Oracle® Enterprise Manager Ops Center

Feature Reference Guide

12*c* Release 2 (12.2.2.0.0)

E38539-08

February 2015

ORACLE

Oracle Enterprise Manager Ops Center Feature Reference Guide, 12*c* Release 2 (12.2.2.0.0)

E38539-08

Primary Authors: Barbara Higgins, Laura Hartman, Shanthi Srinivasan, Owen Allen, Uma Shankar, Richa Agarwala, Karen Orozco Sanchez, Salvador Esparza Perez

Contents

3 Jobs

6 Storage

7 Networks

8 Plans and Profiles

10 Reports

11 Hardware

12 Operating System Management

13 Operating System Provisioning

14 Operating System Updates

18 Oracle Solaris Zones

19 Oracle VM Server for SPARC

20 Oracle VM Server for x86

21 Server Pools

22 Virtual Datacenters

Preface

The *Oracle® Enterprise Manager Ops Center Feature Reference Guide* describes all features of the Oracle Enterprise Manager Ops Center software.

Audience

This document is intended for users who require a detailed description of features and functionality.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information, see the Oracle Enterprise Manager Ops Center Documentation Library at http://docs.oracle.com/cd/E40871_01/index.htm.

Oracle Enterprise Manager Ops Center provides online Help. Click Help at the top-right corner of any page in the user interface to display the online help window.

For the latest releases of Oracle documentation, check the Oracle Technology Network at: http://www.oracle.com/technetwork/documentation/index.html#em

Conventions

The following text conventions are used in this document:

Convention	Meaning
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.

Convention	Meaning
`monospace`	Monospace type indicates commands within a paragraph, code in examples, text that appears on the screen, or text that you enter.

Part I

Getting Started

Part I contains the following chapters:

1 Introduction

Oracle Enterprise Manager Ops Center is Oracle's comprehensive solution for managing the physical and virtual assets in your data center: operating systems, firmware, BIOS configurations, bare metal server and virtual guest provisioning, hardware monitoring, automatic My Oracle Support service request generation, and performance and energy management.

Oracle Enterprise Manager Ops Center in Your Datacenter

Various sites and various users within each site value different aspects of the Oracle Enterprise Manager Ops Center software:

- As a monitoring tool, the software discovers and identifies all assets in its environment, and displays the status of assets and details of a specific asset. When an incident occurs, you can track the progress of the investigation or service request.
- As a provisioning tool, the software deploys new assets in a manner consistent with existing assets because you use profiles to set the attributes and you use plans to deploy the profiles to one target or many targets. In a similar way, when assets must be upgraded, you use the update profiles and plans to perform the operations.
- As a virtualization manager, the software creates and manages virtual operating systems, virtual servers such as Oracle VM Server for SPARC and Oracle Solaris Zone, and virtual data centers in the cloud. To support these virtual assets, Oracle Enterprise Manager Ops Center manages and provides resources for storage and networks.

Not every product feature is relevant to your site's activities or for your role. What you see in the user interface is affected by several factors:

- **The role attached to your user account**: The actions for your role are available in the Actions pane. The required roles for using a feature are listed in each chapter of this document. When you must accomplish a task and the necessary action is not available to you, your administrator can add the role to your user account. When you log in again, the action is available.
- **The current connection mode**: In Connected mode, actions that rely on OS and firmware images and packages use the latest images and packages downloaded from Oracle and vendor sites. If your site uses the product software in Disconnected mode, the images and packages in your local knowledge base do not change until your site acquires them. Also in Connection mode, you can create service request from incidents with full asset information and warranty status. In

Disconnected mode, you must contact My Oracle Support and provide this information.

Changing the connection mode can be done easily and temporarily. See the *Oracle Enterprise Manager Ops Center Administration Guide* for the procedure for changing the connection mode.

- **The scope of Oracle Enterprise Manager Ops Center**: The product software is designed to manage assets of a data center, from small to large. However, the product software is also installed in an Oracle Engineered System, which is a complete set of integrated hardware and software designed to reach a specific level of capability, capacity, and scale. In this case, the product software is managing the components of the engineered system and the virtual assets that the system supports. Some actions are not relevant to an engineered system and so are not visible in the user interface.

About This Document

After the assets have been discovered and brought under the management of the software, as described in Chapter 2, "Asset Management", learn about the assets by selecting each one or each type and viewing the information in the center pane and its tabs.

