SPARC SuperCluster T4-4 Zones With Oracle® Database on Database Domains

Configuration Guide

Part No.: E38377-04
October 2014

Please
Recycle

Adobe PostScript

Contents

Using This Documentation

This document describes how to set up zones and Real Application Clusters (RACs) on the Database Domains in the SPARC SuperCluster T4-4.

Product Notes

For late-breaking information and known issues about this product, refer to the product notes, which can be viewed through a browser by viewing the following directory on the first SPARC T4-4 server installed in the SPARC SuperCluster T4-4:

```
/opt/oracle/node/doc/E21659_01/index.html
```

Related Documentation

Documentation	Links
All Oracle products	http://www.oracle.com/documentation
Oracle Solaris OS and systems software library	http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_swhttp://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw

Feedback

Provide feedback about this documentation at:

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

Access to Oracle Support

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

Planning to Set Up Zones on Database Domains

Once the SPARC SuperCluster T4-4 is set up, you can modify it to deploy zones in the Application Domains and Database Domains. These topics describe how to plan for configuring zones on Database Domains.

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

- For Oracle Solaris 10, refer to *Oracle Solaris Administration: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones* at:

  `http://docs.oracle.com/cd/E23823_01/index.htmlhttp://download.oracle.com/docs/cd/E23823_01`

- For Oracle Solaris 11, refer to *Oracle Solaris Administration: Network Interfaces and Network Virtualization* and *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* at:

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

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

Important Cautions

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

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

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

Naming Conventions

This document uses these naming conventions:

- `ssc`*number* for the SPARC SuperCluster T4-4
- `r`*number* for the cluster number
- `cn`*number* for the SPARC T4-4 server (compute node) number
- `db`*number* for the Database Domain number
- `z`*number* for the zone number

Understanding Domain Configurations

The SPARC SuperCluster T4-4 has either two SPARC T4-4 servers (Half Rack) or four SPARC T4-4 servers (Full Rack). Each SPARC T4-4 servers will have from 1 to 4 domains, where the domains are either Database Domains or Application Domains running either Oracle Solaris 10 or Oracle Solaris 11.

At some point later in these procedures, you will need to know the number of IB HCAs that are associated with the Database Domain that you are adding zones to. The following table provides this information.

TABLE: IB HCAs Associated With Domains

	Number of IB HCAs Associated with This Domain			
	First Domain	**Second Domain**	**Third Domain**	**Fourth Domain**
One Total Domain	4	N/A	N/A	N/A
Two Total Domains	2	2	N/A	N/A
Three Total Domains	2	1	1	N/A
Four Total Domains	1	1	1	1

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 3
- "Memory Available for Domains" on page 4

Cores Available for Domains and Zones

Each SPARC T4-4 server has four processors, and each processor has eight cores, for a total of 32 cores per SPARC T4-4 server.

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

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

The following table provides information on the number of cores that are available for specific domains and zones that are created within those domains. See "Understanding Domain Configurations" on page 3 for more information on the different types of domains.

TABLE: Cores Available for the Domains and Zones

	Total Number of Cores Available for This Domain	Number of Cores Set Aside for Global Zone	Number of Cores Available for Nonglobal Zones
4 IB HCA Domain	32	4	28
2 IB HCA Domain	16	4	12
1 IB HCA Domain	8	2	6

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

For each zone that you create, use a *minimum* of one core per zone. However, depending on the workload that you expect on a zone, a larger number of cores per zone might be preferable, thereby reducing the total number of zones on each SPARC T4-4 server. Carefully consider the expected workload on each zone that you create, so that you allot the appropriate number of cores to those zones.

Memory Available for Domains

Each SPARC T4-4 server has a total of 64 DIMM slots. The following memory sizes are available for SPARC SuperCluster T4-4:

- 4 GB DIMMs (total memory of 256 GB for each SPARC T4-4 server)

- 8 GB DIMMs (total memory of 512 GB for each SPARC T4-4 server)
- 16 GB GIMMs (total memory of 1024 GB for each SPARC T4-4 server)

Zones and Cluster Planning for Database Domains

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

- The size of the Database Domain, which will determine the number of cores available for that domain and the maximum number of zones that can be created in that Database Domain
- The number of zones that you want to create in that Database Domain
- The number of cores that you want to assign to each zone in that Database Domain, keeping in mind that two or four cores will always be set aside for the global zone (the Database Domain)
- The clusters that are used for the zones on the Database Domains, where:
  - Each zone that is a member of that cluster should use the same number of cores.
  - Clusters should be created where the zones that are members of a cluster should be on separate Database Domains, and the cluster spans across both SPARC T4-4 servers so that there is not a single point of failure.

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

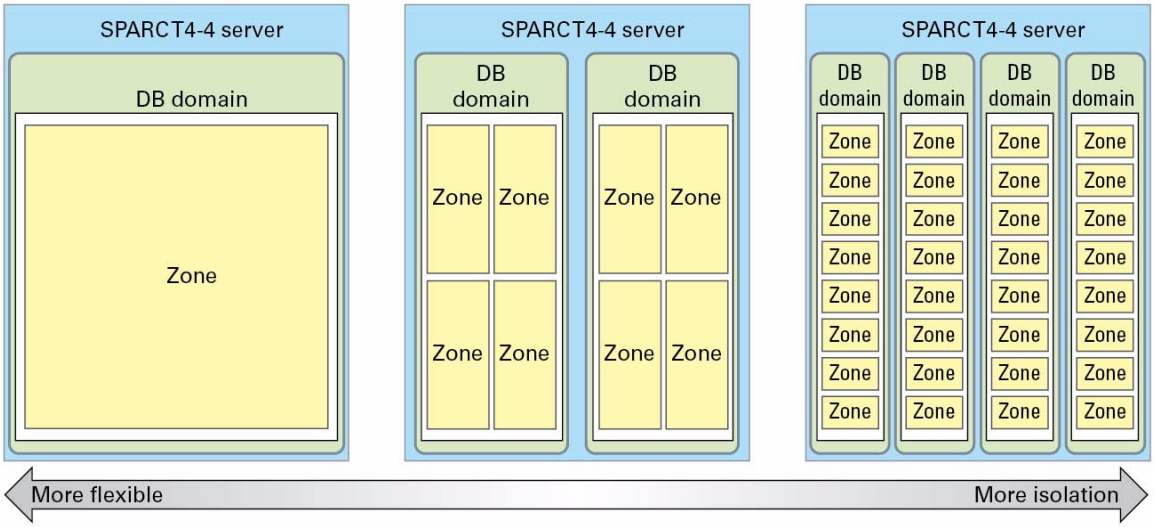

Guidelines for Planning the Number of Zones and Clusters

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

- When you first install the operating system instances on a domain, that domain is automatically designated as the global zone. Any 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 four-node cluster, where there is one zone on each Database Domain on the four SPARC T4-4 servers, and those zones are clustered together, even if two of the Database Domains take up 25% of the servers and the other two Database Domains take up 50% of the servers. 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.
- When creating zones on a Database Domain, if the last zone that you are creating on that domain uses the last remaining core available in that domain, then that zone will be configured with no CPU resource controls. For example, if you are creating zones on a server with four total domains, so that you could have a maximum of six single-core zones on each domain, the sixth and final zone in the domain would be configured with no CPU resource controls. This zone would share the core and vCPUs with the global zone in this case.
- To minimize remote memory accesses, you should allocate physical memory to each of the zones in multiples of 16 GB to align with the size of the memory banks on the server.
- While it is possible to create a cluster where all the zones are in the same Database Domain or are on the same SPARC 4-4 server, that configuration is not recommended because it is not a highly-available configuration. It is advisable to create clusters where the zones are on separate Database Domains, and the cluster spans across several SPARC T4-4 servers, 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 you have three Exadata Storage Servers 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 you do not have to manually partition the Exadata Storage Servers. For high redundancy, you should have five Exadata Storage Servers available for every cluster with zones on the Database Domains.

  For example, if you have a SPARC SuperCluster T4-4 Half Rack, which has three Exadata Storage Servers, and 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, if you have a SPARC SuperCluster T4-4 Full Rack (with six Exadata Storage Servers) and an Oracle Exadata Storage Expansion Full Rack connected to your SPARC SuperCluster T4-4 (with 18 Exadata Storage Servers), and 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, and 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 8
- "Understanding an Example Scenario" on page 13

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 (ssc01r1 and ssc01r2), then you might use these names for the three disk groups for each cluster:

   - First cluster (ssc01r1):
     - DATA_ssc01r1
     - RECO_ssc01r1
     - DBFS_ssc01r1
   - Second cluster (ssc01r2):
     - DATA_ssc01r2
     - RECO_ssc01r2
     - DBFS_ssc01r2

Record your Exadata Storage Server disk group layout in the worksheet in "Recording Your Exadata Storage Server Disk Group Layout" on page 36.

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

- "Determine the Available Existing Storage Capacity" on page 10
- "Determine the Number of Clusters in the System and Maximum Amount of Space Available to Each Cluster" on page 12
- "Determine the Amount of Space for Each Disk Group in Each Cluster" on page 13

Determine the Available Existing Storage Capacity

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

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

- High-capacity drives – 3 TB
- High-performance drives – 600 GB

In addition, the total number of Exadata Storage Servers varies, depending on the type of SPARC SuperCluster T4-4 and Exadata Storage Expansion Rack that you have:

Rack Type	Quarter Rack	Half Rack	Full Rack
SPARC SuperCluster T4-4	N/A	4	8
Oracle Exadata Storage Expansion Rack	4	9	18

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

The following tables provide capacity information for the SPARC SuperCluster T4-4 and the Oracle Exadata Storage Expansion Rack.

TABLE: Exadata Storage Server Storage Capacity in the SPARC SuperCluster T4-4

	Half Rack		Full Rack	
Capacity Type	**HP Disks**	**HC Disks**	**HP Disks**	**HC Disks**
Raw disk capacity	21.6 TB	108 TB	43.2 TB	216 TB
Usable mirrored capacity (ASM normal redundancy)	9.5 TB	48 TB	19 TB	94.5 TB
Usable triple mirrored capacity (ASM high redundancy)	6.5 TB	32.3 TB	13 TB	64.5 TB
Raw flash capacity	1.1 TB		2.2 TB	

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

Capacity Type	**3 TB Disks (X3-2)**	**4 TB Disks (X4-2)**
Raw disk capacity	• Quarter Rack: 144 TB • Half Rack: 324 TB • Full Rack: 648 TB	• Quarter Rack: 192 TB • Half Rack: 432 TB • Full Rack: 864 TB
Usable mirrored capacity (ASM normal redundancy)	• Quarter Rack: 64 TB • Half Rack: 144 TB • Full Rack: 288 TB	• Quarter Rack: 85.3 TB • Half Rack: 192 TB • Full Rack: 384 TB
Usable triple mirrored capacity (ASM high redundancy)	• Quarter Rack: 43 TB • Half Rack: 97 TB • Full Rack: 193.5 TB	• Quarter Rack: 57.3 TB • Half Rack: 129 TB • Full Rack: 384 TB
Raw flash capacity	• Quarter Rack: 6.4 TB • Half Rack: 14.4 TB • Full Rack: 28.8 TB	• Quarter Rack: 12.8 TB • Half Rack: 28.8 TB • Full Rack: 57.6 TB

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

Capacity Type	600 GB Disks (X3-2)	1.2 TB Disks (X4-2)
Raw disk capacity	• Quarter Rack: 28 TB • Half Rack: 64 TB • Full Rack: 128 TB	• Quarter Rack: 57 TB • Half Rack: 129 TB • Full Rack: 258 TB
Usable mirrored capacity (ASM normal redundancy)	• Quarter Rack: 13 TB • Half Rack: 29 TB • Full Rack: 58 TB	• Quarter Rack: 26 TB • Half Rack: 58 TB • Full Rack: 116 TB
Usable triple mirrored capacity (ASM high redundancy)	• Quarter Rack: 8 TB • Half Rack: 19 TB • Full Rack: 39 TB	• Quarter Rack: 16 TB • Half Rack: 39 TB • Full Rack: 78 TB
Raw flash capacity	• Quarter Rack: 6.4 TB • Half Rack: 14.4 TB • Full Rack: 28.8 TB	• Quarter Rack: 12.8 TB • Half Rack: 28.8 TB • Full Rack: 57.6 TB

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

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

- The logical divisions for the databases in your system, where clusters that are created, and zones that are members of those clusters, might have different functions or audiences.
- The number of zones that are available and the number of clusters that can be created with those zones as members, which is partly determined by several factors, such as the number of Database Domains on each SPARC T4-4 server, the size of each Database Domain, and so on.

See "Zones and Cluster Planning for Database Domains" on page 5 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 how you might decide on the Exadata Storage Server layout partly based on the amount of existing storage capacity that you have 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

Consider the following example scenario, where you have the following:

- One SPARC SuperCluster T4-4 Half-Rack system, with three Exadata Storage Servers
- Twelve high-capacity (3 TB) disk drives in each Exadata Storage Server
- Each Exadata Storage Server is set at normal (double-mirrored) redundancy

Using the information from the section "Determine the Available Existing Storage Capacity" on page 14, you would see that you have **48 TB** of usable storage capacity available to you for the Exadata Storage Servers in your SPARC SuperCluster T4-4.

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

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

If you have 48 TB of usable disk capacity available to you, and you want to create two clusters for the zones that you will be creating, you could use the existing 48 TB of space in one of several ways:

- Use the entire 48 TB of existing space for the two clusters, where you could:
  - Create two equal-sized clusters, where the total amount of space used by the disk groups in each cluster would be 24 TB, or
  - Create two clusters of different sizes, where the total disk space used by the disk groups for the two clusters equals 48 TB (for example, one cluster with the total disk space used by the disk groups equaling 12 TB and the other cluster with the total disk space used by the disk groups equaling 36 TB)

  Keep in mind that by using the entire 48 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 48 TB of usable space for the two clusters that you will be creating, while reserving some of the usable space for future expansion.

For this example scenario, we will create two equal-sized clusters, so when we begin allotting disk sizes to each disk group within each cluster, we must verify that the total amount of space used by all three disk groups does not exceed 24 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 will be using for the disk groups for each cluster is roughly 3.2 TB, which is far below the maximum amount of 24 TB that is available for each cluster based on the calculations from the previous section.

Determine the Grid Disk Sizes

Finally, determine how you will use the grid disks to build up each disk group to the appropriate size.

As mentioned earlier in this example scenario, we will be using three Exadata Storage Servers, with twelve disk drives in each Exadata Storage Server, for the clusters in this system. We will also have to share the three Exadata Storage Servers across each cluster because we have fewer than three Exadata Storage Servers available for each cluster (see "Planning the Exadata Storage Server Disk Group Layout" on page 8 for more information). Therefore, we will be spanning across 36 disk drives total for the DATA and RECO disk groups, so you would create 36 grid disks total for those disk groups. As mentioned earlier, the DBFS disk group uses ten of the twelve disks in each Exadata Storage Server, so that disk group would span across 30 disk drives total, and you would create 30 grid disks total for that disk group.

You could therefore use the following values for the grid disk sizes for each disk group:

- DATA – Dividing the 2 TB of space that you want for the DATA disk group between the 36 grid disks amounts to slightly more than 55 GB for each grid disk:

  2000 GB / 36 = 55.56

Therefore, allotting 55 GB to each grid disk will provide roughly 2 TB of usable space for the DATA disk group for each cluster (1.98 TB).

- RECO – Dividing the 1 TB of space that you want for the RECO disk group between the 36 grid disks amounts to slightly less than 28 GB for each grid disk:

  1000 GB / 36 = 27.78

  Therefore, allotting 28 GB to each grid disk will provide roughly 1 TB of usable space for the RECO disk group for each cluster (1.008 TB),

- DBFS – Dividing the 200 GB of space that you want for the DBFS disk group between the 30 grid disks available for that disk group amounts to slightly less than 7 GB for each grid disk:

  200 GB / 30 = 6.66

  Therefore, allotting 7 GB to each grid disk will provide roughly 200 GB of usable space for the DBFS disk group for each cluster (210 GB).

Recording Your Existing and Planned Configuration Information

Use the following worksheets to record the existing configuration that you have for your SPARC SuperCluster T4-4, which will affect the number and size of the zones that you can create. You will also use the following worksheets to record the decisions that you made for the number of zones and clusters that you will be deploying, and the layout that you will use for the Exadata Storage Servers.

- "Recording Your Existing Configuration" on page 16
- "Recording Your Zone Configuration Information" on page 20
- "Recording Your Cluster Configuration Information" on page 34
- "Recording Your Exadata Storage Server Disk Group Layout" on page 36

Recording Your Existing Configuration

Several zone setup decisions and procedures are affected by the configuration choices that you made when you first purchased your SPARC SuperCluster T4-4. Record your existing configuration in the following tables, and use this information as a reference for certain zone setup decisions and procedures.

For each SPARC T4-4 server, provide the following information:

- The total number of domains, both Application Domains and Database Domains, that you have on that server. This information will be used to determine the maximum number of cores that you can use for each zone that you create on the Database Domains on that server.

  For example, if you have four total domains on a SPARC T4-4 server, then the maximum number of cores that you can assign to each zone created on any of those domains would be six. See "Guidelines for Planning the Number of Zones and Clusters" on page 6 for more information.
- Whether each domain is an Application Domain or a Database Domain. This information will be used mainly to help you remember specifically which domains are Database Domains on each server.
- The specific location of each Database Domain (domain 1, domain 2, domain 3 or domain 4). This information will also be used to determine whether or not you must manually create VNETs for the host management network for certain Database Domains. See "Host Management Network After Zone Creation" on page 46 for more information on the situations where you must manually create VNETs for certain Database Domains.
- The 1 GbE host management host name and IP address for each domain.

The following sections provide worksheets to record your existing configuration:

SPARC T4-4 Server 1 Configuration Information

Configuration Information	Entry
How many *total* domains do you have on this server?	
What type of domains do you have (Application Domain or Database Domain) and what is the 1 GbE host management host name and IP address for each domain?	
Domain 1	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 2	Domain type: 1 GbE host management host name: 1 GbE host management IP address:

Configuration Information	Entry
Domain 3	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 4	Domain type: 1 GbE host management host name: 1 GbE host management IP address:

SPARC T4-4 Server 2 Configuration Information

Configuration Information	Entry
How many *total* domains do you have on this server?	
What type of domains do you have (Application Domain or Database Domain) and what is the 1 GbE host management host name and IP address for each domain?	
Domain 1	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 2	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 3	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 4	Domain type: 1 GbE host management host name: 1 GbE host management IP address:

SPARC T4-4 Server 3 Configuration Information (Full-Rack Only)

Configuration Information	Entry
How many *total* domains do you have on this server?	

Configuration Information	Entry
What type of domains do you have (Application Domain or Database Domain) and what is the 1 GbE host management host name and IP address for each domain?	
Domain 1	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 2	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 3	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 4	Domain type: 1 GbE host management host name: 1 GbE host management IP address:

SPARC T4-4 Server 4 Configuration Information (Full-Rack Only)

Configuration Information	Entry
How many *total* domains do you have on this server?	
What type of domains do you have (Application Domain or Database Domain) and what is the 1 GbE host management host name and IP address for each domain?	
Domain 1	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 2	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 3	Domain type: 1 GbE host management host name: 1 GbE host management IP address:
Domain 4	Domain type: 1 GbE host management host name: 1 GbE host management IP address:

Recording Your Zone Configuration Information

Use the tables in the following sections to record the following information for each SPARC T4-4 server:

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

Note that the worksheet for each SPARC T4-4 server has entries for 28 zones, which is the maximum number of zones that you can have on any Database Domain. However, you can have from 1 to 28 zones on each Database Domain, and the number of cores that you have allocated to those zones will vary, depending on the size of that zone. See "Guidelines for Planning the Number of Zones and Clusters" on page 6 for more information.

Also refer to the information that you entered in the "Recording Your Existing Configuration" on page 16 worksheets for each SPARC T4-4 server to determine the maximum number of zones that you can have on each Database Domain, and the maximum number of cores that are available for those zones, depending on the total number of domains on that server.

For example, if you have four domains on a server, then the maximum number of zones that you can have on any of the four domains would be six, and the maximum number of cores available for each of those domains would also be six. You could then choose any of the following zone configurations for each domain:

- One zone, consisting of six cores (48 vCPUs)
- Two zones, consisting of three cores (24 vCPUs) apiece
- Six zones, consisting of one core (8 vCPUs) apiece

Enter the zone configuration information in the following worksheets:

- "SPARC T4-4 Server 1 Zone Configuration Information" on page 21
- "SPARC T4-4 Server 2 Zone Configuration Information" on page 24
- "SPARC T4-4 Server 3 Zone Configuration Information (Full-Rack Only)" on page 28
- "SPARC T4-4 Server 4 Zone Configuration Information (Full-Rack Only)" on page 31

SPARC T4-4 Server 1 Zone Configuration Information

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
1	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
2	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
3	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
4	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
5	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
6	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
7	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
8	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
9	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
10	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
11	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
12	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
13	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
14	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
15	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
16	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
17	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
18	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
19	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
20	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
21	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
22	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
23	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
24	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
25	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
26	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
27	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
28	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

SPARC T4-4 Server 2 Zone Configuration Information

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
1	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
2	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
3	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
4	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
5	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
6	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
7	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
8	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
9	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
10	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
11	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
12	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
13	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
14	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
15	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
16	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
17	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
18	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
19	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
20	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
21	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
22	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
23	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
24	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
25	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
26	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
27	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
28	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

SPARC T4-4 Server 3 Zone Configuration Information (Full-Rack Only)

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
1	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
2	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
3	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
4	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
5	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
6	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
7	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
8	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
9	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
10	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
11	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
12	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
13	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
14	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
15	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
16	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
17	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
18	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
19	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
20	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
21	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
22	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
23	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
24	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
25	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
26	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
27	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
28	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

SPARC T4-4 Server 4 Zone Configuration Information (Full-Rack Only)

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
1	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
2	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
3	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
4	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
5	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
6	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
7	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
8	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
9	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
10	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
11	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
12	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
13	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
14	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
15	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
16	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
17	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
18	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
19	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
20	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
21	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
22	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Zone Number	Number of Cores and vCPUs Assigned to This Zone	Host Names	IP Addresses
23	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
24	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
25	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
26	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
27	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:
28	Cores: vCPUS (cores x 8):	Host management: Client access: InfiniBand: VIP:	Host management: Client access: InfiniBand: VIP:

Recording Your Cluster Configuration Information

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

- Type of cluster (two-node or four-node)
- Name of cluster (for example, ssc01r1 or ssc01r2)
- Zones that are members of this cluster, and the Database Domains where they reside
- Host name and IP addresses that will be used for SCAN for each RAC

Note that clusters should span across multiple SPARC T4-4 servers for high availability, so if you are creating a cluster, and a zone in a Database Domain is a member of that cluster on one SPARC T4-4 server, then other zones that are members of that cluster should be in similar Database Domains on different SPARC T4-4 servers.

