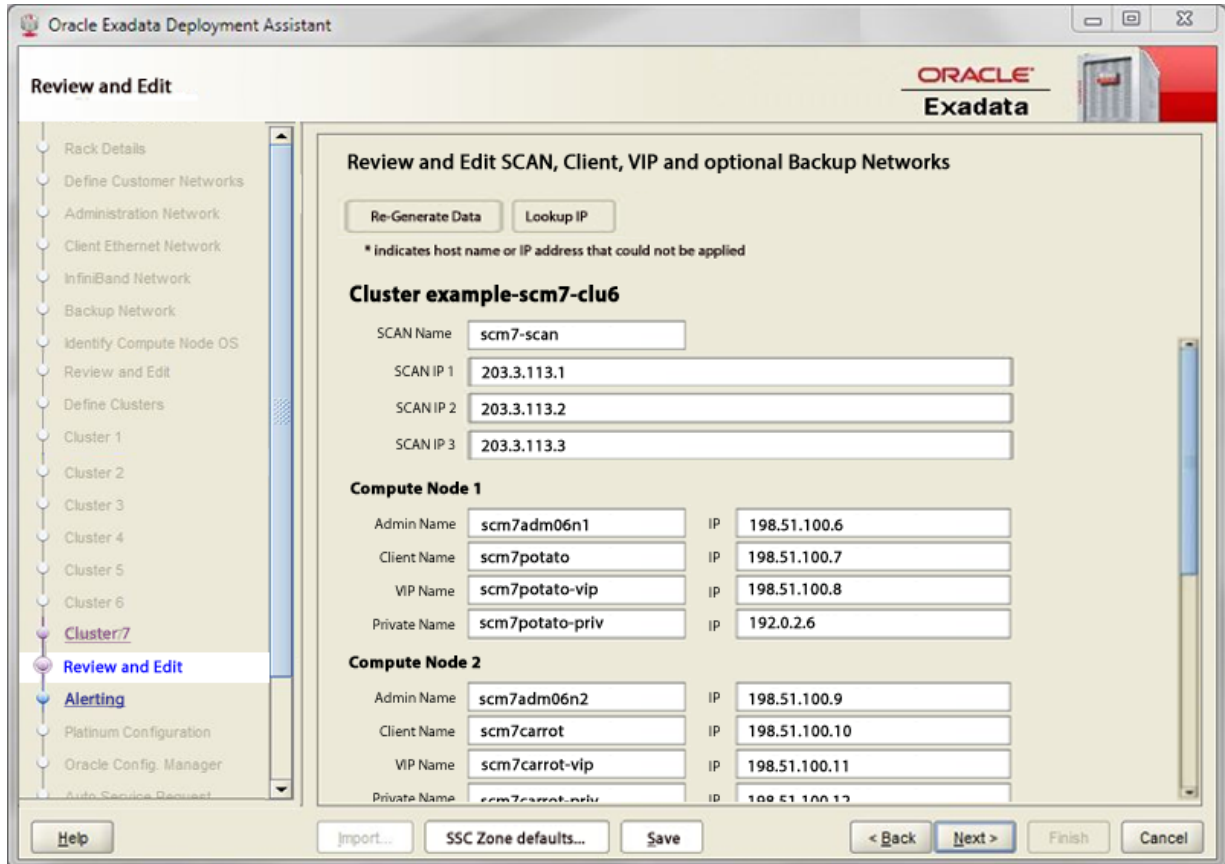
For example, if the first available starting IP address for the client network for the first new cluster is 172.16.8.150, and the Pool Size field next to the Start IP field for that network shows 7, then IP addresses 172.16.8.150 to 172.16.8.156 are assigned to the client access network for this cluster. 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 77.

▼ **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, 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 79.