When you are ready to perform a task such as discovering a type of hardware or upgrading a server's operating system, go to the How To tab of the documentation library to find the example procedure, which demonstrates one set of options. If you do not find an example of the procedure you want to perform, look in this document for the procedure. Use this document to learn about all the options for the product's operations so that you can determine how you will perform the procedure.

This document is divided into parts and, within each part, the chapters describes the capabilities of the product software's features. Where it is practical, all information about a feature is discussed in the same chapter but there are also links to other documents in the library.

About the Document Library

All documentation for the Oracle Enterprise Manager Ops Center 12*c* Release 2 software is located at the site: `http://docs.oracle.com/cd/E40871_01/index.htm`.

You can use this library from any browser or you can connect to one of the documents or a specific chapter from within the product's user interface. Click the Help option in the title bar to display the Help menu, containing the list of documents, as shown in Figure 1–1.

Figure 1–1 Accessing the Document Library

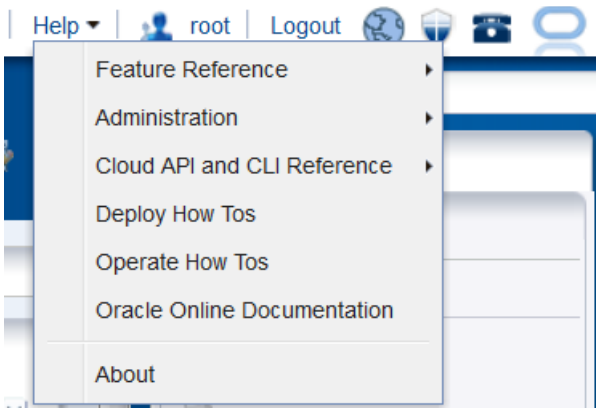

The documentation library contains the documents in Table 1–1. The Deploy How To tab at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` and the Operate How To tab at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm` have links to end-to-end examples and to workflows that combine the examples into deployment and operation scenarios.

Use the site's Search feature to search throughout the library of product documents or to search a specific document. The site can also convert the documents to PDF, EPUB, and Mobipocket file formats.

Table 1–1 Documents for Oracle Enterprise Manager Ops Center

Document	Content
Concepts Guide	An overview of the product software's architecture, its architecture, and its features. This document also explains the product's user interface.
Readme	Links to the installation and upgrade information and a description of known issues in the current release.
Release Notes	Information about the current version, procedures for installation, and known issues.
Installation Guide for Oracle Solaris Operating Systems	Information about planning for a new installation of the product software and the procedure for installing the software on an Oracle Solaris server.
Installation Guide for Linux Operating System	Information about planning for a new installation of the product software and the procedure for installing the software on a Linux server.
Ports and Protocols	Lists the ports used by the product software, the protocol for each port, and its purpose. It also includes the websites that the product software uses.
Upgrade Guide	Information about updating an existing installation of the product software to the current version.
Administration Guide	Procedures for configuring each component of the product software, for configuring the software for high availability, for managing users and roles, and for maintaining the product database. This guide also has procedures for obtaining operating system updates, enabling Auto Service Requests (ASR), using the OCDoctor script, and upgrading the software.
Feature Reference Guide	Descriptions of the product features in detail and with procedures.
Feature Reference Appendix Guide	Reference information such as parameters and variables for customizing the product software.
Command Line Interface Guide	Instructions for using the product's command-line interface and man pages for each command.
Security Guide	Descriptions and procedures for a secure Oracle Enterprise Manager Ops Center deployments.

Table 1–1 (Cont.) Documents for Oracle Enterprise Manager Ops Center

Document	Content
Certified System Matrix Guide	Supported hardware, operating systems, virtualization technologies, databases, and browsers.
Cloud Infrastructure API and CLI Reference Guide	API and CLI commands to manage programmatically the allocated virtual resources for a virtual datacenter account and to create and manage vServers.
System Monitoring Plug-in for Oracle Enterprise Manager Ops Center	Procedure for installing and configuring the plug-in that enables Oracle Enterprise Manager Ops Center to connect to Enterprise Manager Cloud Control.

User Preferences and Role Preferences

The *Oracle Enterprise Manager Ops Center Concepts Guide* describes the features of the user interface. Some of the actions and the abilities can be changed, either by an individual user or by an administrator for all users with a specific role.