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

- Zone 1 in Database Domain 1 in SPARC T4-4 server 1
- Zone 1 in Database Domain 1 in SPARC T4-4 server 2
- Zone 1 in Database Domain 1 in SPARC T4-4 server 3
- Zone 1 in Database Domain 1 in SPARC T4-4 server 4

Note that the worksheet has entries for ten clusters, but you could have more than ten clusters in your SPARC SuperCluster T4-4, depending on the number of Database Domains that have zones that are members of each cluster, the size of the zones in each cluster, and so on. Copy the worksheet if you have more than ten clusters in your system.

Cluster Number and Name	**Type of Cluster (Two-Node or Four-Node)**	**Zone Membership Information**	**SCAN Host Names and IP Addresses for Each RAC**
1	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
2	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
3	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
4	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
5	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
6	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
7	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:

Cluster Number and Name	Type of Cluster (Two-Node or Four-Node)	Zone Membership Information	SCAN Host Names and IP Addresses for Each RAC
8	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
9	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:
10	Type of cluster: Spans these servers:	Zone number: Database Domain number:	Host name: IP address:

Recording Your Exadata Storage Server Disk Group Layout

Use the following table to record your Exadata Storage Server disk group layout for each cluster. Note that the worksheet has entries for ten clusters, but you could have more than ten clusters in your SPARC SuperCluster T4-4, 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 will have more than ten clusters in your system.

Note that the name that you assign to each disk group should designate 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 8. For example, if the name of a particular cluster is ssc01r1, then the names of the disk groups for that cluster should be something similar to the following:

- DATA_ssc01r1
- RECO_ssc01r1
- DBFS_ssc01r1

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

Understanding Network Considerations

Understanding the Networking Setup Before Zone Creation

Before you begin the process of setting up zones on the Database Domains, you should first understand how the networks were set up when you had your SPARC SuperCluster T4-4 installed. Getting a clear understanding of the existing network setup is helpful when understanding how the network setup will change when you configure zones on the Database Domains.

When your SPARC SuperCluster T4-4 server was installed, the SPARC T4-4 servers and the domains created within each server were assigned host names and IP addresses that allowed for access to the host management, client access, and InfiniBand networks. The following figure shows a graphical representation of this initial setup.

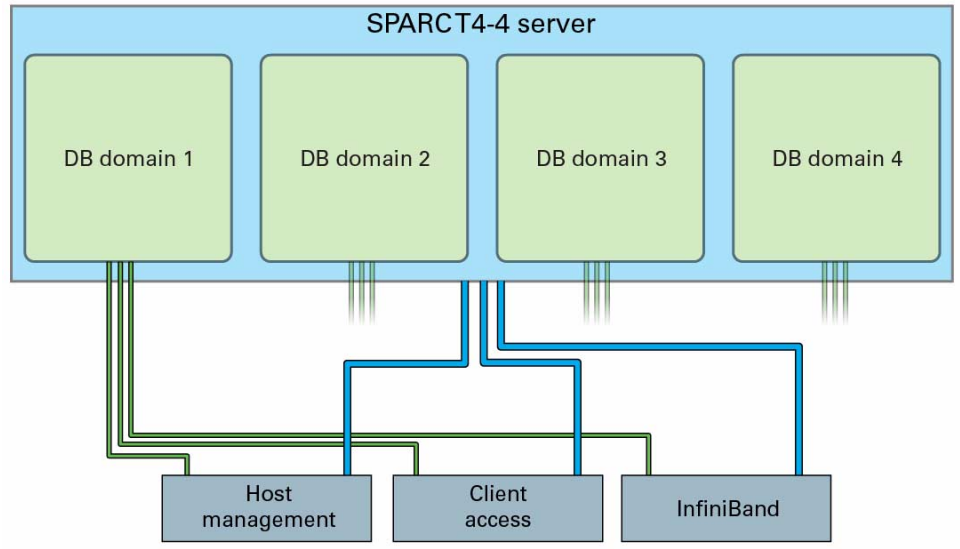

The following sections provide more detailed information on how the networks would be configured at that initial installation:

- "Host Management Network Before Zone Creation" on page 40
- "Client Access Network Before Zone Creation" on page 41
- "InfiniBand Network Before Zone Creation" on page 43

Host Management Network Before Zone Creation

There are four host management ports (NET0–NET3) on each SPARC T4-4 server. The method that you use to connect to the host management network through these four host management ports will vary, depending on the total number of domains, both Application Domains and Database Domains, that you have set up on each SPARC T4-4 server:

- One total domain – Two out of the four 1-GbE host management ports are part of one IPMP group for this domain:
  - NET0
  - NET3

- Two total domains – Two 1-GbE host management ports are part of one IPMP group in the first domain, and the other two ports are part of the second IPMP group in the second domain. Following are the domains that are created and the 1 GbE host management ports associated with each:
  - First domain: NET0–NET1
  - Second domain: NET2–NET3
- Three total domains – Two 1-GbE host management ports are part of one IPMP group in the first domain, and the other two ports are part of the second IPMP group in the second domain. For the third domain, virtual network devices (VNETs) are needed, and the two VNETs created for the third domain are part of the third IPMP group. Following are the domains that are created and the 1 GbE host management ports associated with each:
  - First domain: NET0–NET1
  - Second domain: NET2–NET3
  - Third domain: NET0-NET1, through VNETs
- Four total domains – Two 1-GbE host management ports are part of one IPMP group in the first domain, and the other two ports are part of the second IPMP group in the third domain. For the second and fourth domains, virtual network devices (VNETs) are needed, and the two VNETs created for the each domain are part of the third and fourth IPMP groups. Following are the domains that are created and the 1 GbE host management ports associated with each:
  - First domain: NET0–NET1
  - Second domain: NET2-NET3, through VNETs
  - Third domain: NET2–NET3
  - Fourth domain: NET0-NET1, through VNETs

Client Access Network Before Zone Creation

Access to the client access network is achieved through the dual-port 10-GbE network interface cards (NICs) installed in express module slots in each SPARC T4-4 server. There are four dual-port 10-GbE NICs in each SPARC T4-4 server. In each case, IPMP grouping is used to provide redundancy to the 10-GbE client access network, where two physical ports are part of one IPMP group and a single data address is used to access these two physical ports. For the situations where the two physical ports are on separate dual-port 10-GbE NICs, that data address allows traffic to continue flowing to the ports in the IPMP group, even if one of the two 10-GbE NICs fail.

The dual-port 10-GbE NICs, and the ports on each NIC, that are assigned to each domain will vary, depending on the number of domains on each SPARC T4-4 server:

- One total domain – Two out of the four dual-port 10-GbE NICs are used in this situation, and one port from each dual-port 10-GbE NIC is part of one IPMP group for this domain:
  - Express module slot 2, port 0
  - Express module slot 14, port 1
- Two total domains – Two of the four dual-port 10-GbE NICs are associated with each domain in this situation, and one port on each of the separate dual-port 10-GbE NICs would be part of one IPMP group for each domain:
  - First domain:
    - Express module slot 2, port 0
    - Express module slot 10, port 1
  - Second domain:
    - Express module slot 6, port 0
    - Express module slot 14, port 1
- Three total domains – The number of dual-port 10-GbE NICs, and the ports used by each NIC, will vary for the three domains:
  - First domain: Two dual-port 10-GbE NICs are associated with this domain, and one port on each of the two NICs would be part of one IPMP group:
    - Express module slot 2, port 0
    - Express module slot 10, port 1
  - Second domain: One dual-port 10-GbE NIC is associated with this domain, and the two ports on the one NIC is part of one IPMP group:
    - Express module slot 6, port 0
    - Express module slot 6, port 1
  - Third domain: One dual-port 10-GbE NIC is associated with this domain, and the two ports on the one NIC is part of one IPMP group:
    - Express module slot 14, port 0
    - Express module slot 14, port 1
- Four total domains – One dual-port 10-GbE NIC is associated with each domain in this situation, and the two ports on each NIC are part of one IPMP group:
  - First domain:
    - Express module slot 2, port 0
    - Express module slot 2, port 1

- Second domain:
  - Express module slot 10, port 0
  - Express module slot 10, port 1
- Third domain:
  - Express module slot 6, port 0
  - Express module slot 6, port 1
- Fourth domain:
  - Express module slot 14, port 0
  - Express module slot 14, port 1

InfiniBand Network Before Zone Creation

Connection to the InfiniBand network is achieved through the InfiniBand host channel adapters (HCAs) installed in express module slots in each SPARC T4-4 server. There are four InfiniBand HCAs in each SPARC T4-4 server.

The InfiniBand HCAs, and the ports on each HCA, that are assigned to each domain will vary, depending on the number and type of domains on each SPARC T4-4 server:

- One total domain – The HCAs and the ports on those HCAs that are used to connect to the InfiniBand network will vary, depending on the type of domain:
  - Application Domain: The InfiniBand HCA installed in express module slot 1, and ports 1 and 2 on that InfiniBand HCA, are used for connection to the InfiniBand network.
  - Database Domain: The following InfiniBand HCAs and ports are used for connection to the InfiniBand network:
    - Express module slot 1, ports 1 and 2
    - Express module slot 9, ports 1 and 2
    - Express module slot 7, ports 1 and 2
    - Express module slot 12, ports 1 and 2
- Two total domains – The HCAs and the ports on those HCAs that are used to connect to the InfiniBand network will vary, depending on the type of domain:
  - Application Domains:
    - Application Domain as the first domain:
      -- Express module slot 1, port 1
      -- Express module slot 9, port 2

- Application Domain as the second domain:
  -- Express module slot 7, port 1
  -- Express module slot 12, port 2
- Database Domains:
  - Database Domain as the first domain:
    -- Express module slot 1, ports 1 and 2
    -- Express module slot 9, ports 1 and 2
  - Database Domain as the second domain:
    -- Express module slot 7, ports 1 and 2
    -- Express module slot 12, ports 1 and 2

- Three total domains – The HCAs and the ports on those HCAs that are used to connect to the InfiniBand network will vary, depending on the type of domain:
  - Application Domains:
    - Application Domain as the first domain:
      -- Express module slot 1, port 1
      -- Express module slot 9, port 2
    - Application Domain as the second domain: Express module slot 7, ports 1 and 2
    - Application Domain as the third domain: Express module slot 12, ports 1 and 2
  - Database Domains:
    - Database Domain as the first domain:
      -- Express module slot 1, ports 1 and 2
      -- Express module slot 9, ports 1 and 2
    - Database Domain as the second domain: Express module slot 7, ports 1 and 2
    - Database Domain as the third domain: Express module slot 12, ports 1 and 2
- Four total domains – The HCAs and the ports on those HCAs that are used to connect to the InfiniBand network are the same, regardless of the type of domain:
  - First domain: Express module slot 2, ports 1 and 2
  - Second domain: Express module slot 10, ports 1 and 2
  - Third domain: Express module slot 6, ports 1 and 2
  - Fourth domain: Express module slot 14, ports 1 and 2

Understanding the Networking Setup After Zone Creation

After you create zones on the Database Domains, the existing networking setup that was configured before the zone creation will remain in place with no modifications (see "Understanding the Networking Setup Before Zone Creation" on page 39 for more information on the networking setup at the initial installation of your SPARC SuperCluster T4-4). When you create zones on the Database Domains, additional virtual networking interfaces are created to accommodate those zones, allowing the zones access to the host management, client access and InfiniBand networks. The following figure shows a graphical representation of how the zones would connect to these networks.

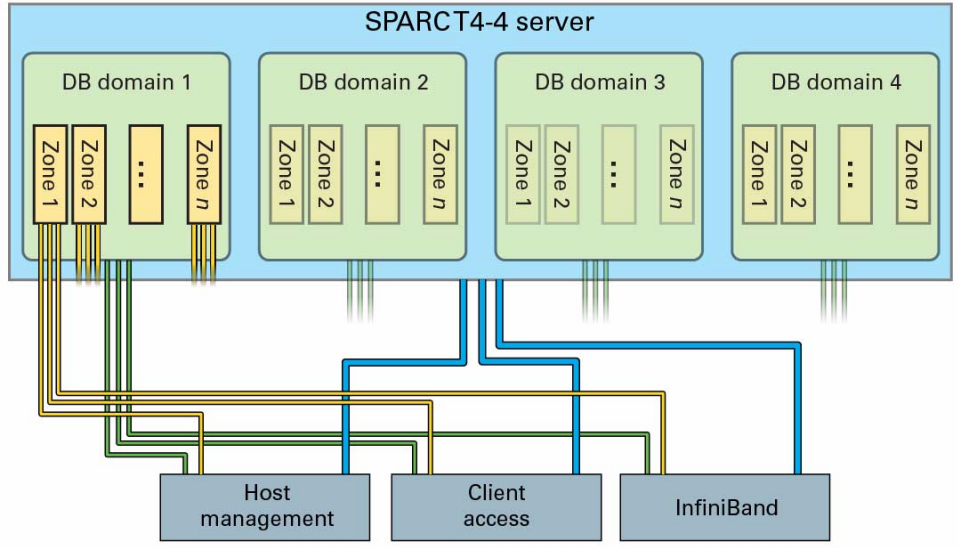

These topics describe how the networks will be set up for the zones that you will be creating:

- "Host Management Network After Zone Creation" on page 46
- "Client Access Network After Zone Creation" on page 47
- "InfiniBand Network After Zone Creation" on page 48

Host Management Network After Zone Creation

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

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

Whether or not VNICs are created automatically for you as part of the zone creation process, or if you have to manually create VNETs to use as access to the host management network for the zones you want to create on the Database Domains, depends on the total number of domains that you have on each SPARC T4-4 server and the zones that you want to create on certain domains.

- "Zones on a One-Domain System" on page 46
- "Zones on a Two-Domain System" on page 46
- "Zones on a Three-Domain System" on page 47
- "Zones on a Four-Domain System" on page 47

Zones on a One-Domain System

The domain in a one-domain system uses physical interfaces to connect to the host management network (see "Host Management Network Before Zone Creation" on page 40 for more information). Therefore, when zones are created on domains in a one-domain system, VNICs are automatically created as part of the zone creation process. You do not have to manually create VNETs for any zones in this situation.

Zones on a Two-Domain System

Both domains in a two-domain system use physical interfaces to connect to the host management network (see "Host Management Network Before Zone Creation" on page 40 for more information). Therefore, when zones are created on domains in a two-domain system, VNICs are automatically created as part of the zone creation process. You do not have to manually create VNETs for any zones in this situation.

Zones on a Three-Domain System

The first two domains on a three-domain system use physical interfaces to connect to the host management network (see "Host Management Network Before Zone Creation" on page 40 for more information). Therefore, when zones are created on the first two domains in a three-domain system, VNICs are automatically created as part of the zone creation process. You do not have to manually create VNETs for any zones in this situation.

However, the third domain in a three-domain system uses VNETs to connect to the host management network. Therefore, if you are creating zones on the third domain in a three-domain system, you must manually create VNETs on the third domain to use as a connection to the host management network for the zones on that domain. See "Creating VNETs" on page 135 for those instructions.

Zones on a Four-Domain System

The first and third domains on a four-domain system use physical interfaces to connect to the host management network (see "Host Management Network Before Zone Creation" on page 40 for more information). Therefore, when zones are created on the first or third domains in a four-domain system, VNICs are automatically created as part of the zone creation process. You do not have to manually create VNETs for any zones in this situation.

However, the second and fourth domains in a four-domain system use VNETs to connect to the host management network. Therefore, if you are creating zones on the second or fourth domain in a four-domain system, you must manually create VNETs on those domains to use as a connection to the host management network for the zones on those domains. See "Creating VNETs" on page 135 for those instructions.

Client Access Network After Zone Creation

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

InfiniBand Network After Zone Creation

When you configure zones on the Database Domains, part of the zone creation process will also automatically create the VNICs that will be used as access to the InfiniBand network for the zones. You should not have to manually configure the VNICs that will be used for the InfiniBand network for the zones.

Understanding the Network Setup for Clusters

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

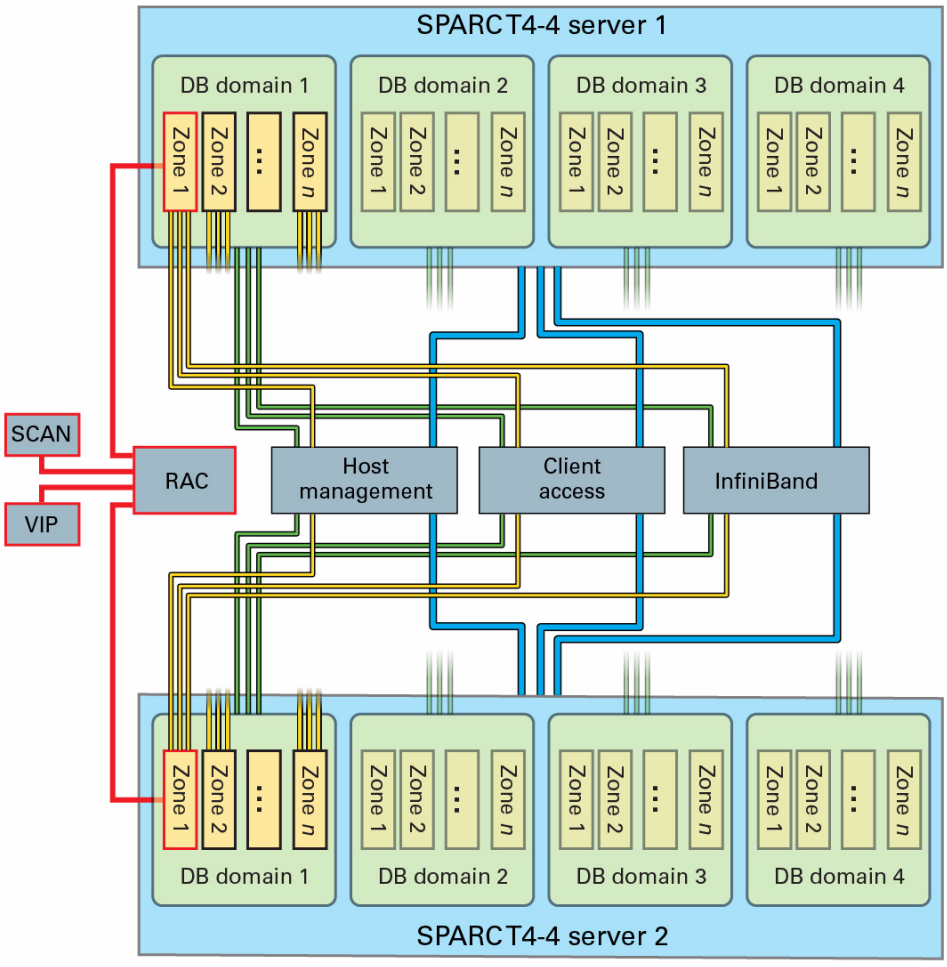

In this respect, access to the clusters that have zones as members is no different than how clusters are accessed when clusterS are used on two or more physical computers (each with an Oracle RDBMS instance) to concurrently access a single database.

Preparing to Configure Zones on Database Domains

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

▼ Prepare to Configure Zones

Complete these tasks before setting up zones and clusters on the Database Domains with zones in the SPARC SuperCluster T4-4.

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 7 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:
  - Administration network
  - 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 SPARC SuperCluster T4-4
   - IP addresses and host names for Oracle ILOM for the following components:
     - SPARC T4-4 servers with Database Domains that you are setting up zones on
     - Exadata Storage Servers
   - IP addresses and host names for the 1 GbE 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 SPARC SuperCluster T4-4 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 53.

Existing Network Worksheet

	Existing Host Names and IP Addresses	
Component	**Host Name**	**IP Address**
Oracle ILOM for SPARC T4-4 Server 1		
Oracle ILOM for SPARC T4-4 Server 2		
Oracle ILOM for SPARC T4-4 Server 3		
Oracle ILOM for SPARC T4-4 Server 4		
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		
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 IB spine switch		

	Existing Host Names and IP Addresses	
Component	**Host Name**	**IP Address**
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		

▼ Update the Base Software

1. **Log in to the first SPARC T4-4 server.**
2. **Determine if your system has the correct level of base software installed.**

```
# pkg list entire
```

- If you see output similar to the following, then you have the correct level of software installed. Go to "Determine the Repository Location" on page 55:

```
NAME (PUBLISHER) VERSION                    IFO
entire 0.5.11-0.175.1.19.6.11.1 i--
```

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

- Solaris 11 Update 1 OS
- SRU 19.6
- Build 11.1 or later

- If the SRU version is less than SRU 19.6, then you do not have the correct level of software installed. Follow these instructions to get the latest version of the software:

a. Log in to MOS.

http://support.oracle.com

b. Select the Patches & Updates tab.