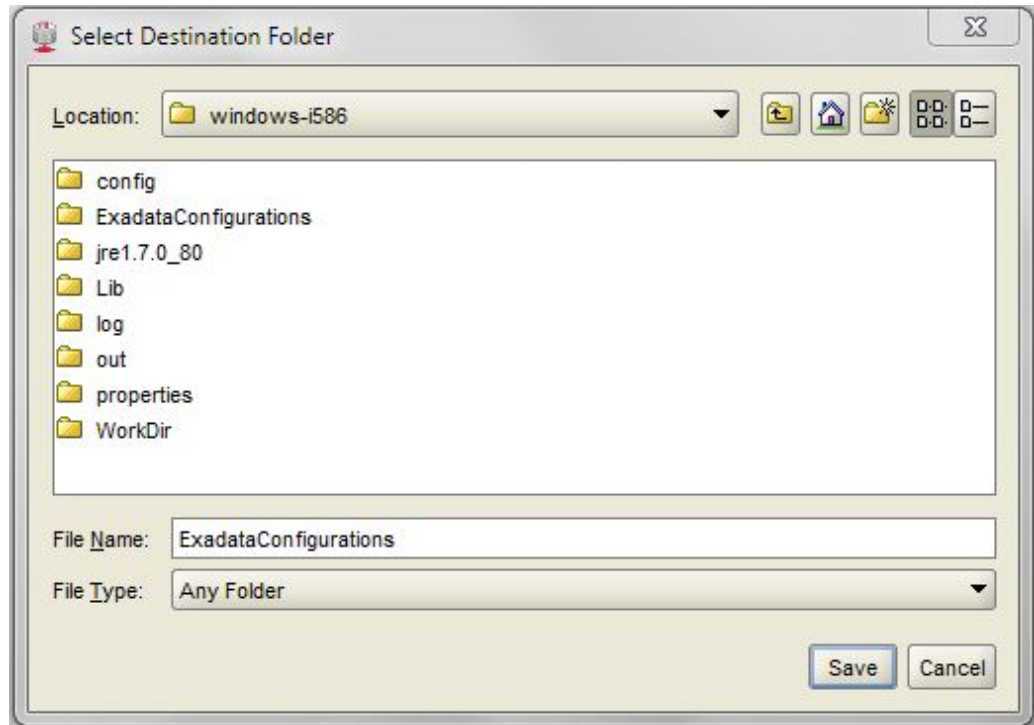
▼ Verify Remaining Configuration Information

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 79.

▼ 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 83.

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 83](#)
- [“Delete a Template Zone From a Database Domain” on page 85](#)

▼ Create a Template Zone on a Database Domain

1. **Log in to the first Database Domain in your Oracle SuperCluster M8 or Oracle SuperCluster M7 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, *scm7dbadmin01-01_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 scm7dbadmin01.us.example.com running
/zoneHome/scm7dbadmin01.us.example.com solaris excl
- scm7dbadmin01-01_T installed /zoneHomeT/scm7dbadmin01-01_T solaris excl
```

You do not have to create a template zone on the Database Domain in this case. Go to [“Running the `install.sh` Script” on page 87](#).

- 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 46](#) 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/
scm7dbadmin01-01-20121129093659.log and traces in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
scm7dbadmin01-01-20121129093659.trc and o/p in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
scm7dbadmin01-01-20121129093659.out
INFO: Begin Template Creation
INFO: Checking for scm7dbadmin01-01_T
INFO: Checking for scm7dbadmin01-01_T
INFO: Creating a template Virtual Guest for future use
INFO: 2012-11-29 09:37:14 : Executing zfs create rpool/scm7dbadmin01-01_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/scm7dbadmin01-01_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 : scm7dbadmin01-01_T Data file :: none

The name for the template zone for this particular domain is *DB-domain_T* (for example, scm7dbadmin01-01_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 46](#) 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
-  scm7dbadmin01-01_T installed /zoneHomeT/scm7dbadmin01-01_T solaris excl
```

The template zone is shown as *DB-domain_T* (for example, in the example output above, scm7dbadmin01-01_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.

```
./ssc_exavm -destroy -template
INFO: Logging all actions in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
scm7dbadmin01-01-20121130163540.log and traces in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
scm7dbadmin01-01-20121130163540.trc and o/p in /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/
scm7dbadmin01-01-20121130163540.out
INFO: Destroying Template Virtual Guest
INFO: Destroying Template ...
INFO: Destroying Virtual Guests ...
INFO: 2012-11-30 16:35:40 : Executing zoneadm -z scm7dbadmin01-01_T uninstall -F

INFO: Executing /opt/oracle.supercluster/zonetools/ssc_exavm/tmp/vmtrash/24064.SysCall.0.cmd

INFO: Running System command ...
R R R R R R R R R R R R R R R R R R R R R R R R
SUCCESS: Done
INFO: 24 : Completed Command Execution
INFO: 2012-11-30 16:36:04 : Completed Command Execution
INFO: 2012-11-30 16:36:04 : Executing sleep 5

... (OUTPUT ABBREVIATED) ...

INFO: Running System command ...
R
SUCCESS: Done
INFO: 1 : Completed Command Execution
INFO: 2012-11-30 16:36:10 : Completed Command Execution
```

Running the `install.sh` Script

These topics describe how to run the `install.sh` script.