To see the current specifications or to change the specifications, click **Setup** in the title bar as shown in Figure 1–2 and then click **My Preferences** to view information about how the your account has been specified.

Figure 1–2 Setting User Preferences

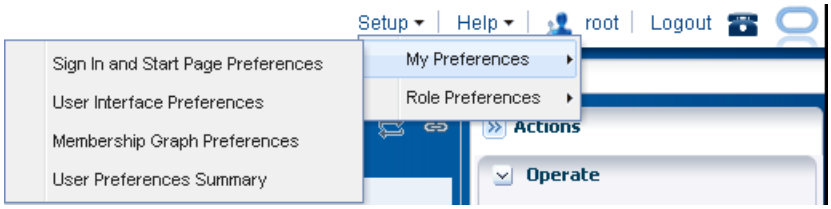

About the Current User Interface Preferences

Figure 1–3 shows the User Preferences Summary window, which displays the current specifications for your start page, time intervals, and each asset type's default tab in the center pane. The current specification for the display of the Membership Graph and the Sign In and Start Page have separate windows.

- Sign In and Start Page Preferences
- Membership Graph Preferences

To change the specifications in the User Preferences Summary window, click **User Interface Preferences**. Make changes, then log out and log in again.

Figure 1–3 User Interface Summary

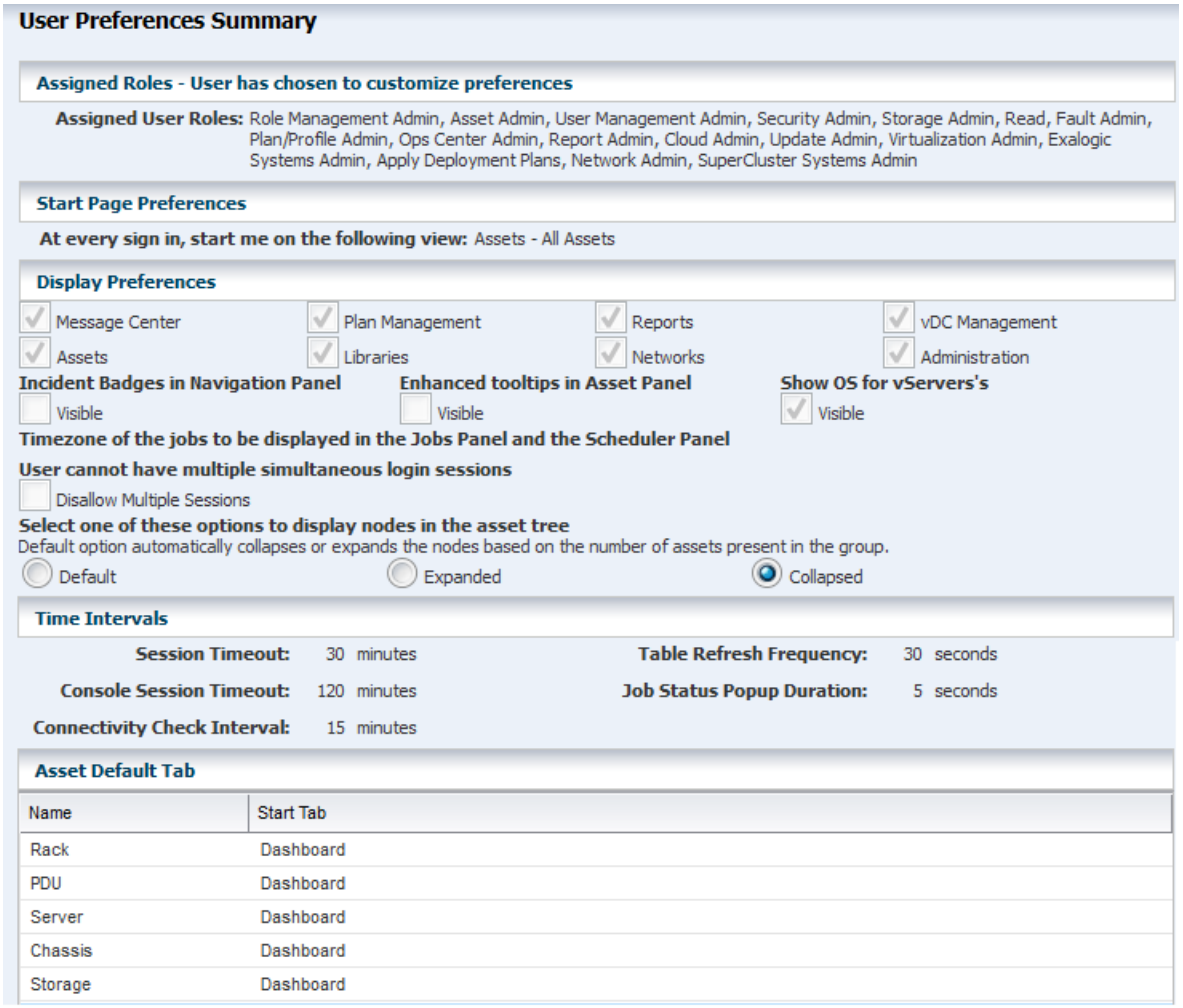

When you change any of the preferences, this window includes a note to indicate that the preferences are not the default specifications.

In the Assigned Roles section, the roles assigned to this user account are listed.

The Start Page Preferences specifies the default view after you log in.

The Display Preferences section provides the following preferences:

- Include or exclude each of the drawers in the Asset pane.
- Show or hide the items in the Asset pane: incident badges, tooltips, and the operating system of a vServer.
- Specify the time zones for the jobs in the Jobs pane.
- Disable simultaneous sessions. The default behavior is to allow a user to log in multiple times. This convenience can be a security risk.
- Change the way assets are expanded or collapsed in the Assets pane. The Default option relies on the number of assets to determine whether the node is expanded or collapsed. You can chose the Expanded option to always show all assets or the Collapsed option to always show only the Asset Type.

The Time Interval section shows the duration of the Session Timeout, the Table Refresh Frequency, the Console Session Timeout, the Job Status Popup Duration, and the Connectivity Check Interval.

The Asset Default Tab section lists each type of asset. For each asset type, you can specify the tab that is the default view displayed in the center pane when an asset of that type is selected.