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

d. In the Product field, search for and select SPARC SuperCluster T4-4.

Select any of the SPARC SuperCluster T4-4 options.

e. In the Release field, select the SPARC SuperCluster T4-4 1.1 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 SPARC SuperCluster T4-4 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 SPARC SuperCluster T4-4

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

1. **Log in to a Database Domain.**
2. **Type:**

```
# pkg publisher
```

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

- If you see one of the the following outputs in the URI field, then your system is set up to access files remotely, from Oracle's IPS repository. Go to "Install or Update Packages From the Remote Repository" on page 56.

```
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 SPARC SuperCluster T4-4. Go to "Install or Update Files From the Local Repository" on page 57.

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

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

▼ Install or Update Packages From the Remote Repository

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

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

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

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

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

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

a. **Go to the following location:**

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

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

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

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

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

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

   b. **Repeat** Step a **for every Database Domain to contain zones.**

      When you have updated the packages on every Database Domain to contain zones, go to "Verify Configuration Tool Installation" on page 62.

4. **Install the necessary packages:**

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

5. **Repeat** Step 2 **through** Step 4 **for every Database Domain to contain zones.**

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

   Go to "Verify Configuration Tool Installation" on page 62.

▼ 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 1 GbE host management network.**

2. **Locate the project for the local repository on the ZFS storage appliance:**

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

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

3. **Select that project for the repository:**

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

   For example:

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

4. **List the file systems in the repository area:**

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

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

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

   - `solaris`
   - `s11repo`
   - `exa-family`

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

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

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

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

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

     a. **Log in to MOS:**

        `http://support.oracle.com`

     b. **Select the `Patches & Updates` tab.**

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

d. In the Product field, search for and select SPARC SuperCluster T4-4.

Select any of the SPARC SuperCluster T4-4 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 SPARC SuperCluster T4-4.

Select any of the SPARC SuperCluster T4-4 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 (1.1 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 SPARC T4-4 server.**

      For example, copy the ZIP file to the `/var/tmp` directory on the first SPARC T4-4 server.

   b. **Unzip the ISO image ZIP file.**

   c. **Locate the ISO image in this directory.**

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

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

      where:

      - *tmp-dir* is the name of the temporary directory that you use as a mount point.

      - *iso-name*`.iso` is the name of the ISO image.

      - *path-to-repo-on-ZFSSA* is the path to the repository on the ZFS storage appliance. This path appears after the `file://` output from the `pkg publisher` command that you entered in the section "Determine the Repository Location" on page 55.

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

10. **Install the packages onto each Database Domain with zones where you will be setting up zones.**

    a. **Log in to the first Database Domain with zones.**

    b. **Install the necessary packages:**

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

11. **Repeat these steps for every Database Domain.**

12. **When you have installed the packages on every Database Domain that will contain zones, go to "Verify Configuration Tool Installation" on page 62.**

▼ Verify Configuration Tool Installation

1. **Log in to the first Database Domain to contain zones.**
2. **Download the latest Oracle OneCommand utility.**
   a. **Locate the My Oracle Support note that provides information on the latest version of the patch with the Oracle OneCommand utility.**

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

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

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

Caution – You must download patch 19766190 or later.

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

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

      The options are:

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

   d. **Click the Download button to download the patch.**
   e. **Unzip the patch file.**

      Several files become available.

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

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

4. **Verify that the necessary tools and scripts are available in the `bin` and `zonetools` subdirectories.**
   - To verify the tools and scripts are in the `bin` subdirectory, type:

```
# ls bin
```

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

```
# ls zonetools
```

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

5. **Repeat Step 1 through Step 4 for all other Database Domains to contain zones.**
6. **Use the OEDA to create the configuration files to be used by the Java OneCommand utility later on in this process.**

   Go to "Creating Configuration Files Using OEDA" on page 65.

Creating Configuration Files Using OEDA

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

▼ Verify Exadata Storage Server Disk Space for Additional Zones

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

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

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

1. **Determine the amount of free space available on the Exadata Storage Servers:**

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

You should see output similar to the following:

```
CD_00_etc25celadm01 366.6875G ((offset=162.046875G,size=366.6875G))
CD_01_etc25celadm01 366.6875G ((offset=162.046875G,size=366.6875G))
CD_02_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_03_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_04_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_05_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_06_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_07_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_08_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_09_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_10_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
CD_11_etc25celadm01 393.8125G ((offset=164.046875G,size=393.8125G))
FD_00_etc25celadm01 0
FD_01_etc25celadm01 0
FD_02_etc25celadm01 0
FD_03_etc25celadm01 0
FD_04_etc25celadm01 0
FD_05_etc25celadm01 0
FD_06_etc25celadm01 0
FD_07_etc25celadm01 0
FD_08_etc25celadm01 0
FD_09_etc25celadm01 0
FD_10_etc25celadm01 0
FD_11_etc25celadm01 0
FD_12_etc25celadm01 0
FD_13_etc25celadm01 0
FD_14_etc25celadm01 0
FD_15_etc25celadm01 0
```

2. **Look at the CD_*number* entries to determine the amount of available space on each of the Exadata Storage Servers.**

3. **Use the information on the amount of available space on each of the Exadata Storage Servers to determine how much space you can use for the disk groups for each new cluster you are creating for the new zones.**

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

   - DATA
   - RECO
   - DBFS (if necessary)

4. **Create the configuration files using OEDA.**

   See "Creating Configuration Files Using OEDA: Manual Procedures" on page 67.

Creating Configuration Files Using OEDA: Manual Procedures

The first time you create configuration files using the latest OEDA, you must manually input all the necessary information into the forms. Once you have gone through this manual process one time, you can then use the generated configuration files and the automated procedures described in "Creating Configuration Files Using OEDA: Automated Procedures" on page 105.

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

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

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

Follow these procedures to use OEDA to create your configuration files for a system with a software build prior to 1.1.0.build-03:

Oracle Exadata Deployment Assistant Overview

The Oracle Exadata Deployment Assistant is normally used when the system is initially installed and configured, and its purpose is to gather networking information for the Database Domains and for the physical components in the SPARC SuperCluster T4-4, such as the SPARC T4-4 servers, Exadata Storage Servers, switches and PDUs. However, in this section, you will be using the Oracle Exadata Deployment Assistant to gather networking information for the clusters and single-instance database zones that you want to set up.

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

- The Database Machine Prefix would normally refer to the name of the SPARC SuperCluster T4-4 as a whole, and would be used as a prefix for the hostnames for the components inside the SPARC SuperCluster T4-4. However, for the procedures in this section, the Database Machine Prefix refers to the name for the cluster that you will be setting up, and will be used as a prefix for the zones that are a part of that cluster.
- The Oracle Database Deployment Assistant would normally assign hostnames to the compute nodes (the SPARC T4-4 servers) in several areas, such as the Administration Network. However, in some areas in this section, information that is being entered for compute nodes will actually apply to the zones that are part of

this particular cluster, not to the compute nodes themselves. The instructions in this section will point out when information that needs to be entered for compute nodes actually applies to zones rather than compute nodes.

- In other areas in this process, information that is being entered for a compute node actually does apply to the compute node and not to a zone in the cluster, such as the ILOM host name and IP address for the compute node. The instructions in this section will point out when information that needs to be entered for compute nodes actually does apply for the compute nodes, and not for zones in this cluster.
- When the Oracle Database Deployment Assistant is first run on your SPARC SuperCluster T4-4 during the initial installation and configuration, part of the process would include assigning host names and IP addresses to the switches and PDUs. Because you are using this same tool in this situation to gather networking information to be used for the clusters and single-instance database zones, certain screens will have networking information for the switches and PDUs that will be based on the information you are entering for this particular cluster, and would therefore be different from the information that was originally entered for those components when your system was first set up. The networking information for the switches and PDUs are not cluster-dependant, so you will be asked to manually change the networking information for those components back to their original values several times in these procedures.

▼ Set Up the OEDA

1. **Prepare to configure zones.**

   See "Preparing to Configure Zones on Database Domains" on page 51.

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

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

   The OEDA files are part of the `Ocmd` ZIP file that you downloaded earlier in this process. See "Verify Configuration Tool Installation" on page 62 for more information.

   There are different OEDA ZIP files for these operating systems:

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

4. **Unzip the `Ocmd` ZIP file for your operating system on the first Database Domain that will deploy zones.**

5. **Go into the directory that was created after you unzipped the `Ocmd` ZIP file for your operating system.**

6. **Start the OEDA.**

   Go to "Start the OEDA" on page 70.

▼ 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.

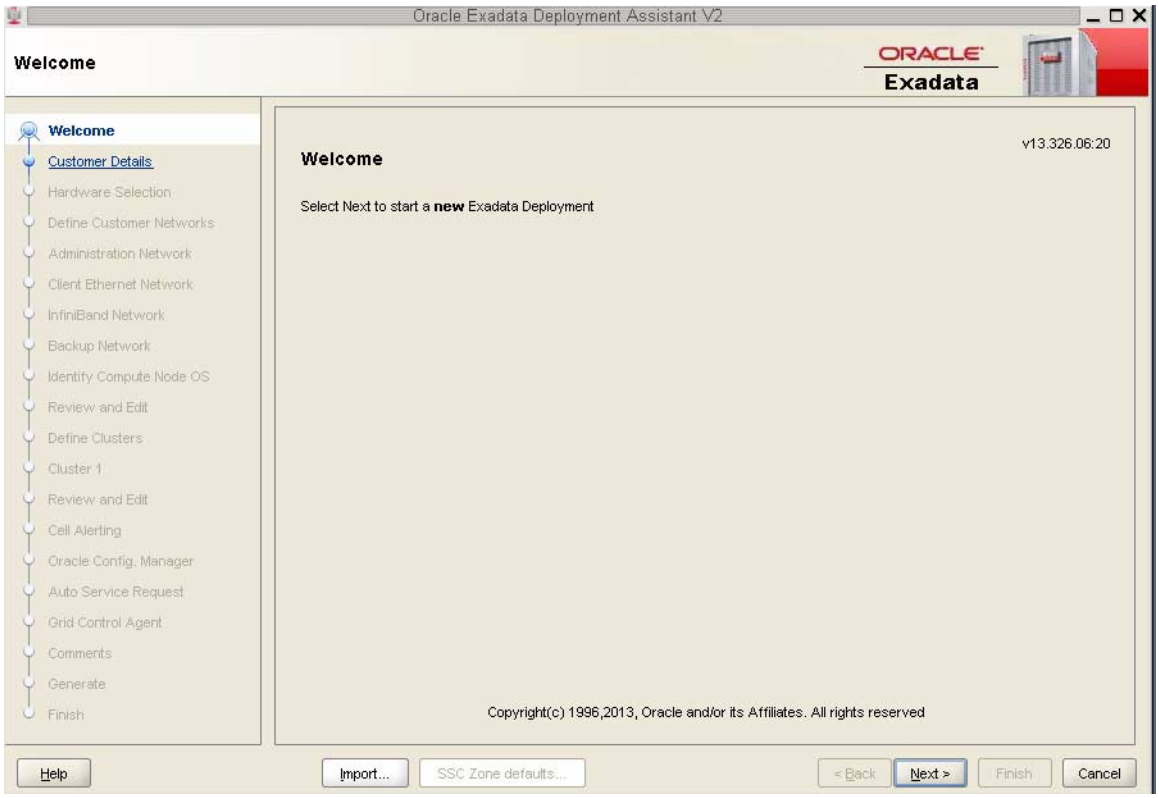

2. **In the Welcome window, click Next to begin filling out the forms in the Oracle Exadata Deployment Assistant.**

   The Customer Details page is displayed. Go to "Complete the Customer Details Page" on page 71.

▼ Complete the Customer Details Page

1. **Complete the fields in the Customer Details page.**

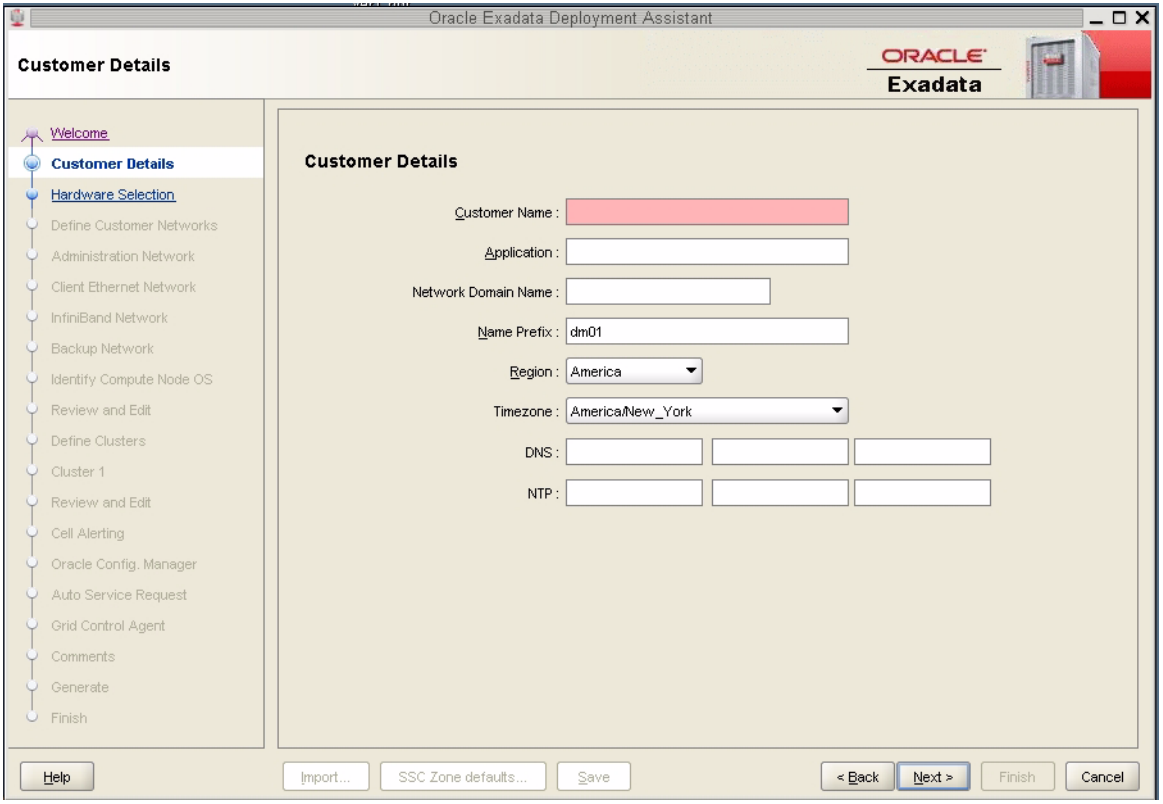

The value that you enter in the Name Prefix will be the prefix that will be used for the prefix for the host names for the zones or clusters that you will be creating. This value should be eight characters or fewer. This prefix will be used for this set of zones that you are entering information for.

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

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

---

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

   The Hardware Selection page is displayed. Go to "Complete the Hardware Selection Page" on page 73.

▼ Complete the Hardware Selection Page

1. **Review the Hardware Selection page.**

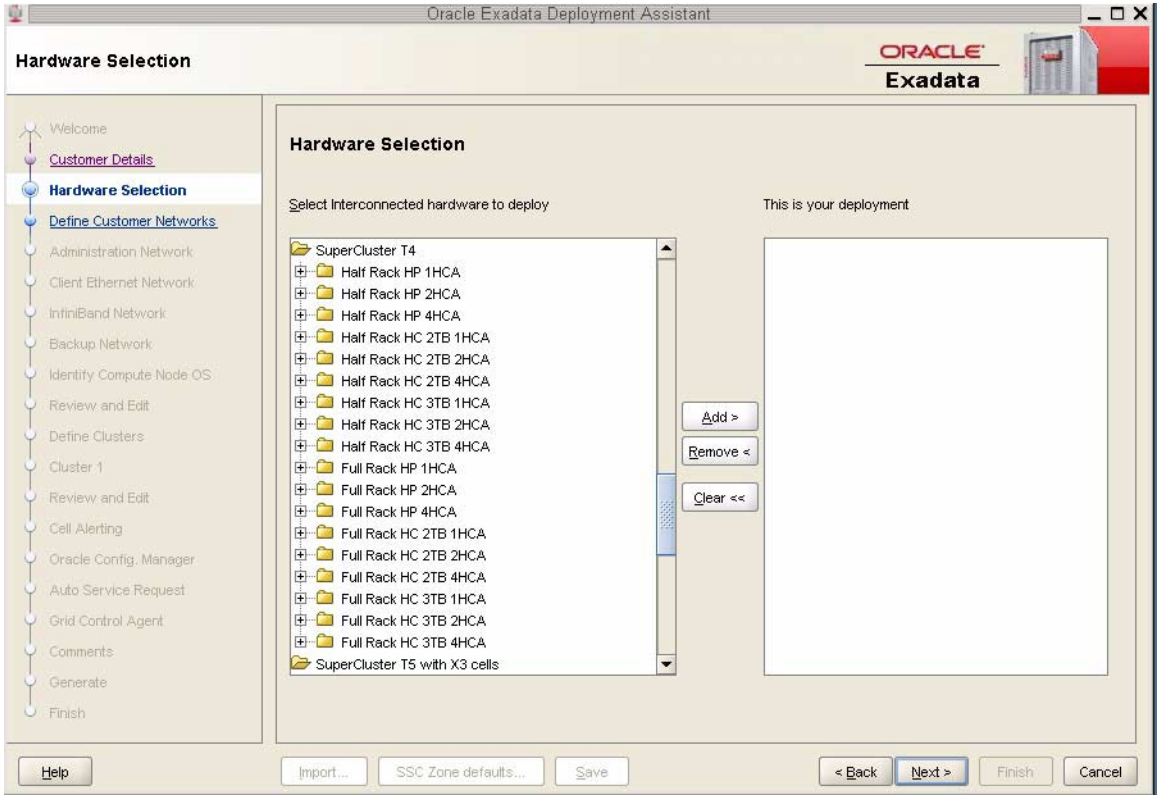

2. **Open the proper folder in the Hardware Selection page for your configuration.**

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

   - The type of SPARC SuperCluster T4-4 that you have (Half Rack or Full Rack).
   - How many total domains are on your SPARC T4-4 servers, and which domain you are setting zones up on.
   - The number of IB HCAs that are associated with that Database Domain. See "Understanding Domain Configurations" on page 3 for more information.
   - Whether you have high performance disks (HP) or high capacity (HC) disk in the Exadata Storage Servers.

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

   - A Half Rack version of the SPARC SuperCluster T4-4

- You have three total domains, and you are setting up zones on the first domain
- There are two IB HCAs associated with this domain
- You have high performance (HP) disks in the Exadata Storage Servers

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

`Half Rack HP 2HCA`

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

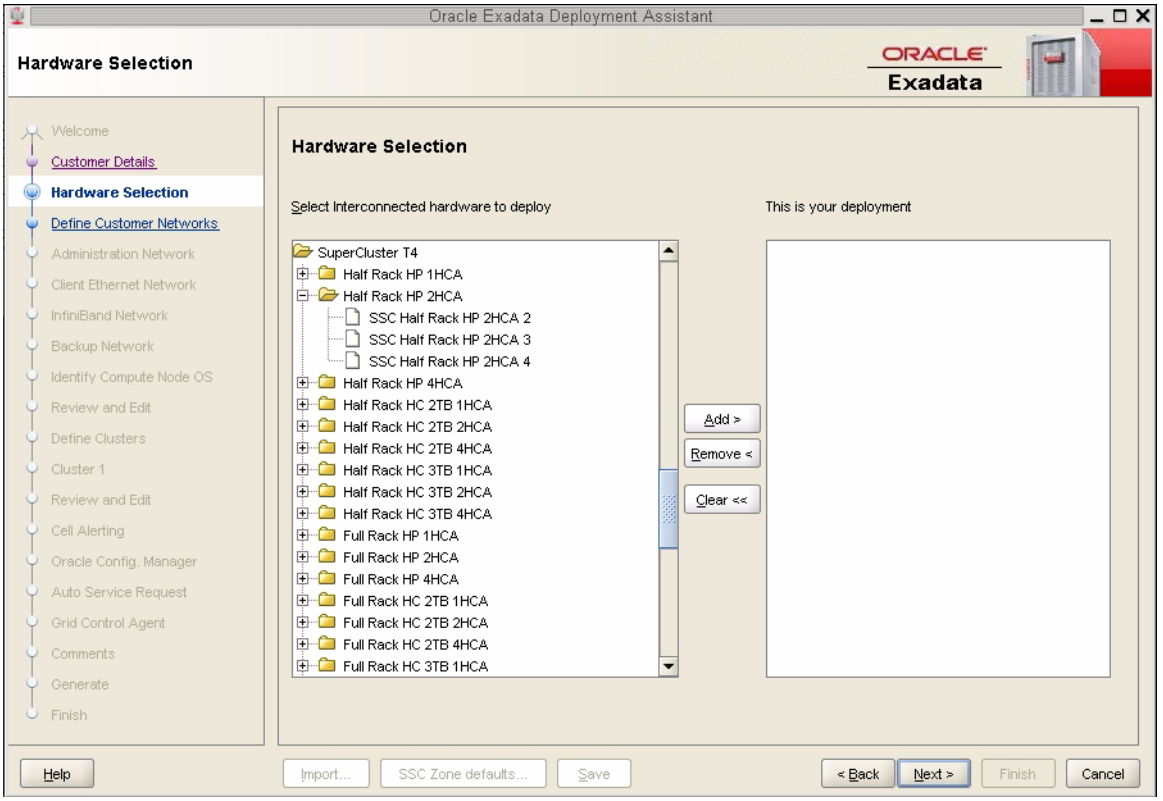

3. **Make the appropriate selection within the folder that you opened.**

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

   For example, within the `Half Rack HP 2HCA` folder, you would see the following options:

   - `SSC Half Rack HP 2HCA 2`
   - `SSC Half Rack HP 2HCA 3`

- `SSC Half Rack HP 2HCA 4`

The number at the end of each option denotes the number of domains with 2 IB HCAs that you could be part of this cluster. Therefore:

- If you are creating a cluster that spans across *two* domains with 2 IB HCAs, you would select `SSC Half Rack HP 2HCA 2`
- If you are creating a cluster that spans across *three* domains with 2 IB HCAs, you would select `SSC Half Rack HP 2HCA 3`
- If you are creating a cluster that spans across *four* domains with 2 IB HCAs, you would select `SSC Half Rack HP 2HCA 4`

When you have selected the proper configuration for your system, click the Add button to add that selection to the `This is your deployment` window, then click Next. The Networking page is displayed.

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

   The Define Customer Networks page is displayed. Go to "Complete the Define Customer Networks Page" on page 75.

▼ Complete the Define Customer Networks Page

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

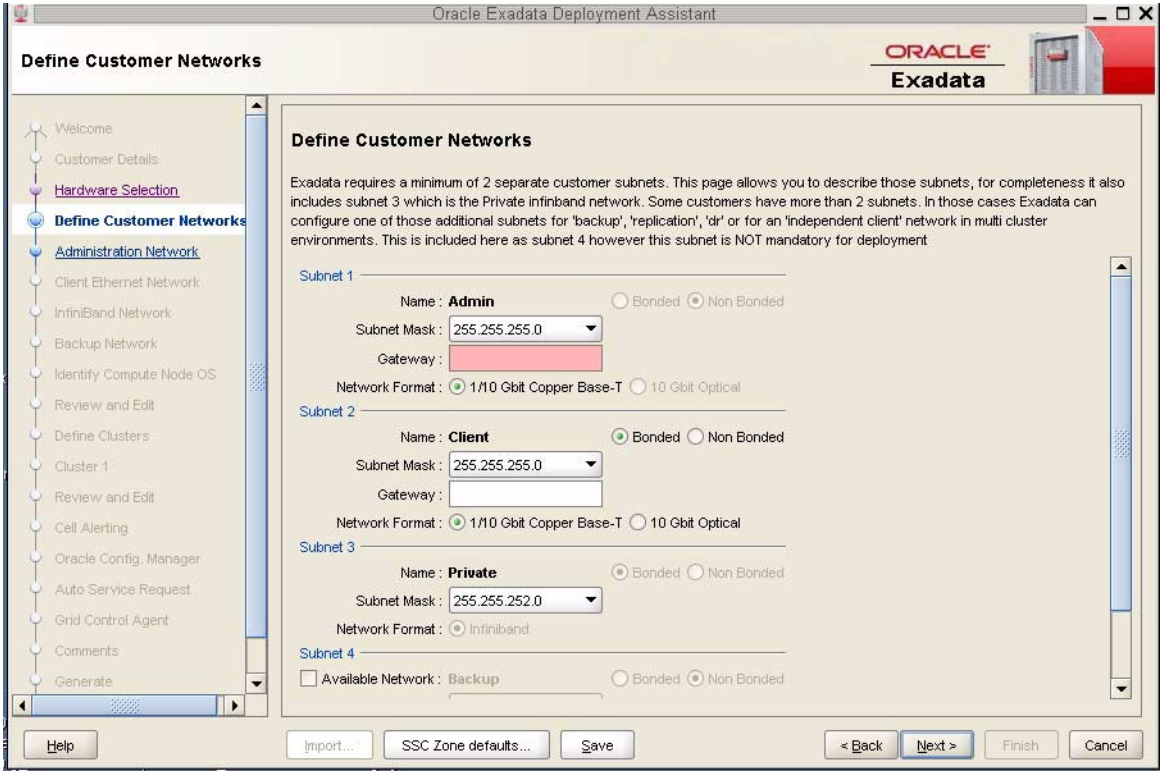

a. In the Subnet 1 (Admin) area, provide the following information:

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

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

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

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

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

2. Click the Next button.

The Administration Network page is displayed. Go to "Complete the Administration Network Page" on page 77.

▼ Complete the Administration Network Page

1. Complete the fields in the Administration Network page.

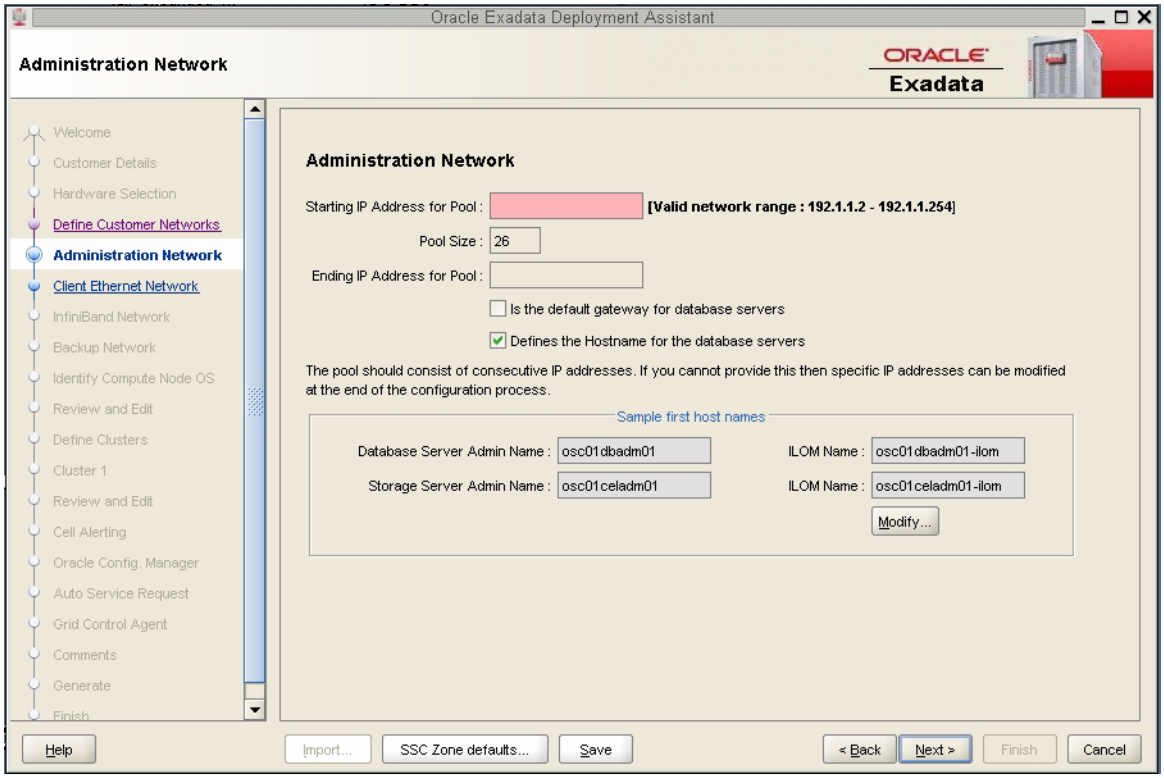

Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.

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

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

   The Admin Network Format Masks page is displayed.

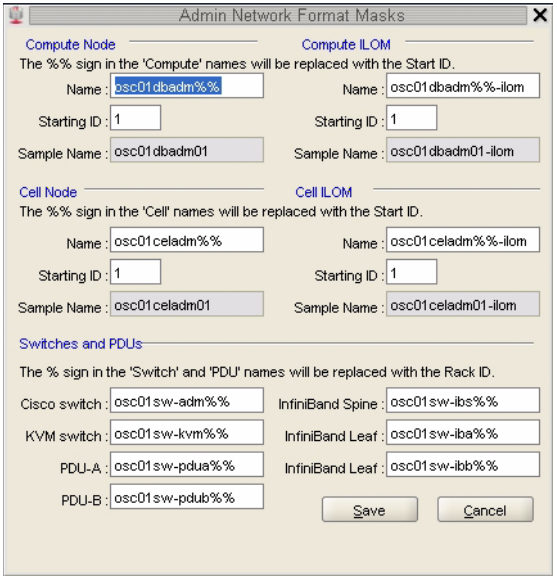

3. **Modify the host names in this page for the zones in this cluster.**

a. In the Compute Node area, enter a unique name for this set of zones in the Name field in the Compute Node area.

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

For example, if you entered ssc01r1 for the Name Prefix in the Customer Details page, you should see the following text in the Name field in the Compute Node area:

```
ssc01r1dbadmin%%
```

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

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

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

cluster-name`z#db#cn%`

where:

- *cluster-name* is the Name Prefix that you entered in the Customer Details page.
- `z#` is used for the zone numbers for this cluster. The zones numbers would be the same as the number for the cluster that they reside in. For example, if this is the first cluster that you are configuring (r1), then all the zones in this cluster would be z1 zones.
- `db#` is used for the Database Domain numbers on each SPARC T4-4 server that will house the zones that will be members of this cluster.
- `cn%` is used for the SPARC T4-4 servers (compute nodes) that house the Database Domains that will contain the zones that you are creating. Note that the % part of the name allows for a number to be automatically generated at the end of this process.

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

- This is the first cluster and set of zones that you will be setting up on your system
- The zones for this cluster will all reside on the first domain on the SPARC T4-4 servers

- This will be a two-node cluster (two zones in this cluster), and the zones will reside on the Database Domains in SPARC T4-4 servers 1 and 2 in the SPARC SuperCluster T4-4 Half Rack

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

```
ssc01r1z1db1cn%
```

where:

- `ssc01r1` is the Database Machine Prefix that you chose earlier
- `z1` indicates that these zones are for the first cluster that you will be configuring
- `db1` indicates that this set of zones that will be members of this cluster all reside on the first domain on each SPARC T4-4 server
- The `cn` before the % allows for the compute node numbers to be automatically generated at the end of the process

---

Note – There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T4-4 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T4-4 server, and the second zone resides on the second Database Domain on the second SPARC T4-4 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

---

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

The Name field in the Compute ILOM area will automatically populate, depending on the information that you entered in the Name field in the Compute Node area. For example, if you entered `ssc01db1cn%-z1` in the Name field in the Compute Node area, then the following information will automatically populate in the Name field in the Compute ILOM area:

```
ssc01db1cn%-z1-ilom
```

However, the Oracle ILOM information applies to the SPARC T4-4 server as a whole, not to a cluster with zones that reside on that server. You must change this field so that the format matches the host name for Oracle ILOM that has already been assigned to your SPARC T4-4 servers, where the %% in the field denotes a numerical wild card that will be automatically assigned.

Note – For some SPARC SuperCluster T4-4s, the Oracle ILOM host names assigned to the SPARC T4-4 servers use a single-digit number, such as ssc01cn**1**-ilom, instead of a double-digit number, such as ssc01cn**01**-ilom. Each percentage sign (%) wild card generates a single digit, so if the Oracle ILOM host names assigned to your SPARC T4-4 servers use a single-digit number, remove one of the two percentage signs so that host names that are automatically generated use only a single digit.

For example, if the host names for Oracle ILOM for your SPARC T4-4 servers are *SuperCluster-name*cn*T4-4-number*-ilom (for example, ssc01cn1-ilom), then you would delete the necessary text in this field until you have this:

```
ssc01cn%-ilom
```

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

The Name field in the Cell Node area will automatically populate, depending on the Database Machine Prefix that you chose earlier. For example, if you used ssc01r1 as the Database Machine Prefix, then the following information will automatically populate in the Name field in the Cell Node area:

```
ssc01r1celadmin%%
```

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

Note – For some SPARC SuperCluster T4-4s, the 1 GbE host management host names assigned to the Exadata Storage Servers use a single-digit number, such as ssc01cel**1**, instead of a double-digit number, such as ssc01cel**01**. Each percentage sign (%) wild card generates a single digit, so if the 1 GbE host management host names assigned to your Exadata Storage Servers use a single-digit number, remove one of the two percentage signs so that host names that are automatically generated use only a single digit.

For example, if the host names for the 1 GbE host management for your Exadata Storage Servers are *SuperCluster-name*cel*ESS-number* (for example, ssc01cel1), then you would delete the necessary text in this field until you have this:

```
ssc01cel%
```

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

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

d. Change all information for the switches and PDUs.

In the Switches and PDUs area, manually change all of the information in the beginning of each field, before the `sw`, so that the information matches the original name that was assigned to the SPARC SuperCluster T4-4 at the initial deployment. Doing so will create generated names that will match the host names for the 1 GbE host management that were originally assigned to each switch and PDU.

For example, if the original name of the SPARC SuperCluster T4-4 was ssc01, then you would enter the following information for the switches and PDUs:

- Cisco switch: `ssc01sw-admin%`
- KVM switch: `ssc01sw-kvm%`
- PDU-A: `ssc01sw-pdua%`
- PDU-B: `ssc01sw-pdub%`
- IB switch: `ssc01sw-ibs%`
- IB Leaf: `ssc01sw-iba%`
- IB Leaf: `ssc01sw-ibb%`

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

Component/Cluster	Correct Entry
Compute Node Name and Starting ID	Should have a value that applies to the zones in this particular cluster
Compute ILOM Name and Starting ID	Should have a value that applies to the SPARC T4-4 servers
Cell Node Name and Starting ID	Should have a value that applies to the Exadata Storage Servers
Cell ILOM Name and Starting ID	Should have a value that applies to the Exadata Storage Servers
Switches and PDUs Names	Should have values that apply to the switches and PDUs

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

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

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

The Client Ethernet Network page is displayed. Go to "Complete the Client Ethernet Network Page" on page 83.

▼ Complete the Client Ethernet Network Page

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

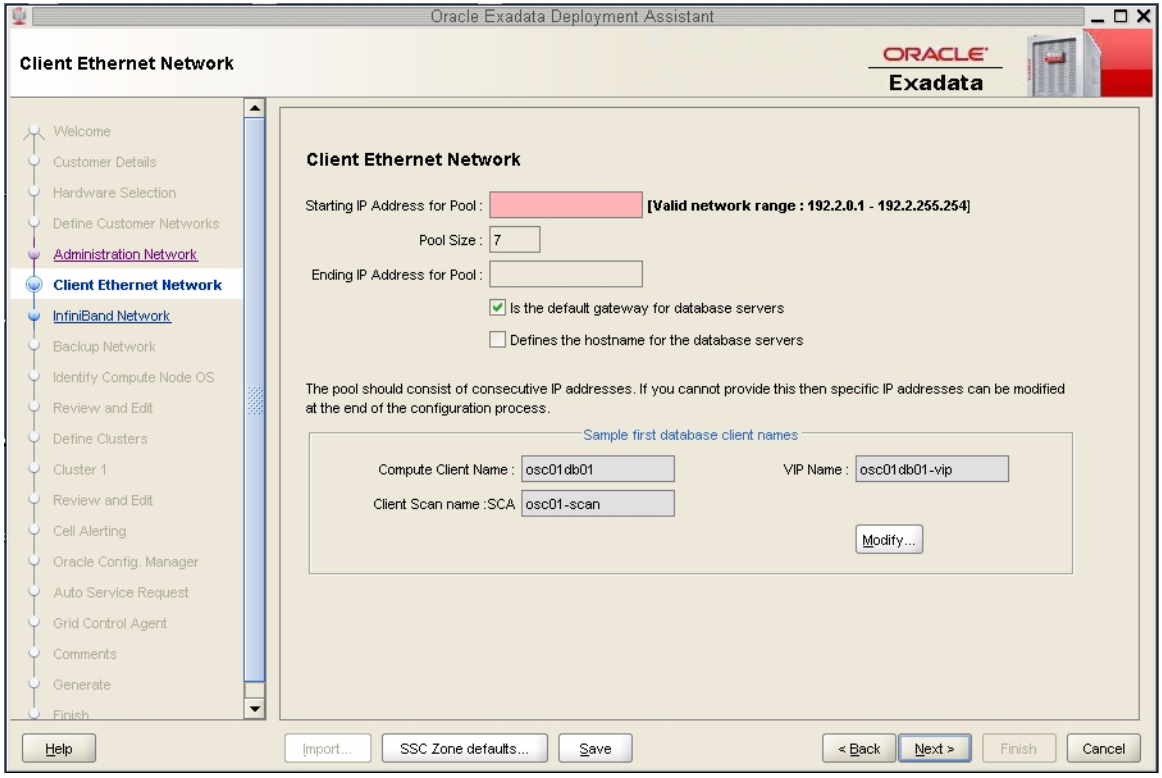

Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.

Field	Entry
Is the default gateway for database servers	Select this option if you want the client ethernet network to be the default gateway for the cluster.
Defines the Hostname for the database servers	Select this option.

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

   The Client Network Format Masks page is displayed.

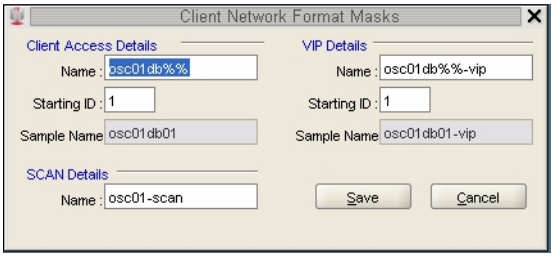

3. **Modify the host names in this page for the zones in this cluster.**

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

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

      For example, if you entered `ssc01r1` for the Database Machine prefix in the Customer Details page, you should see this entry in the Name field in the Client Access Details area:

      ```
      ssc01r1client%%
      ```

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

      cluster-name`z#db#cn%-z#-client`

      For example:

      ```
      ssc01r1z1db1cn%-client
      ```

Note – There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T4-4 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T4-4 server, and the second zone resides on the second Database Domain on the second SPARC T4-4 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

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

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

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

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

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

The IB Network page is displayed. Go to "Complete the IB Network Page" on page 85.

▼ Complete the IB Network Page

1. Complete the fields in the IB Network page.

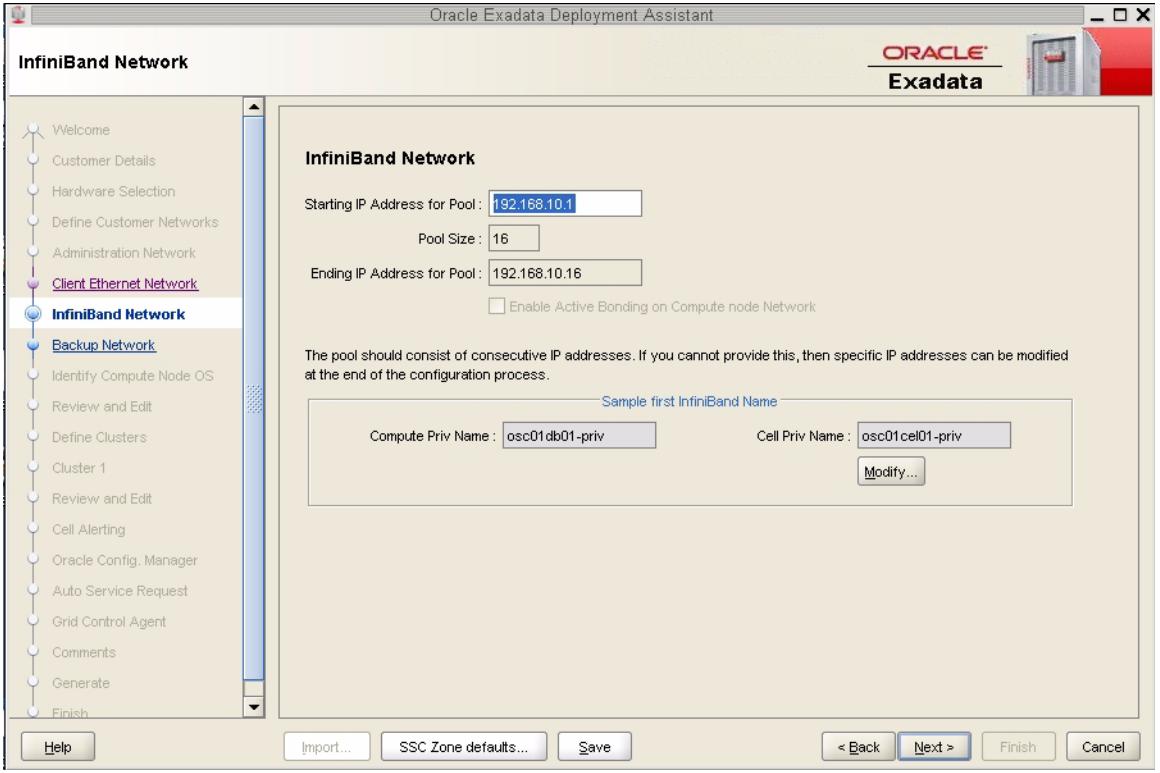

Field	Entry
Starting IP Address for Pool	Type the starting IP address for this cluster.

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

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

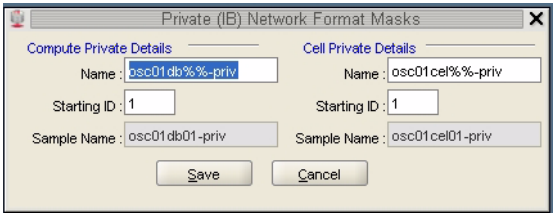

3. **Modify the host names in this page for the zones in this cluster.**

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

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

For example, if you entered `ssc01r1` for the Database Machine prefix in the Customer Details page, you should see this entry in the Name field in the Compute Private Details area:

```
ssc01r1db%%-priv
```

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

cluster-name`z#db#cn%-priv`

For example:

```
ssc01r1z1db1cn%-priv
```

Note – There is a valid configuration where you would have a cluster where the zones reside on different Database Domains on the SPARC T4-4 servers (for example, a two-node cluster where the first zone resides on the first Database Domain on the first SPARC T4-4 server, and the second zone resides on the second Database Domain on the second SPARC T4-4 server). In this case, enter the Database Domain number for the first zone in this field for now - you will make manual modifications for the zone on the second Database Domain later on in this process.

---

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

The Name field in the Cell Private Details area will automatically populate, depending on the Database Machine Prefix that you chose earlier. For example, if you used ssc01r1 as the Database Machine Prefix, then the following information will automatically populate in the Name field in the Cell Private Details area:

```
ssc01r1cel%%-priv
```

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

For example, if the host names for the IB network for your Exadata Storage Servers are *supercluster-name*cel*ESS-number*-priv (for example, ssc01cel01-priv), then you would delete the necessary text in this field until you have this:

```
ssc01cel%%-priv
```

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

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

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

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

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

The Backup Network page is displayed.

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

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

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

Go to "Complete the Identify Compute Node OS and Enable Capacity-on-Demand Configuration Page" on page 89.

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

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

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

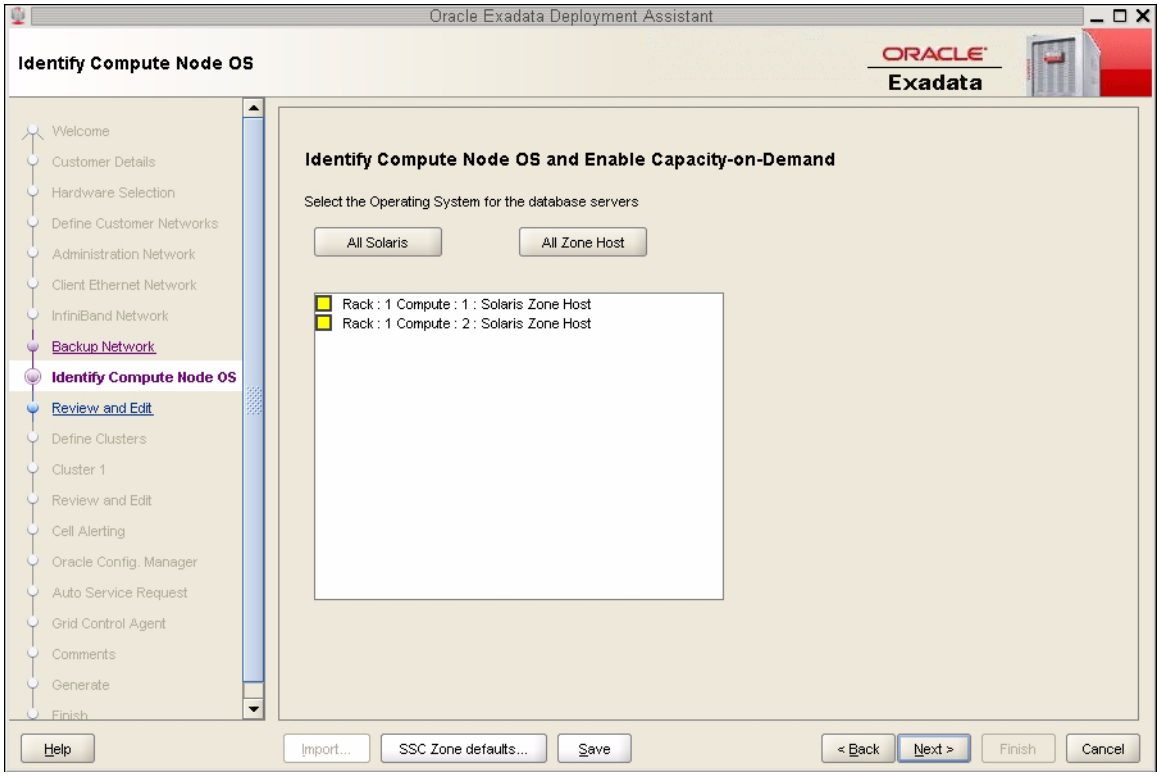

2. **Click Next.**

   The Management and Private Networks page is displayed. Go to "Review the Information in the Management and Private Networks Page" on page 89.

▼ Review the Information in the Management and Private Networks Page

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

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

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

- Database Domains with zones (shown as `Compute Node # - Zone Host`)
- Exadata Storage Servers in the SPARC SuperCluster T4-4
- Switches and PDUs in the SPARC SuperCluster T4-4

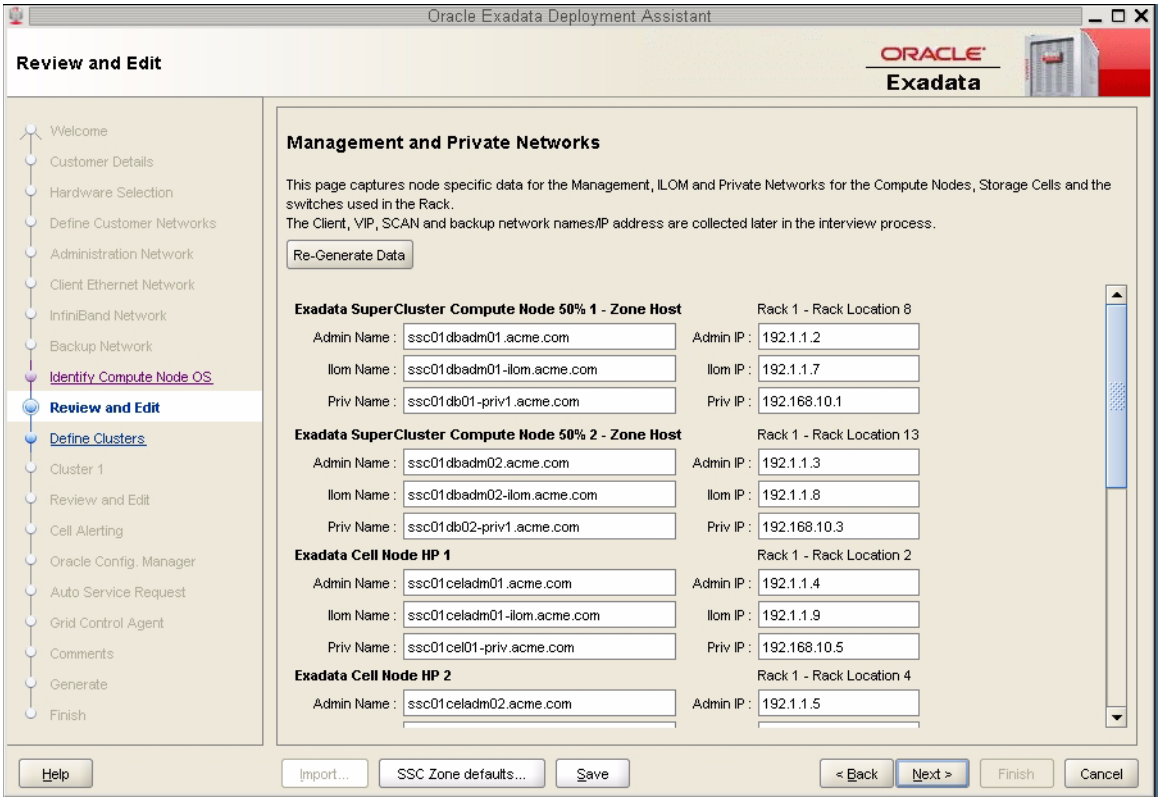

2. **Click Next after reviewing the information in this page.**

   The Define Clusters page is displayed.

3. **Complete the Define Clusters page.**

   Go to "Complete the Define Clusters Page" on page 91.

▼ Complete the Define Clusters Page

1. **Determine the number of clusters that you will create and select the appropriate number in the Number of Clusters to create area.**

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

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

   In this example, you would click on Cluster 1. Note that that cluster was automatically named `cluster-clu1` in the Cluster Name field, as shown in the following figure.

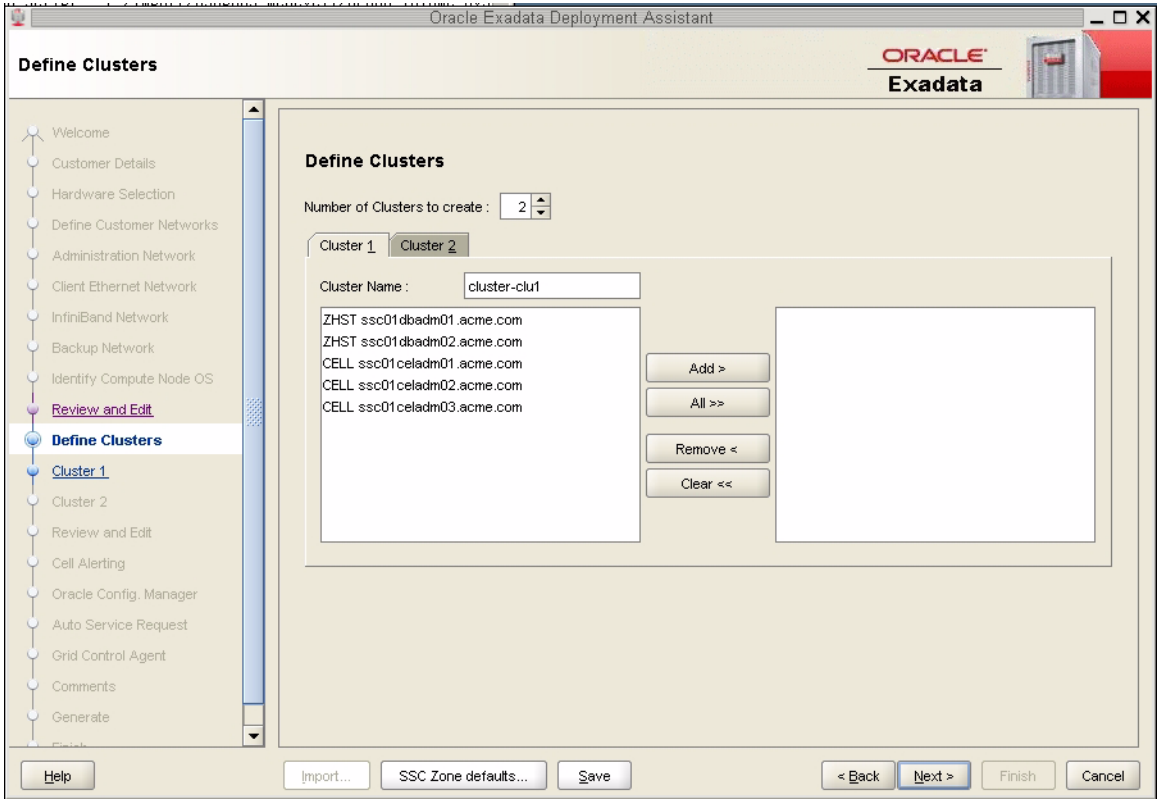

   b. **Set a unique name within your organization for the cluster.**

   By default, a cluster name of `cluster-clu`*number* will be assigned to the cluster (such as `cluster-clu1`). You may run into an issue with Enterprise Manager if you have multiple systems with the same default cluster names across the systems, so you must change the default cluster name to a unique name within your organization for the cluster (for example, *organization-systemname-clusternumber*, such as `acme-ssc01-clu1`).

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

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

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

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

Note – If you are sharing Exadata Storage Servers across clusters, you must first confirm that the Exadata Storage Servers that you are adding to this cluster have enough available space to use for these new zones. See "Verify Exadata Storage Server Disk Space for Additional Zones" on page 65 for more information.

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

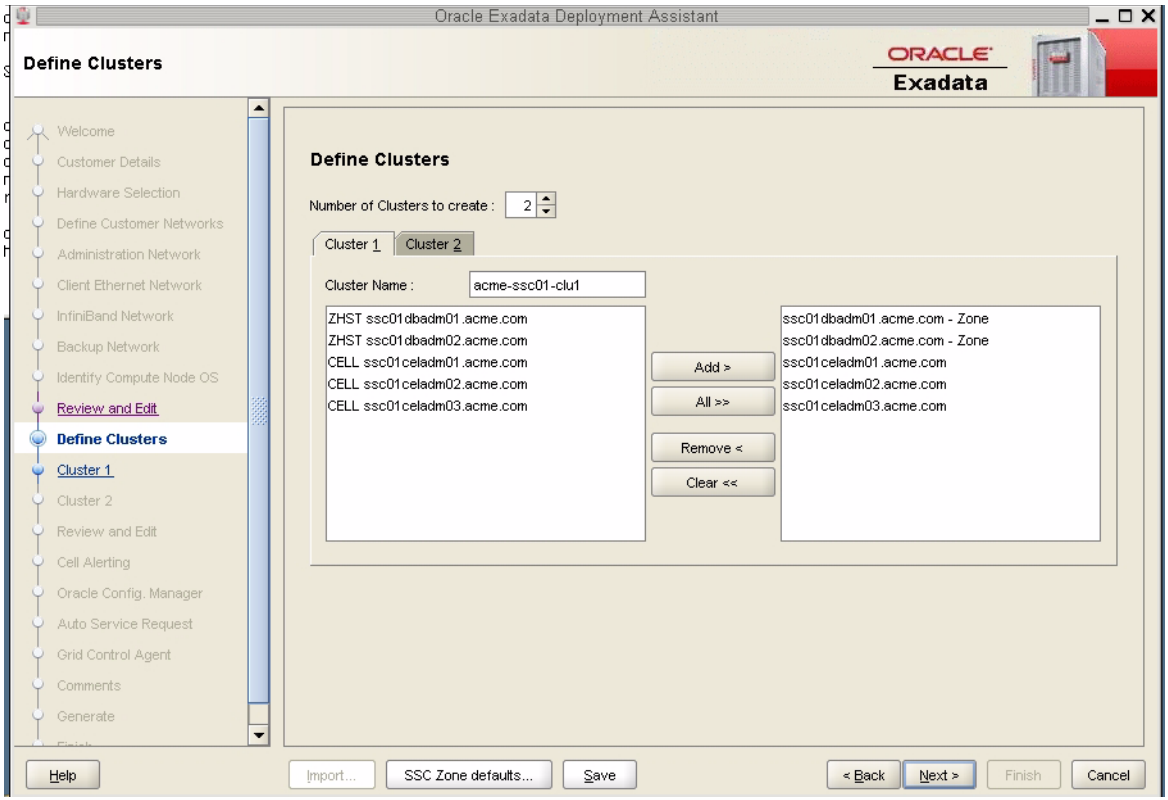

3. **Repeat** Step 2 **to all of the new clusters that you will be creating to assign a unique name and select the Exadata Storage Servers for each cluster.**

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

   The Cluster page for the first new cluster is displayed.

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

   Go to "Set Zone Default Configurations" on page 93.

▼ Set Zone Default Configurations

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

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

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

   The SSC Zone Defaults page appears.

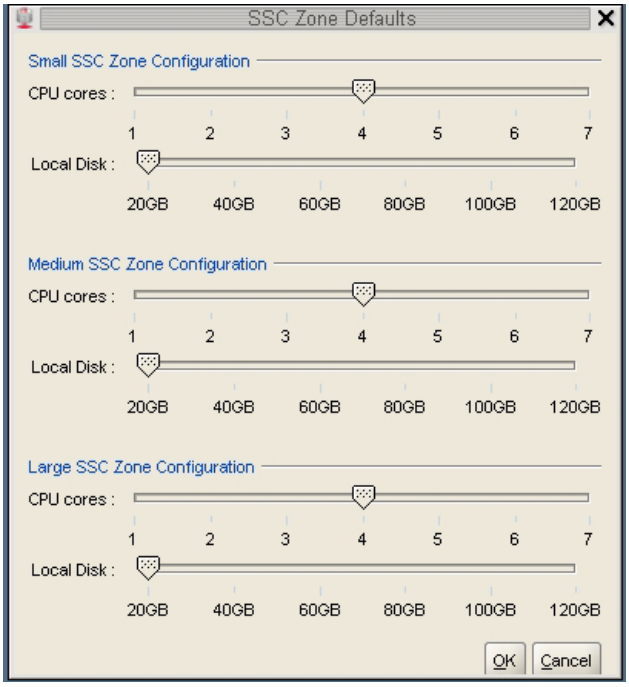

2. **Review the default values for the number of CPU cores and the disk space to be allocated for the small, medium, and large zone configurations.**

   The following default values are used for the zone configurations:

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

- 7 cores
- 20 GB of disk space

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

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

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

   - The following values are available for both options:
     - **Number of CPU cores**: Between 1 and 7 cores
     - **Amount of disk space**: Between 20 GB and 120 GB of disk space
   - When deciding on the number of cores to be allocated, keep in mind that two to four cores will be reserved for the global zone (the Database Domain) and the remaining cores are available for the zones in that Database Domain. See "Determining the Cores Available for Domains and Zones" on page 3 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 96. The values you assign to each of these types of zone configurations are completely user-definable, so you can assign the same values to all three types of zone configurations if you wish.

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

   - 4 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 96.

▼ 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.**

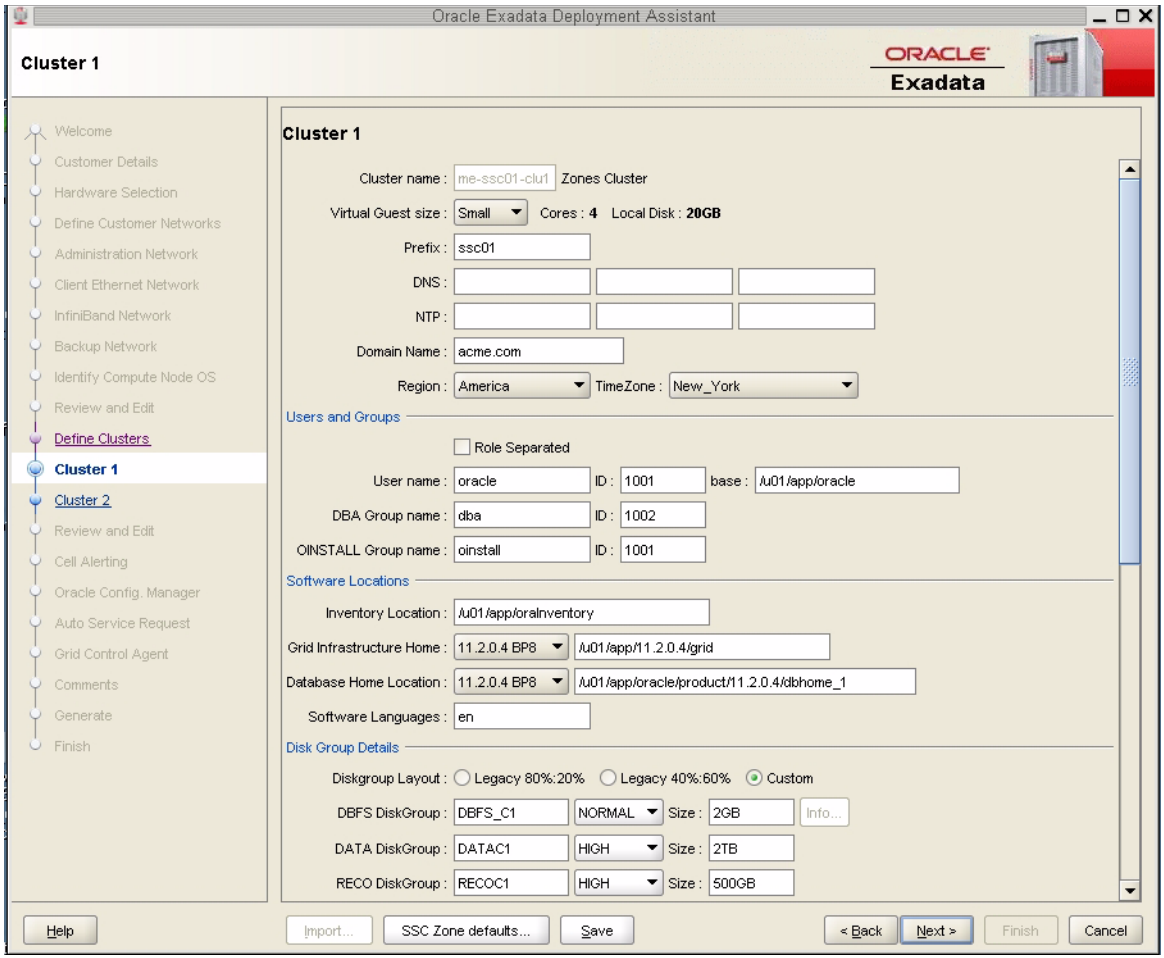

2. **Determine the number of cores that are available in the Database Domains with zones or Database Domains (without zones) for this cluster.**

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 3 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 93.

Note – You must make a selection in the Virtual Guest size area, even if the zone configuration size that you want is shown on the screen by default. The values for the number of cores and the amount of local disk space will not update to the values that you assigned in the section "Set Zone Default Configurations" on page 93 unless you make a selection in the Virtual Guest size area.

---

5. **In the Software Locations area, select the database versions for this cluster in the following areas:**
   - Grid Infrastructure Home
   - Database Home Location
6. **In the Disk Group Details area, provide the following information:**
   - A selection in the Diskgroup Layout area:
     - `Legacy 80%:20%` : Assigns the following sizes to the disk groups:

       - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)

       - DATA Diskgroup: 80% size

       - RECO Diskgroup: 20% size

     - `Legacy 40%:60%` : Assigns the following sizes to the disk groups:

       - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)

       - DATA Diskgroup: 40% size

       - RECO Diskgroup: 60% size

     - `Custom` : Allows you to assign your own sizes to the disk groups
   - Unique name for the disk groups (DBFS, DATA, and RECO) for this cluster. For example, for the first new cluster, which is cluster 1:
     - DBFS Diskgroup: `DBFS_DGC1`
     - DATA Diskgroup: `DATAC1`
     - RECO Diskgroup: `RECOC1`

- 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 65 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.**
8. **Change the Start IP address to the first available, unused IP address for each of the networks.**

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