- [“Set Up Public Key Authentication for ZFS Storage Controllers” on page 87](#)
- [“Set Up the `install.sh` Script” on page 88](#)
- [“Run the `install.sh` Script” on page 90](#)
- [“NTP Issue” on page 92](#)

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

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

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

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

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

1. Ensure that the machine running the OEDA has Oracle JRE 1.7 or later.
2. Locate the `workDir` directory.

By default, the `WorkDir` directory is located in the directory where you downloaded the OEDA patch, as described in [“Preparing to Configure Zones on Database Domains” on page 35](#).

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 `install.sh` script.

By default, the `install.sh` script is located in the directory where you downloaded the OEDA patch, as described in [“Preparing to Configure Zones on Database Domains” on page 35](#).

5. Change directories to the following directory on the Database Domain where you will be running the `install.sh` script:

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

6. Verify that the `iscsi.zpool.sh` script is in the `/opt/oracle.supercluster/bin` directory.

- If you see the `iscsi.zpool.sh` script in the `/opt/oracle.supercluster/bin` directory, go to [Step 7](#).
- If you do not see the `iscsi.zpool.sh` script in the `/opt/oracle.supercluster/bin` directory:

a. Locate the `iscsi.zpool.sh` script.

See [“Verify Configuration Tool Installation” on page 46](#) for the location of that script.

b. Copy the `iscsi.zpool.sh` script into the `/opt/oracle.supercluster/bin` directory on the Database Domain where you will be running the `install.sh` script.

7. Run the `install.sh` script.

Go to [“Run the `install.sh` Script” on page 90](#).

▼ Run the `install.sh` Script



Caution - Ensure that you backup all existing databases before running the `install.sh` script.



Caution - You must use the latest version of the OEDA 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 the `install.sh` script 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 the `install.sh` script 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 49](#) would be:

- `example-scm7-cluster-clu6.xml`
- `example-scm7-cluster-clu7.xml`

2. **Using the first new cluster configuration file, display the list of functions that will be performed by the `install.sh` script:**

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

For example:

```
./install.sh -cf /path-to-config-file/example-scm7-cluster-clu6.xml -l
```

Information similar to the following is displayed:

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

In the next step, you will run the appropriate steps in the `install.sh` script. **Do not run any of the following steps in the `install.sh` script in the next step:**

- Calibrate Cells (step 6 in the example output above)
- Create Cell Disks (step 7 in the example output above)
- Configure Cell Alerting (not shown in the example output above)

4. Run the appropriate steps in the `install.sh` script, 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 example-scm7-cluster-clu6.xml -s 1
```

Run all of the steps in the `install.sh` script **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 example-scm7-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 85](#) if necessary (for example, if the `install.sh` script 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 93](#).

NTP Issue

When your Oracle SuperCluster M8 or Oracle SuperCluster M7 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 93](#)

▼ 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. **Identify the first and second IB networks and net IDs.**

In the following example,

- `stor_ipmp0_0` is the first IB network, and `net22` is the *first_net_ID*.
- `stor_ipmp0_1` is the second IB network, and `net23` is the *second_net_ID*.

```
# dladm show-part | grep stor_ipmp0_
stor_ipmp0_0 8503 net22      up      f---
stor_ipmp0_1 8503 net23      up      f---
```

3. **Add links to the zones in this Database Domains:**

```
# zonecfg -z zonename
zonecfg:zonename> add anet
zonecfg:zonename:anet> set linkname=first_net_ID
zonecfg:zonename:anet> set lower-link=first_net_ID
zonecfg:zonename:anet> set pkey=8503
zonecfg:zonename:anet> end
zonecfg:zonename> add anet
zonecfg:zonename:anet> set linkname=second_net_ID
zonecfg:zonename:anet> set lower-link=second_net_ID
```

```
zonecfg:zonename:anet> set pkey=8503
zonecfg:zonename:anet> end
zonecfg:zonename> exit
```

where:

- *zonename* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone
- *first_net_ID* is the first network ID shown for *stor_ipmp0_0* from [Step 2](#)
- *second_net_ID* is the second network ID shown for *stor_ipmp0_0* from [Step 2](#)

For example:

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

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

```
# zoneadm -z zonename reboot
```

For example:

```
# zoneadm -z scm7zdbadm02 reboot
```

5. Log in to the zone:

```
# zlogin zonename
```

For example:

```
# zlogin scm7zdbadm02
```

6. Create the IP addresses in the zone for the new links:

```
# ipadm create-ip first_net_ID
# ipadm create-ip second_net_ID
```

where:

- *first_net_ID* is the first network ID from [Step 2](#)
- *second_net_ID* is the second network ID from [Step 2](#)

For example:

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

7. 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
stor_ipmp0/v4	static	ok	192.168.28.2/22
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.

8. Create the new IPMP group inside the zone:

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

where:

- *first_net_ID* is the first network ID shown in the output from [Step 2](#)
- *second_net_ID* is the second network ID 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 7](#)

For example:

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

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

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

where *second_net_ID* is the second network ID shown in the output from [Step 2](#).

For example:

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

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

```
# ipmpstat -g
```

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

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

11. Repeat [Step 3](#) through [Step 10](#) 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 97](#)

▼ 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 zonename halt
```

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

```
# zoneadm -z scm7zdbadm02 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 zonename
zonecfg: zonename> select anet linkname=net_device#1
zonecfg: zonename:anet> set vlan-id=VLAN_ID
zonecfg: zonename:anet> set lower-link=net_device#1
zonecfg: zonename:anet> end
zonecfg: zonename> select anet linkname=net_device#2
zonecfg: zonename:anet> set vlan-id=VLAN_ID
zonecfg: zonename:anet> set lower-link=net_device#2
zonecfg: zonename:anet> end
zonecfg: zonename> verify
```

```
zonecfg: zonename> commit
zonecfg: zonename> exit
```

where:

- *zonename* 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 scm7zdbadm02
zonecfg: scm7zdbadm02> select anet linkname=net4
zonecfg: scm7zdbadm02:anet> set vlan-id=13
zonecfg: scm7zdbadm02:anet> set lower-link=net4
zonecfg: scm7zdbadm02:anet> end
zonecfg: scm7zdbadm02> select anet linkname=net7
zonecfg: scm7zdbadm02:anet> set vlan-id=13
zonecfg: scm7zdbadm02:anet> set lower-link=net7
zonecfg: scm7zdbadm02:anet> end
zonecfg: scm7zdbadm02> verify
zonecfg: scm7zdbadm02> commit
zonecfg: scm7zdbadm02> exit
```

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

```
# zoneadm -z zonename reboot
```

For example:

```
# zoneadm -z scm7zdbadm02 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

D

- determining
 - amount of space for each disk group, 24
 - available storage capacity, 22
 - maximum amount of available space for RACs, 23
 - number of RACs, 23
- disk groups
 - considerations when planning layout, 20
 - determining amount of disk space, 24
 - planning layout, 21
 - recording planned layout, 33

E

- Exadata Storage Servers
 - considerations when planning disk group layout, 20
 - example layout scenario, 24
 - planning disk group layout, 21
 - recording planned disk group layout, 33
 - rules and restrictions for disk group layout, 20

I

- iscsi-lun.sh script, locating, 47

O

- Oracle Exadata Deployment Assistant
 - creating the configuration files, 54, 88
 - downloading, 46

- Oracle OneCommand utility, downloading, 46

P

- PDomains
 - overview, 14
 - preparing to set up zones, 35

R

- RACs
 - determining maximum allowable disk space, 23
 - determining number of RACs, 23
 - planning number of zones and RACs, 16
 - recording planned configuration, 32
 - rules and recommendations, 19
- recording
 - planned Exadata Storage Server disk group layout, 33
 - planned RAC configuration, 32
 - planned zone configuration, 26
- repository
 - determining location of, 39
 - installing or updating files from local repository, 42
 - installing or updating files from remote repository, 40
- restrictions
 - Exadata Storage Server disk group layout, 20
 - setting up zones and RACs, 19
- rules
 - Exadata Storage Server disk group layout, 20
 - setting up zones and RACs, 19

S

- ssc_exavm tool
 - creating template zone, 84
 - deleting template zone, 85

T

- template zone
 - creating, 84
 - deleting, 85
 - understanding, 83

Z

- zonemanifest.sh script, locating, 47
- zones
 - number allowed per domain, 18
 - planning number of zones and RACs, 16
 - recording planned zone configuration, 26
 - rules and recommendations, 18