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

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

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

   For example, if the first available starting IP address for the client network for the first new cluster is `172.16.8.150`, and the Pool size field next to the Start IP field for that network shows 7, then IP addresses `172.16.8.150` to `172.16.8.156` is assigned to the client access network for this cluster. Therefore, `172.16.8.157` is the next available IP address for the client network for the next cluster, if necessary.
9. **Review the rest of the information in this page and make changes as necessary.**
10. **Click Next when you have completed the page for this cluster.**

    The Cluster page for the next cluster that you want to create is displayed.
11. **Repeat Step 1 through Step 10 for each cluster.**
12. **Click Next when you have set up the last new cluster.**

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

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

    Go to "Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page" on page 100.

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

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

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

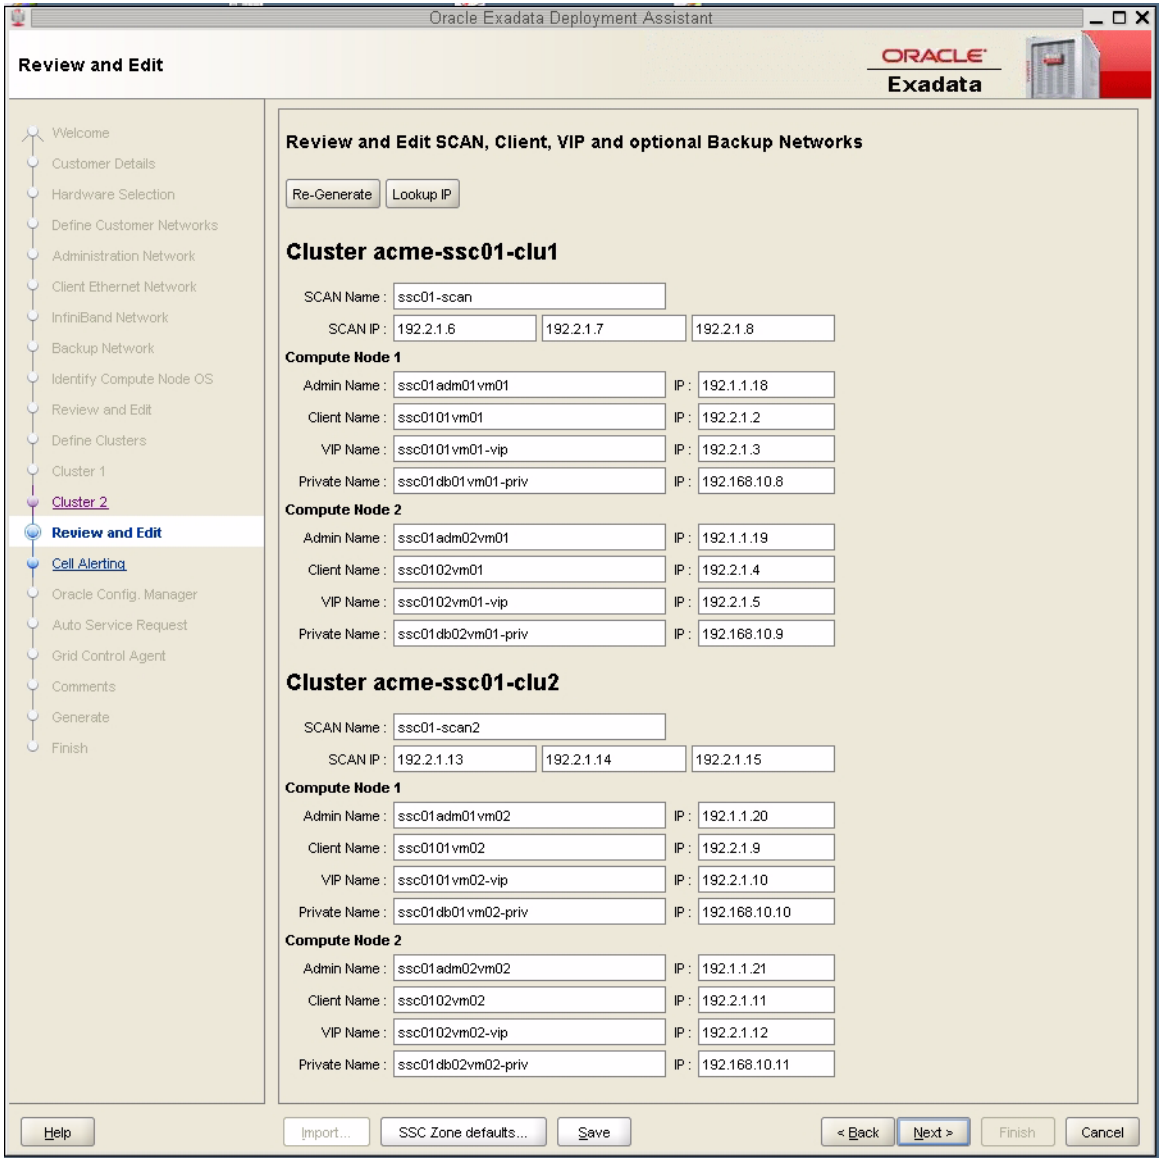

2. **Click Next.**

   The Cell Alerting page is displayed.

Note – If an error message appears when you click Next in the Review and Edit SCAN, Client, VIP and Optional Backup Networks page, indicating that there are duplicate IP addreses, 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 102.

▼ 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 102.

▼ Generate the Configuration Files

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

   A pop-up window appears.

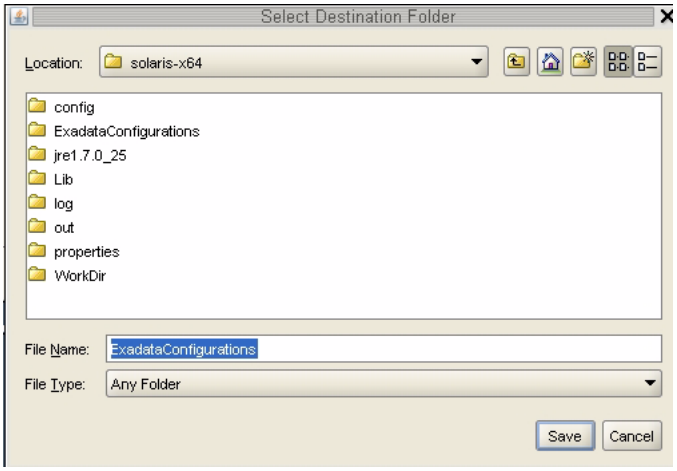

2. **Navigate to the appropriate directory, enter the name for the configuration folder for this deployment, and click Save.**

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

3. **Click Finish.**

   The OEDA closes.

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

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

   - Individual configuration files for every zone that was set up on Database Domains with zones (shown as *name-prefix-zone-name*`-zoneConfig.xml`)
   - Individual configuration files for every cluster that was set up, either for zones or for Database Domains without zones (shown as *customer-name*`-`*name-prefix*`-cluster-clu`*cluster-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, and you provided the following information during your initial installation:

   - Customer name: Acme
   - Name prefix: ssc01

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

- Individual configuration files for the zones that were set up on Database Domains with zones, such as:
  - `ssc01zdbadm01-zoneConfig.xml`: Configuration file for the first zone
  - `ssc01zdbadm02-zoneConfig.xml`: Configuration file for the second zone
  - `ssc01zdbadm03-zoneConfig.xml`: Configuration file for the third zone
- Individual configuration files for the clusters that were set up, such as:
  - `Acme-ssc01-cluster-clu1.xml`: Configuration file for the first cluster (the first cluster with zones)
  - `Acme-ssc01-cluster-clu2.xml`: Configuration file for the second cluster (the second cluster with zones)
  - `Acme-ssc01-cluster-clu3.xml`: Configuration file for the third cluster (the third cluster with zones)
- A single master configuration file that contains all the configuration information for your system, such as:
  - `Acme-ssc01.xml`

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

---

5. Determine whether you want to create more clusters at this time.

- If you want to create more clusters at this time, now that you have gone through one complete pass at entering information with the new OEDA, you can now use the more automated procedures provided in "Creating Configuration Files Using OEDA: Automated Procedures" on page 105, even though your system is technically at build level 1.5.4 or earlier.
- If you do not want to create more clusters at this time, go to "Creating the Template Zone on Each Database Domain" on page 129.

Creating Configuration Files Using OEDA: Automated Procedures

Note – The procedures in this section apply only if you have gone through the manual procedures for inputing information and generating configuration files using the latest OEDA, as described in "Creating Configuration Files Using OEDA: Manual Procedures" on page 67. If you are unsure

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

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

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

- "Complete the Cluster Page" on page 121
- "Complete the Cluster Review and Edit SCAN, Client, VIP and Optional Backup Networks Page" on page 125
- "Verify Remaining Configuration Information" on page 127
- "Generate the Configuration Files" on page 127

▼ Locate the Necessary Files to Set Up Zones

1. **Locate the folder that contains the most recent configuration files.**

   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-clu`*cluster-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, and you provided the following information during your initial installation:

   - Customer name: Acme
   - Name prefix: ssc01

   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:
     - `ssc01zdbadm01-zoneConfig.xml`: Configuration file for the first zone
     - `ssc01zdbadm02-zoneConfig.xml`: Configuration file for the second zone
     - `ssc01zdbadm03-zoneConfig.xml`: Configuration file for the third zone

     And so on for the configuration files for the remaining five zones.
   - Individual configuration files for the five clusters that were set up:

- `Acme-ssc01-cluster-clu1.xml`: Configuration file for the first cluster (the first cluster with zones)
- `Acme-ssc01-cluster-clu2.xml`: Configuration file for the second cluster (the second cluster with zones)
- `Acme-ssc01-cluster-clu3.xml`: Configuration file for the third cluster (the third cluster with zones)
- `Acme-ssc01-cluster-clu4.xml`: Configuration file for the fourth cluster (the fourth cluster with zones)
- `Acme-ssc01-cluster-clu5.xml`: Configuration file for the fifth cluster (the cluster for the Database Domains without zones)

- A single master configuration file that contains all the configuration information for your system:
  - `Acme-ssc01.xml`

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

▼ Set Up the OEDA

1. **Prepare to configure zones.**

   See "Preparing to Configure Zones on Database Domains" on page 51.

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

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

   The OEDA files are part of the `Ocmd` ZIP file that you downloaded earlier in this process. See "Verify Configuration Tool Installation" on page 62 for more information.

   There are different OEDA ZIP files for these operating systems:

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

4. **Unzip the `Ocmd` ZIP file for your operating system on the first Database Domain that will deploy zones.**

5. **Go into the directory that was created after you unzipped the `Ocmd` ZIP file for your operating system.**

6. **Start the OEDA.**

   Go to "Start the OEDA" on page 108.

▼ 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.

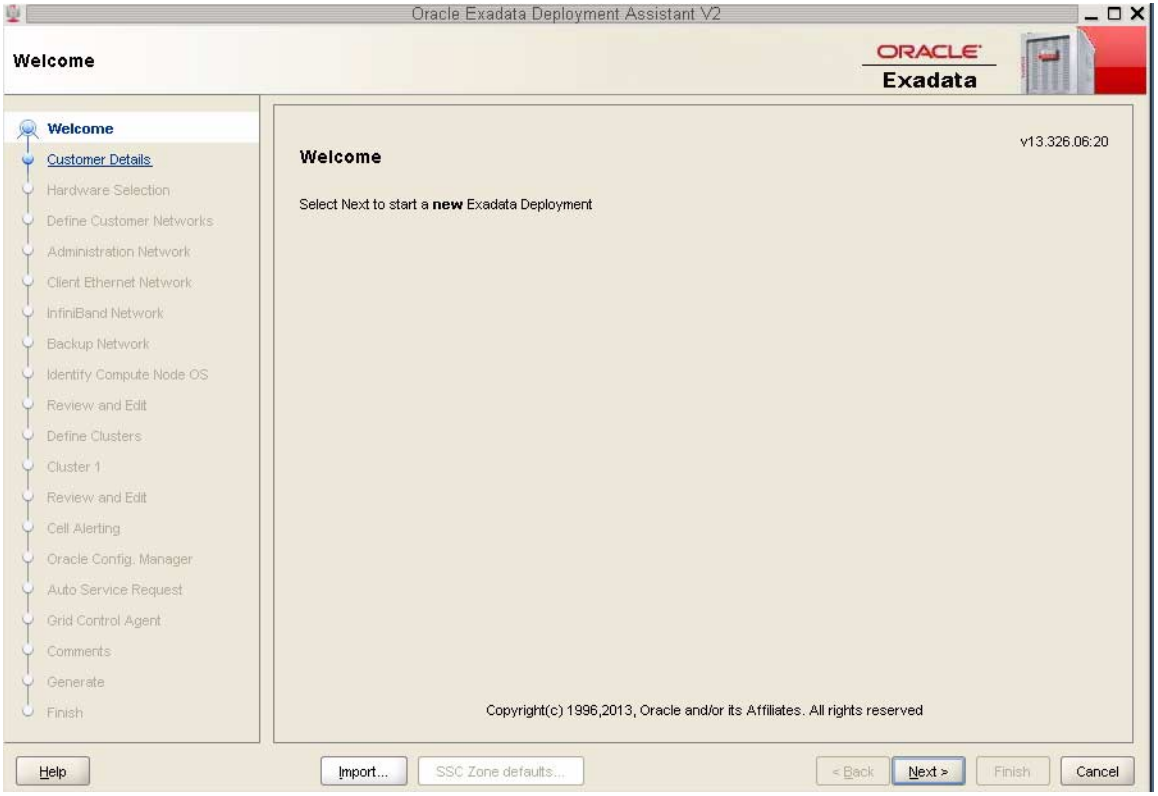

2. **Import the OEDA XML input file.**

   Go to "Import the Most Recent OEDA Configuration File" on page 109.

▼ 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.**

   See "Locate the Necessary Files to Set Up Zones" on page 106 for more information.

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

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

   - **Customer name**: Acme

   - **Name prefix**: ssc01

   Then the configuration file for this cluster would be named `Acme.ssc01.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.

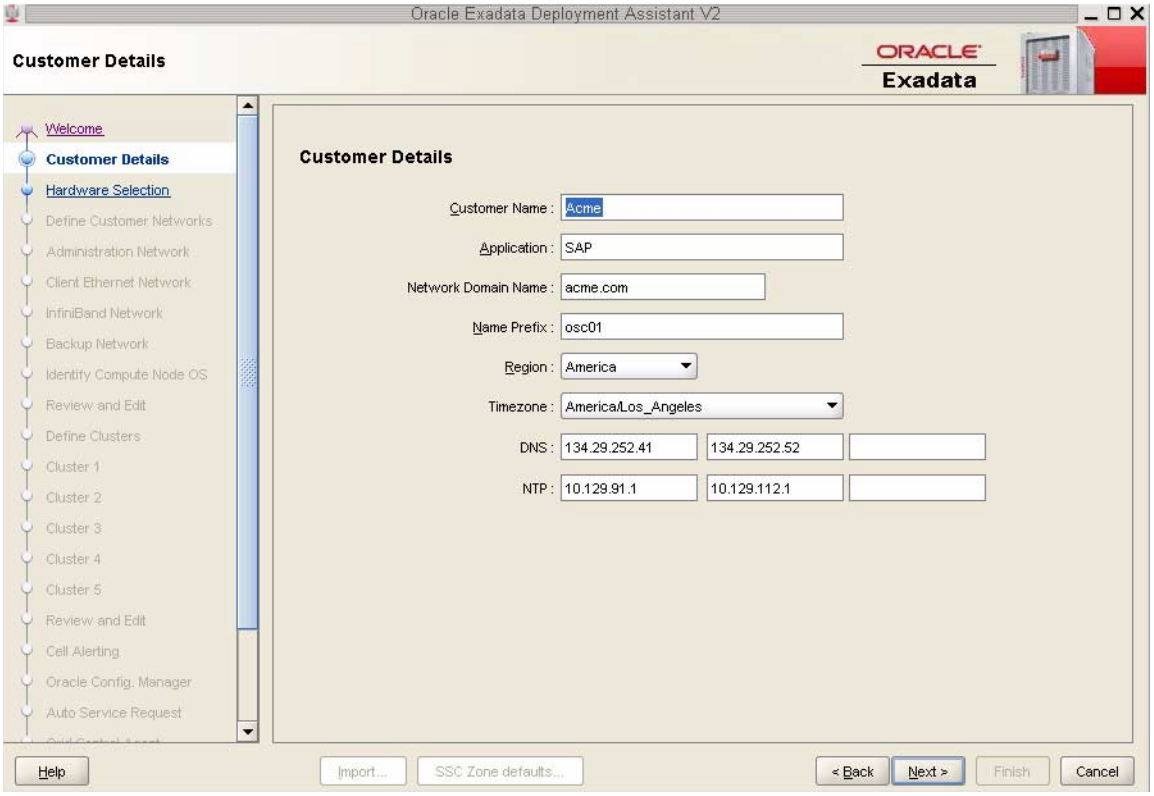

5. **Verify the existing configuration information.**

   Go to "Review Existing Configuration Information" on page 110.

▼ 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 SPARC SuperCluster T4-4. 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 111.

▼ 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 SPARC SuperCluster T4-4 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:

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

Then you should see the following screen:

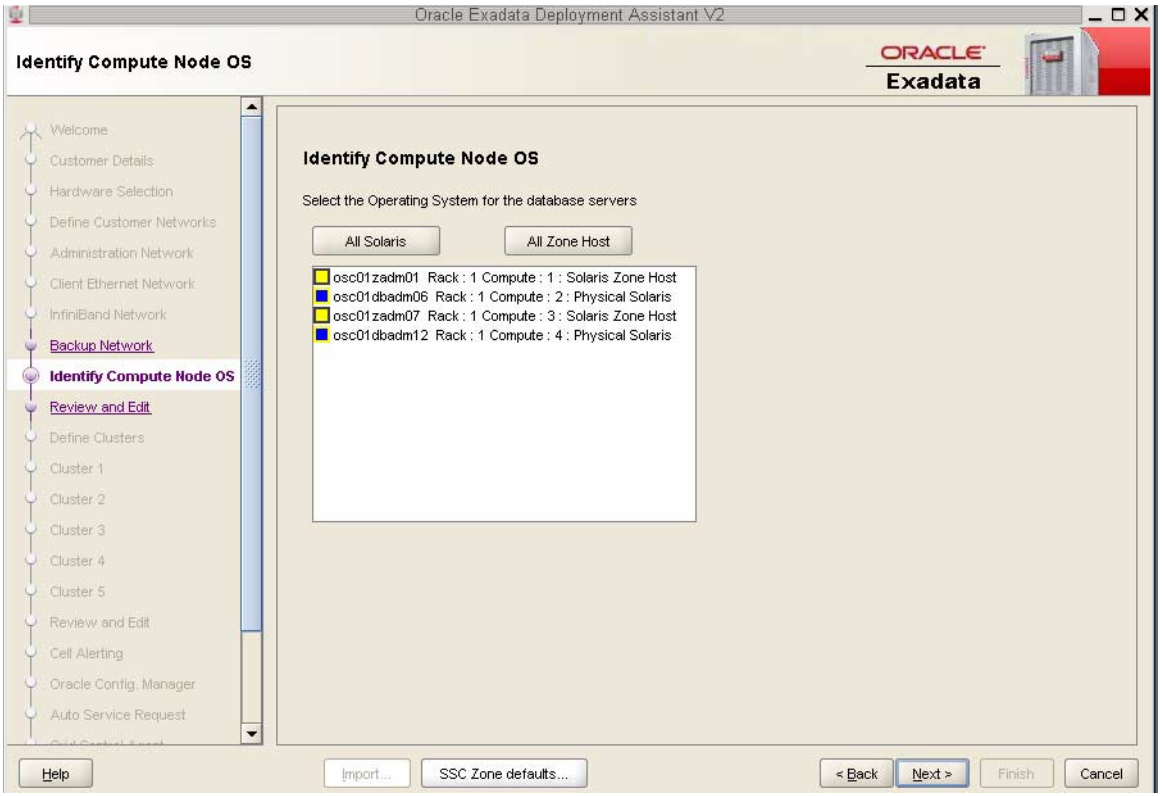

In this screen, these domains are displayed in this manner:

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

---

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

---

2. **Click Next.**

   The Review and Edit page is displayed.

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

   Go to "Review the Information in the Management and Private Networks Page" on page 113.

▼ Review the Information in the Management and Private Networks Page

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

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

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

- Database Domains with zones (shown as `Compute Node # - Zone Host`)
- Database Domains (without zones) (shown as `Compute Node #`)
- Exadata Storage Servers in the SPARC SuperCluster T4-4
- Switches and PDUs in the SPARC SuperCluster T4-4

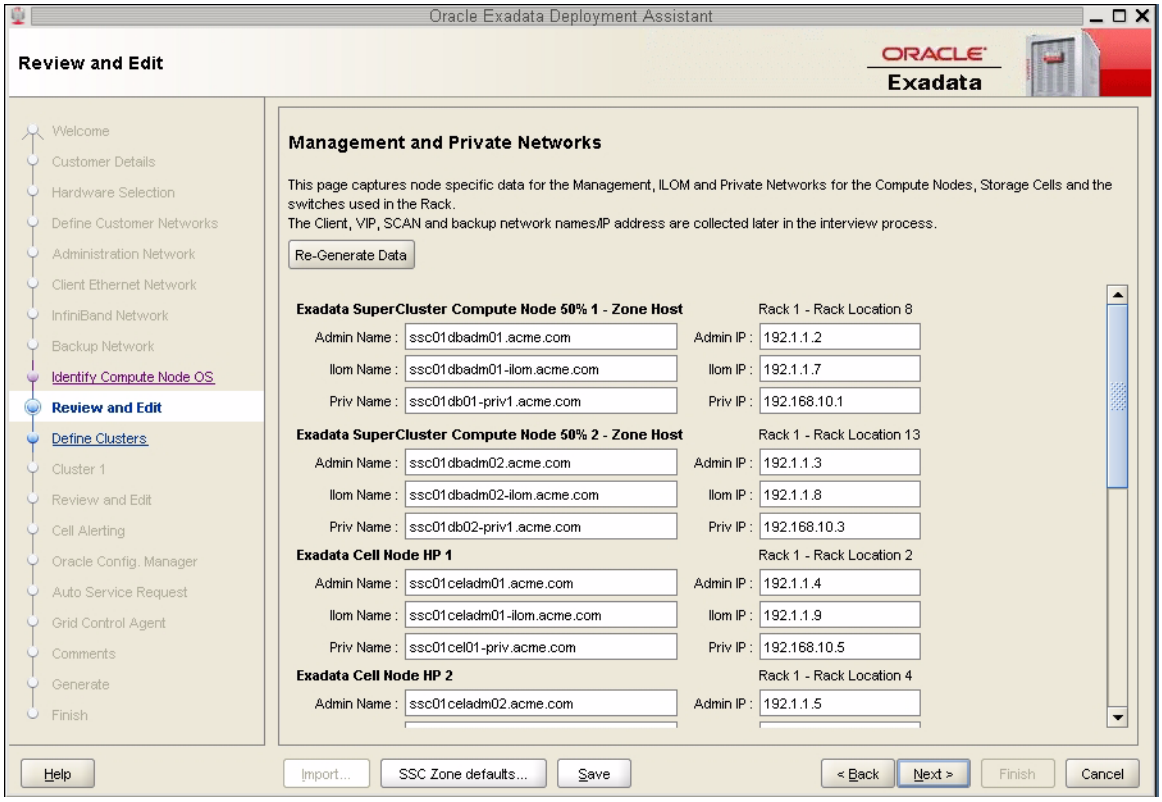

2. **Click Next after reviewing the information in this page.**

   The Define Clusters page is displayed.

3. **Complete the Define Clusters page.**

   Go to "Complete the Define Clusters Page" on page 114.

▼ 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 two clusters that were set up for the Database Domains with zones (two zones within each Database Domain with zones, so two clusters for the zones in the Database Domains with zones). Then you should see the following screen:

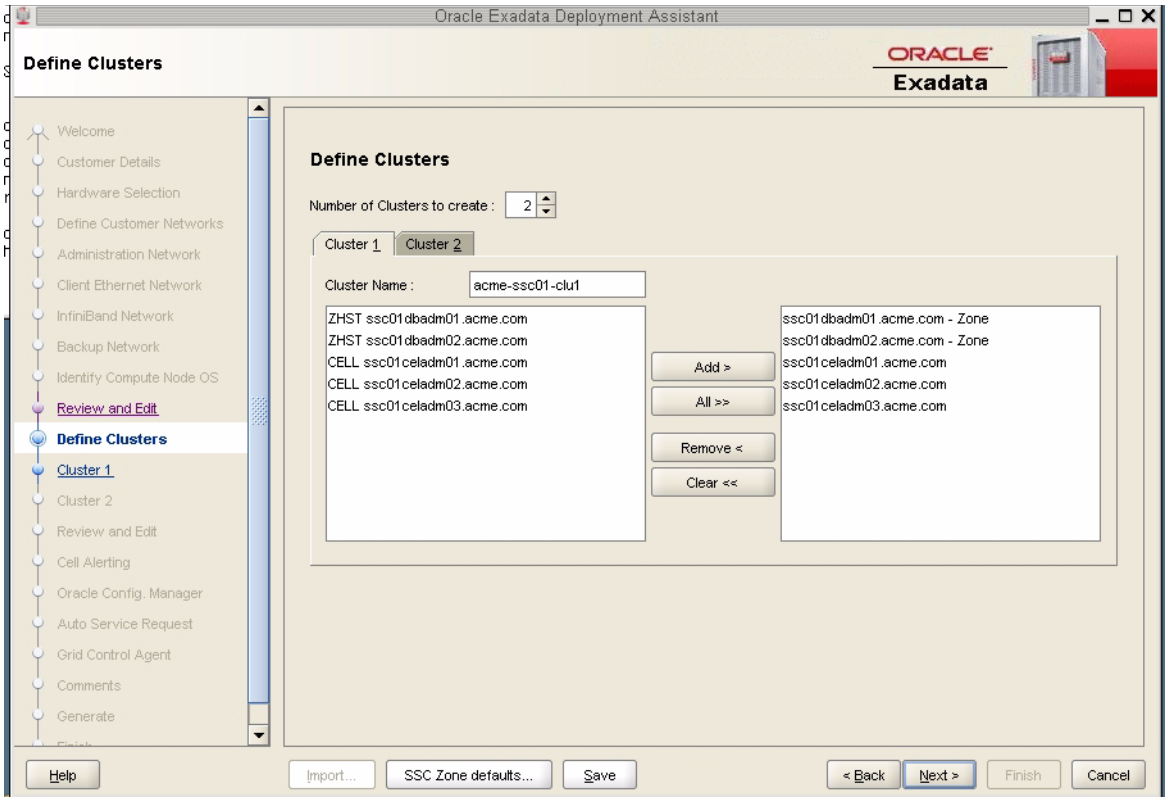

Where Clusters 1 and 2 should display the clusters that were set up for the Database Domains with zones (two zones within each Database Domain with zones, so two clusters for the zones in the Database Domains with zones).

For Clusters 1 and 2, 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 and 2 show the following information:

- `ssc01dbadm01.acme.com`
- `ssc01dbadm02.acme.com`

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

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 2 and note that the cluster name is `acme-ssc01-clu2`.

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

   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 four total clusters:

   - Two 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 3. Note that that cluster was automatically named `cluster-clu3` in the Cluster Name field, as shown in the following figure.

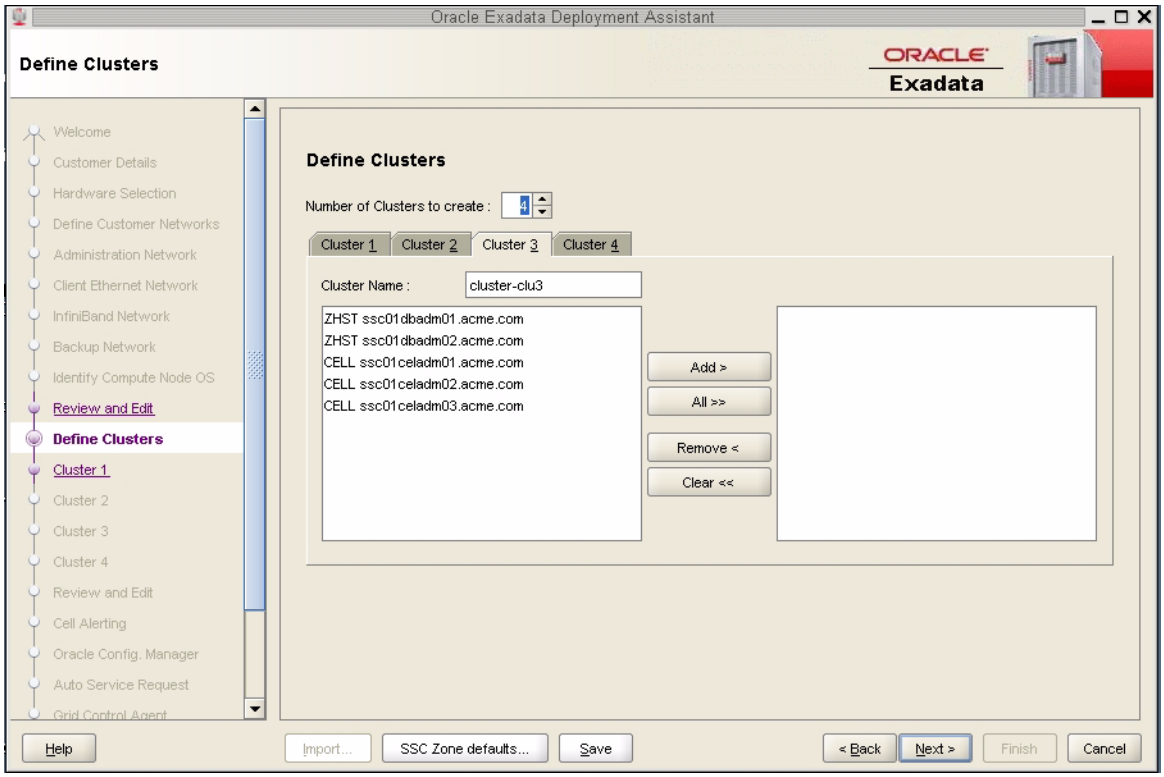

b. Set a unique name within your organization for the cluster.

By default, a cluster name of `cluster-clu`*number* will be assigned to the cluster (such as `cluster-clu3`). You may run into an issue with Enterprise Manager if you have multiple systems with the same default cluster names across the systems, so you must change the default cluster name to a unique name within your organization for the cluster (for example, *organization-systemname-clusternumber*, such as `acme-ssc01-clu3`).

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 65 for more information.

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

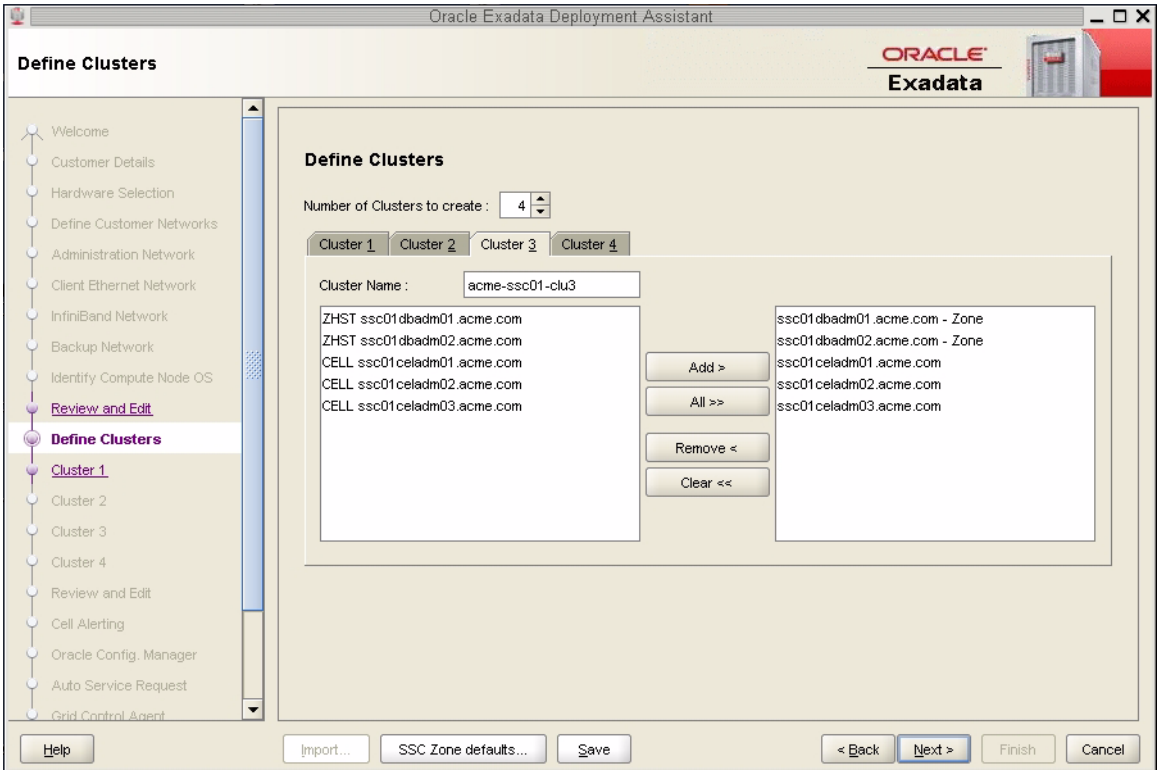

6. **Repeat Step 5 to assign the Exadata Storage Servers to all of the new clusters that you will be creating.**

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 119.

▼ Set Zone Default Configurations

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

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

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

   The SSC Zone Defaults page appears.

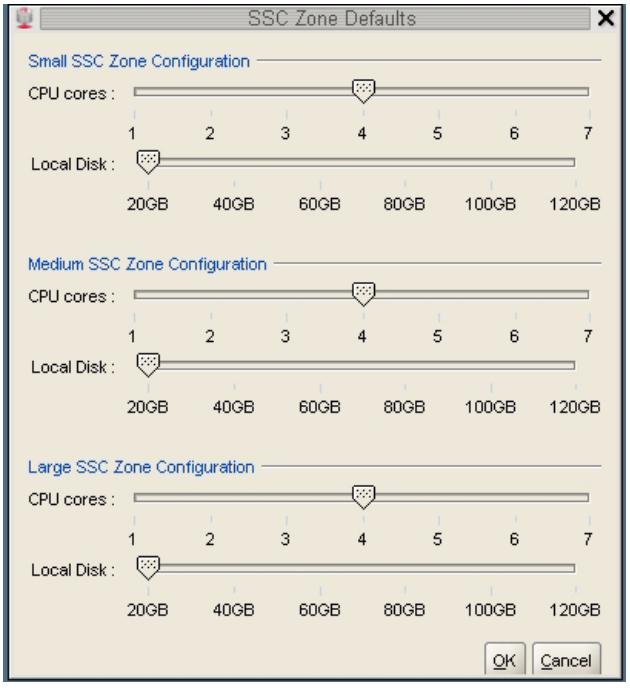

2. **Review the default values for the number of CPU cores and the disk space to be allocated for the small, medium, and large zone configurations.**

   The following default values are used for the zone configurations:

   - Small zone configuration:
     - 1 core

- 20 GB of disk space
- Medium zone configuration:
  - 4 cores
  - 20 GB of disk space
- Large zone configuration:
  - 7 cores
  - 20 GB of disk space

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

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

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

   - The following values are available for both options:
     - **Number of CPU cores**: Between 1 and 7 cores
     - **Amount of disk space**: Between 20 GB and 120 GB of disk space
   - When deciding on the number of cores to be allocated, keep in mind that two to four cores will be reserved for the global zone (the Database Domain) and the remaining cores are available for the zones in that Database Domain. See "Determining the Cores Available for Domains and Zones" on page 3 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 121. 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:

- 4 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 121.

▼ 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, assume you had two clusters created previously. Review the cluster configuration page for Cluster 3, which is the first new cluster that you are creating. Skip over the cluster configuraton pages for the first two clusters, which were already set up as part of your initial installation, until you get to the cluster configuration page for Cluster 3.

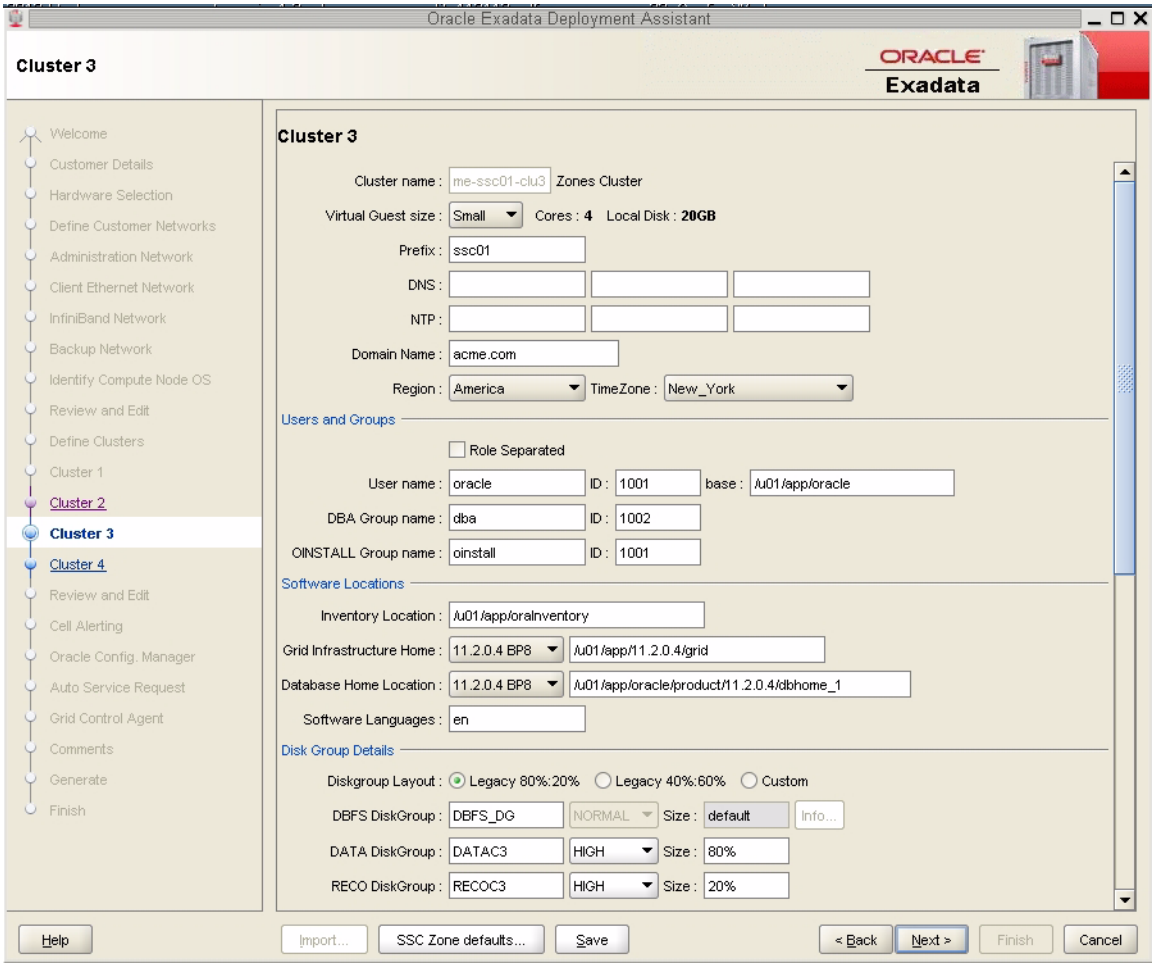

2. **Determine the number of cores that are available in the Database Domains with zones or Database Domains (without zones) for this cluster.**

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 3 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 119.

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 119, unless you make a selection in the Virtual Guest size area.

---

5. **In the Software Locations area, select the database versions for this cluster in the following areas:**
   - Grid Infrastructure Home
   - Database Home Location
6. **In the Disk Group Details area, provide the following information:**
   - A selection in the Diskgroup Layout area:
     - `Legacy 80%:20%` : Assigns the following sizes to the disk groups:

       - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)

       - DATA Diskgroup: 80% size

       - RECO Diskgroup: 20% size

     - `Legacy 40%:60%` : Assigns the following sizes to the disk groups:

       - DBFS Diskgroup: Default size (the default size for the DBFS disk group in this selection is the size of the operating system disk slice on disks 0 and 1, which is usually between 29 GB and 31 GB)

       - DATA Diskgroup: 40% size

       - RECO Diskgroup: 60% size

     - `Custom` : Allows you to assign your own sizes to the disk groups
   - Unique name for the disk groups (DBFS, DATA, and RECO) for this cluster. For example, for the first new cluster, which is cluster 6:
     - DBFS Diskgroup: `DBFS_DGC6`
     - DATA Diskgroup: `DATAC6`
     - RECO Diskgroup: `RECOC6`

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

---

Note – If you are sharing Exadata Storage Servers across clusters, follow the instructions in "Verify Exadata Storage Server Disk Space for Additional Zones" on page 65 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 125.

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

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

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

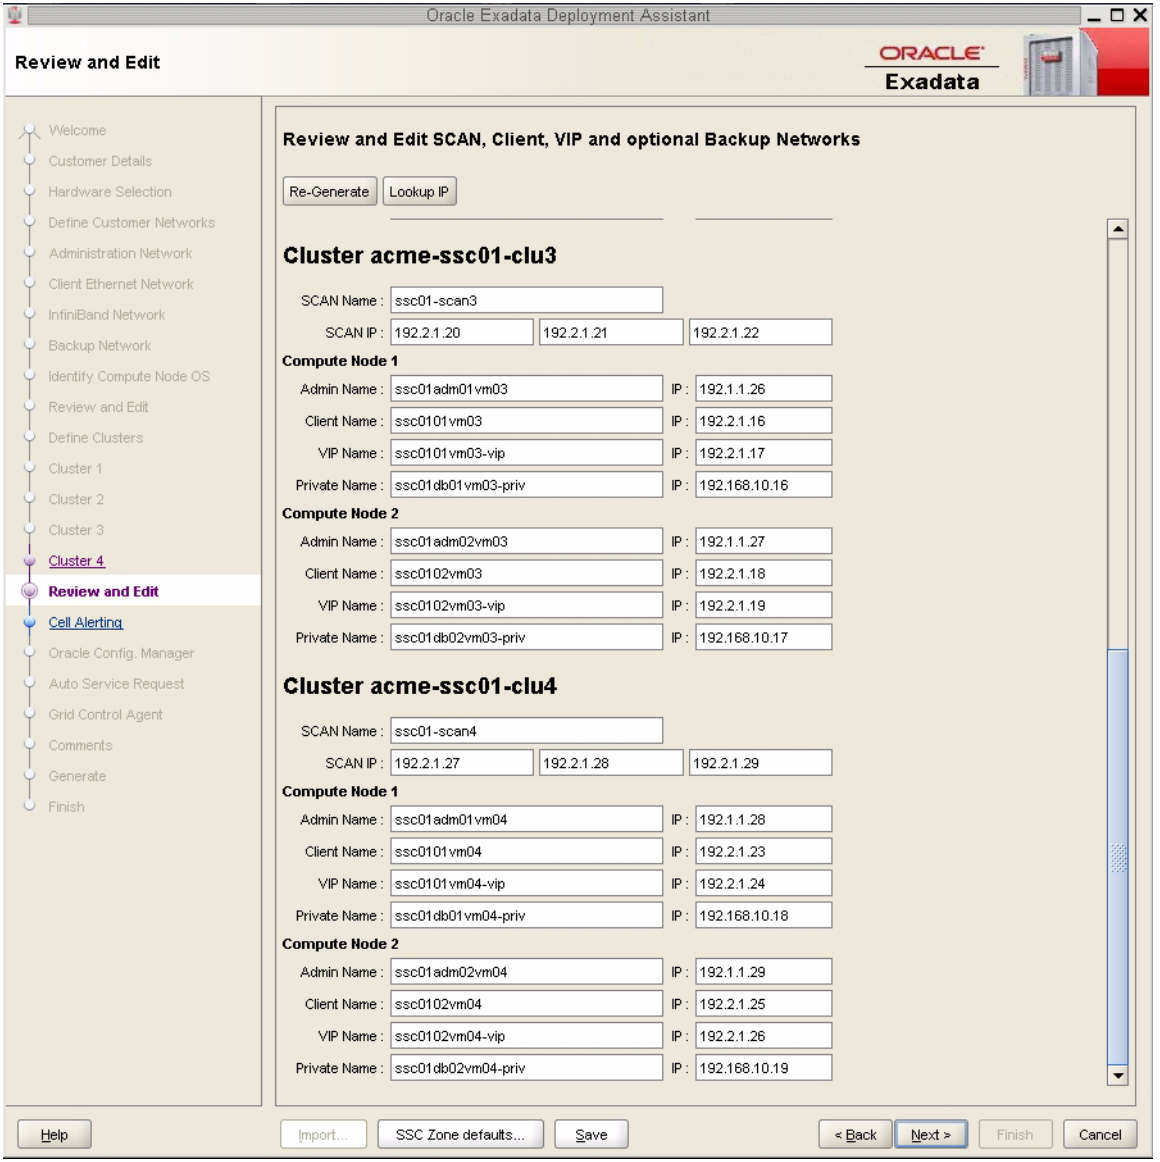

2. **Click Next.**

   The Cell Alerting page is displayed.

Note – If an error message appears when you click Next in the Review and Edit SCAN, Client, VIP and Optional Backup Networks page, indicating that there are duplicate IP addreses, 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 127.

▼ 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 127.

▼ Generate the Configuration Files

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

   A pop-up window appears.

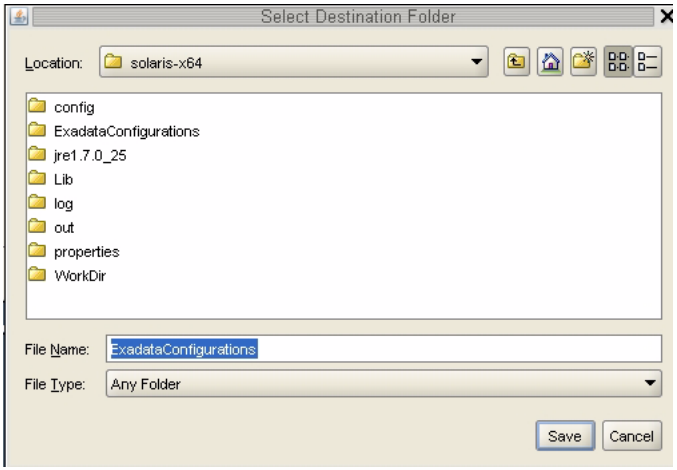

2. **Navigate to the appropriate directory, enter the name for the configuration folder for this deployment, and click Save.**

   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-clu`*cluster-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 129.

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 129
- "Delete a Template Zone From a Database Domain" on page 132

▼ Create a Template Zone on a Database Domain

1. **Log in to the first Database Domain in your SPARC SuperCluster T4-4 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, ssc01dbadmin01-01_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 ssc01dbadmin01.us.oracle.com running
/zoneHome/ssc01dbadmin01.us.oracle.com solaris excl
 - ssc01dbadmin01-01_T installed /zoneHomeT/ssc01dbadmin01-01_T solaris excl
```

  You do not have to create a template zone on the Database Domain in this case. Go to "Creating VNETs" on page 135.

- 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 62 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/ssc01dbadmin01-01-20121129093
659.log and traces in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/ssc01dbadmin01-01-20121129093
659.trc and o/p in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/ssc01dbadmin01-01-20121129093
659.out
INFO: Begin Template Creation
INFO: Checking for ssc01dbadmin01-01_T
INFO: Checking for ssc01dbadmin01-01_T
INFO: Creating a template Virtual Guest for future use
INFO: 2012-11-29 09:37:14 : Executing zfs create rpool/ssc01dbadmin01-01_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/ssc01dbadmin01-01_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 : ssc01dbadmin01-01_T  Data file :: none
```

   The name for the template zone for this particular domain is *DB-domain*_T (for example, ssc01dbadmin01-01_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 62 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
-  ssc01dbadmin01-01_T  installed  /zoneHomeT/ssc01dbadmin01-01_T    solaris  excl
```

   The template zone is shown as *DB-domain*`_T` (for example, in the example output above, `ssc01dbadmin01-01_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/ssc01dbadmin01-01-20121130163
540.log and traces in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/ssc01dbadmin01-01-20121130163
540.trc and o/p in
/opt/oracle.supercluster/zonetools/ssc_exavm/tmp/ssc01dbadmin01-01-20121130163
540.out
INFO: Destroying Template Virtual Guest
INFO: Destroying Template ...
INFO: Destroying Virtual Guests ....
INFO: 2012-11-30 16:35:40 : Executing zoneadm -z ssc01dbadmin01-01_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
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
```

Creating VNETs

The following procedures describe how to create or delete VNETs.

Determining if Additional VNETs Are Needed for a Database Domain

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

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

You must create additional VNETs if you are setting up zones on the following Database Domains on your SPARC T4-4 server:

- The third Database Domain on a SPARC T4-4 server that has three total domains (Database Domains or Application Domains)

- The second and fourth Database Domain on a SPARC T4-4 server that has four total domains (Database Domains or Application Domains)

Refer to the information that you entered in the worksheets in "Recording Your Existing Configuration" on page 16 to determine how many total domains you have on each SPARC T4-4 server, and which of those domains are Database Domains.

- **Using the information provided above, determine if you need to create additional VNETs:**
  - If none of these situations apply for you, you do not have to create additional VNETs. Go to "Running the Java OneCommand" on page 143.
  - If you are setting up zones on Database Domains that are listed above, then you must create additional VNETs for the zones that you will be setting up on these Database Domains. Go to "Create VNETs for Zones" on page 138.

▼ Determine if Additional VNETs Are Needed (Software Commands)

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

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

```
# ipmpstat -g
```

You might see output similar to the following:

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

3. **Note the information displayed in the INTERFACES column from the output.**

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

4. **Type:**

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

- If you do not see any output (if no output was returned after entering the command above), then you do not have to create additional VNETs. Go to "Running the Java OneCommand" on page 143.

- If you see output similar to the following, then you must create additional VNETs for the zones that you will be setting up on these Database Domains:

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

where the `net#` interfaces displayed in the `ipmpstat -g` command is shown with VNETs associated with those `net#` interfaces (in the above example, with vnet0 associated with the net2 interface and vnet1 associated with the net6 interface). Go to "Create VNETs for Zones" on page 138.

▼ Locate the Control Domain

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

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

2. **Type:**

```
# virtinfo -a
```

Output similar to the following appears:

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

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

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

▼ Create VNETs for Zones

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

   See "Preparing to Configure Zones on Database Domains" on page 51 for the location of the `addvnet-wrapper.sh` script.

2. **On the first SPARC T4-4 server, list the Database Domains with zones to locate the Database Domain that you are adding VNETs to:**

```
ldm list
```

3. **Log in to the control domain.**

   See "Locate the Control Domain" on page 137 for instructions on locating the control domain.

4. **From the control domain, add VNETs to the guest domain on this SPARC T4-4 server:**

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

   where:

   - *DB-domain* is the name of the guest domain.
   - *num-zone* is the number of zones that you are setting up on the Database Domains in this particular SPARC T4-4 server.

   For example, if you have the following setup:

   - The name of the guest domain is `ssc01dbadmin02-01`.
   - You will be setting up 4 zones on the Database Domains on this SPARC T4-4 server.

   Then you would type the following command:

```
addvnet-wrapper.sh ssc01dbadmin02-01 4
Adding VNETs on ssc01dbadmin02-01 -----| Okay
Info: found active ldom configuration F8_1_1_ML11242013214449
Info: will add new ldom configuration F8_1_1_ML11242013214449
Adding ldom configuration --------/ Okay

INFO: ssh into ssc01dbadmin02-01 for post-vnet operations
INFO: Renamed link net24 to zonenet24
INFO: Renamed link net25 to zonenet25
INFO: Renamed link net26 to zonenet26
INFO: Renamed link net27 to zonenet27
```

- Viewed from the Hypervisor perspective, using the `ldm` command, the VNETs that you added start with `z`, followed by a nonnegative integer, followed by `mgnet`, and ending with either a `0` or a `1` (for example, `z0mgnet1`).
- Viewed from the Solaris perspective, using the `dladm` command, the network created is in the form of `zonenet`*xx*, where *xx* are consecutive numbers (such as `zonenet12`, `zonenet13`, and so on). Once a zone is deployed, the network picked for that particular zone or instance is mapped to `zonbnet`*xx* (such as `zonbnet12`, `zonbnet13`, and so on) to indicate that they are locked and used for the respective zone.

If you destroy a zone using the instructions in "Delete a Template Zone From a Database Domain" on page 132, the net pool is freed and renamed back to `zonenet`*xx* (such as `zonenet12`, `zonenet13`, and so on) so that any future addition of zones using zone tools can pick those VNETs.

5. **Check the guest domains to ensure that the VNETs are created on the specified domain:**

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

where *DB-domain* is the name of the Database Domain with zones that you added the VNETs to. For example, if the name of the Database Domain with zones that you added the VNETs to is `ssc01dbadmin02-01`, type the following:

```
ldm ls -l -p ssc01dbadmin02-01 | grep '^VNET'
VNET|name=mngnet0|dev=network@0|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fa:f8:e8|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=0
VNET|name=mngnet1|dev=network@1|service=mng-primary-net00@primary|mac-addr=
00:14:4f:fb:a0:cd|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=1
VNET|name=z0mgnet0|dev=network@2|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:da:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=2
VNET|name=z0mgnet1|dev=network@3|service=mng-primary-net1@primary|mac-addr=
00:14:4f:f8:a8:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=3
VNET|name=z1mgnet0|dev=network@4|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:83:cc|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=4
VNET|name=z1mgnet1|dev=network@5|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:26:ab|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=5
```

6. **Repeat Step 2 through Step 5 to add VNETs to the control domain on the second SPARC T4-4 server.**

7. **Run the Java OneCommand.**

   Go to "Running the Java OneCommand" on page 143.

▼ Delete VNETs

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

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

---

1. **Locate the `rmvnet.sh` script.**

   See "Preparing to Configure Zones on Database Domains" on page 51 for the location of the `rmvnet.sh` script.

2. **Identify the VNETs that you want to remove:**

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

   where *DB-domain* is the name of the Database Domain that you added the VNETs to. For example, if the name of the Database Domain that you added the VNETs to is `ssc01dbadmin02-01`, you would type the following:

```
ldm ls -l -p ssc01dbadmin02-01 | grep '^VNET'
VNET|name=mgnet0|dev=network@0|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fa:f8:e8|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=0
VNET|name=mgnet1|dev=network@1|service=mng-primary-net00@primary|mac-addr=
00:14:4f:fb:a0:cd|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=1
VNET|name=z0mgnet0|dev=network@2|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:da:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=2
VNET|name=z0mgnet1|dev=network@3|service=mng-primary-net1@primary|mac-addr=
00:14:4f:f8:a8:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=3
VNET|name=z1mgnet0|dev=network@4|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:83:cc|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=4
VNET|name=z1mgnet1|dev=network@5|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:26:ab|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=5
```

3. **Remove the VNETs.**

   - If you want to remove specific VNETs, type:

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

   where:

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

For example, assume you have the following setup:

- The names of the VNETs that you want to remove are `z1mgnet0` and `z1mgnet1`.
- The name of the Database Domain that you are removing the VNETs from is `ssc01dbadmin02-01`.

Then you would type the following:

```
rmvnet.sh -n z1mgnet0,z1mgnet1 ssc01dbadmin02-01
```

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

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

where:

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

For example, assume you have the following setup:

- You want to remove all VNETs with names that begin with `z1` (in the output from Step 2, that would be `z1mgnet0` and `z1mgnet1`).
- The name of the Database Domain that you are removing the VNETs from is `ssc01dbadmin02-01`.

Then you would type the following:

```
rmvnet.sh -p 'z1+' ssc01dbadmin02-01
```

4. **Verify that the VNETs have been removed:**

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

where *DB-domain* is the name of the Database Domain that you removed the VNETs from. For example, if the name of the Database Domain that you removed the VNETs from is `ssc01dbadmin02-01`, type the following:

```
ldm ls -l -p ssc01dbadmin02-01 | grep '^VNET'
VNET|name=mgnet0|dev=network@0|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fa:f8:e8|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=0
VNET|name=mgnet1|dev=network@1|service=mng-primary-net00@primary|mac-addr=
00:14:4f:fb:a0:cd|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=1
VNET|name=z0mgnet0|dev=network@2|service=mng-primary-net1@primary|mac-addr=
00:14:4f:fb:da:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=2
VNET|name=z0mgnet1|dev=network@3|service=mng-primary-net1@primary|mac-addr=
00:14:4f:f8:a8:06|mode=|pvid=1|vid=|mtu=1500|linkprop=|id=3
```

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

Running the Java OneCommand

These topics describe how to run the Java OneCommand.

▼ Set Up Public Key Authentication for ZFS Storage Controllers

Before running the Java OneCommand, you must set up public key authentication on the ZFS storage controllers for every SPARC T4-4 server where you are setting up zones on the Database Domains.

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

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

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

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

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

```
zfs-controller1-hostname:> configuration preferences keys
zfs-controller1-hostname:configuration preferences keys> create
zfs-controller1-hostname:configuration preferences key (uncommitted)> set type=DSA
zfs-controller1-hostname:configuration preferences key (uncommitted)> set key=
"DSA-key-text"
                          key = DSA-key-text == (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> set comment=
"fw-log1"
                          comment = fw-log1 (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> commit
zfs-controller1-hostname:configuration preferences keys> show
Keys:

NAME     MODIFIED              TYPE   COMMENT
key-000  10/12/2009 10:54:58   DSA    fw-log1
```

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

3. **Repeat** Step 2 **for every Database Domain to contain zones.**

4. **Log out of the first ZFS storage controller:**

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

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

---

5. **Log in to the first Database Domain in your system, if you are not logged in already.**

6. **Verify that you set up the public key authentication correctly:**

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

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

You should be able to log in to the ZFS storage controller successfully and see the CLI prompt.

7. **Repeat** Step 5 **and** Step 6 **for every Database Domain in your system.**

▼ Set Up the Java OneCommand

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

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

   By default, the `WorkDir` directory is located in the directory where you downloaded the Java OneCommand patch, as described in "Preparing to Configure Zones on Database Domains" on page 51.

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

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

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

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

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

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

   By default, the Java OneCommand file (`install.sh`) is located in the directory where you downloaded the Java OneCommand patch, as described in "Preparing to Configure Zones on Database Domains" on page 51.

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

   See "Verify Configuration Tool Installation" on page 62 for the location of that script.

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

   `/opt/oracle.supercluster/bin`

7. **Run the Java OneCommand.**

   Go to "Run the Java OneCommand" on page 146.

▼ Run the Java OneCommand

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

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

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

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

   The configuration files for the new clusters that you created should be named *customer-name-name-prefix*`-cluster-clu`*cluster-number*`.xml`. For example:

   - `Acme-ssc01-cluster-clu3.xml`
   - `Acme-ssc01-cluster-clu4.xml`

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

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

For example:

```
./install.sh -cf /path-to-config-file/Acme-ssc01-cluster-clu3.xml -l
```

Information similar to the following is displayed:

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

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

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

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

- Verify IB and calibrate cells (step 5 in the example output above)
- Create cell disks (step 6 in the example output above)
- Configure cell alerting (not shown in the example output above)

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

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

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

```
./install.sh -cf Acme-ssc01-cluster-clu3.xml -s 1
```

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

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

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

```
./install.sh -cf Acme-ssc01-cluster-clu3.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 132 if necessary (for example, if Java OneCommand fails at step 2 for some reason). However, use caution when destroying zones from a Database Domain.

5. **Repeat** Step 2 **through** Step 4 **for each new cluster configuration file that you generated at the end of the OEDA process.**

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

   Go to "Creating Additional Links on the IB Storage Network for Zones" on page 151.

NTP Issue

When your SPARC SuperCluster T4-4 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 151

▼ Create Additional Links on the IB Storage Network for Zones

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

```
# dladm show-ib
```

Output similar to the following appears:

```
LINK        HCAGUID             PORTGUID            PORT    STATE   PKEYS
net24       21280001FCBEBC      21280001FCBEBE      2       up      FFFF
net22       21280001FCC0EC      21280001FCC0EE      2       up      8503,8504,FFFF
net23       21280001FCBEBC      21280001FCBEBD      1       up      8503,8504,FFFF
net21       21280001FCC0EC      21280001FCC0ED      1       up      FFFF
```

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

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

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

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

where:

- *zone_management_hostname* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone
- *first_net_link* is the first network link shown in the output from Step 2
- *second_net_link* is the second network link shown in the output from Step 2

For example:

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

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

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

For example:

```
# zoneadm -z ssc01zdbadm02 reboot
```

6. **Log in to the zone:**

```
# zlogin zone_management_hostname
```

For example:

```
# zlogin ssc01zdbadm02
```

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

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

where:

- *first_net_link* is the first network link shown in the output from Step 2
- *second_net_link* is the second network link shown in the output from Step 2

For example:

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

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

   a. **Enter the following command:**

```
# ipadm show-addr
```

Output similar to the following appears:

```
ADDROBJ           TYPE     STATE        ADDR
lo0/v4            static   ok           127.0.0.1/8
net13/v4          static   ok           169.254.182.77/24
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.**

9. **Create the new IPMP group inside the zone:**

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

where:

- *first_net_link* is the first network link shown in the output from Step 2
- *second_net_link* is the second network link shown in the output from Step 2
- *ipmp_group* is the name of the IPMP group that you are creating
- *ip_address* is an unused IP address from the same subnet as the `stor_ipmp0` interface that you picked in Step 8

For example:

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

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

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

where *second_net_link* is the second network link shown in the output from Step 2.

For example:

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

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

```
# ipmpstat -g
```

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

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

12. **Repeat Step 4 through Step 11 for each zone that you created in this Database Domain.**

Setting Up VLAN Tagging for Zones

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

- "Set Up VLAN Tagging for Zones" on page 155

▼ Set Up VLAN Tagging for Zones

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

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

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

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

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

```
# zoneadm -z ssc01zdbadm02 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 b, highlighted in the following example output:

```
net5              Ethernet             unknown   0      unknown   ixgbe3
net1              Ethernet             up        1000   full      ixgbe1
net6              Ethernet             unknown   0      unknown   ixgbe4
net0              Ethernet             up        1000   full      ixgbe0
net4              Ethernet             up        10000  full      ixgbe2
net7              Ethernet             up        10000  full      ixgbe5
```

7. **Modify the zone configuration:**

```
# zonecfg -z zone_management_hostname
zonecfg: zone_management_hostname> select anet linkname=net_device#1
zonecfg: zone_management_hostname:anet> set vlan-id=VLAN_ID
zonecfg: zone_management_hostname:anet> set lower-link=net_device#1
zonecfg: zone_management_hostname:anet> end
zonecfg: zone_management_hostname> select anet linkname=net_device#2
zonecfg: zone_management_hostname:anet> set vlan-id=VLAN_ID
zonecfg: zone_management_hostname:anet> set lower-link=net_device#2
zonecfg: zone_management_hostname:anet> end
zonecfg: zone_management_hostname> verify
zonecfg: zone_management_hostname> commit
zonecfg: zone_management_hostname> exit
```

where:

- *zone_management_hostname* is the management host name (or Admin Name that was shown in Oracle Exadata Deployment Assistant) for this zone
- *VLAN_ID* is a VLAN ID that you provide
- *net_device#1* is the first client network device that you saw in the output in Step b
- *net_device#2* is the second client network device that you saw in the output in Step b

For example:

```
# zonecfg -z ssc01zdbadm02
zonecfg: ssc01zdbadm02> select anet linkname=net4
zonecfg: ssc01zdbadm02:anet> set vlan-id=13
zonecfg: ssc01zdbadm02:anet> set lower-link=net4
zonecfg: ssc01zdbadm02:anet> end
zonecfg: ssc01zdbadm02> select anet linkname=net7
zonecfg: ssc01zdbadm02:anet> set vlan-id=13
zonecfg: ssc01zdbadm02:anet> set lower-link=net7
zonecfg: ssc01zdbadm02:anet> end
zonecfg: ssc01zdbadm02> verify
zonecfg: ssc01zdbadm02> commit
zonecfg: ssc01zdbadm02> exit
```

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

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

For example:

```
# zoneadm -z ssc01zdbadm02 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