Preferences By Role

As an administrator, you can set preferences for each role. All user accounts that are assigned the role share the same preferences.

Click **Setup** in the title bar and then click **Role Preferences**. The menu items are the same as for **My Preferences** with the addition of a drop-down list of roles. You select the role and then you select the specifications for all user accounts that have that role.

For example, the default behavior for logging in to the software is to allow the same user to log in multiple times. This is a convenience when monitoring the progress of an operation. However, it can be a security risk so an individual can disable this behavior by checking the **Disable Multiple Sessions** option in the User Interface Preferences window. As an administrator, you can disable this behavior for all users with a certain role by selecting the role and then disabling the option.

Sign In and Start Page Preferences

You can select a section of the Navigation pane as your default view. For example, you can specify that you see Plan Management when you log in. The Assets and Administration sections cannot be hidden. When you select Assets for display, you can choose a default tab to display. Your preferences override the default view for your role and any previous preferences that you.

Membership Graph Preferences

You can change the default orientation of the Membership Graph to left, right, top, or bottom. You can choose the icon size as small or large and the level of depth for the assets to be displayed. Figure 1–4 shows Membership Graph Preferences window.

Figure 1–4 Membership Graph Preferences Window

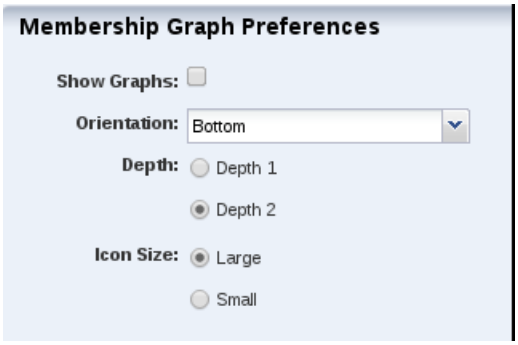

Each time you perform an operation in the Assets pane, the membership graph in the center pane is refreshed. In a datacenter with many assets, you might experience a noticeable delay during the refresh operation. If you are not making changes to the assets or to their relationships, you can hide the membership graph, which eliminates the refresh operation. Clear the **Show Graph** option, as shown in Figure 1–4. You see

the effect of the change after you select an asset.

Time Interval Preferences

Select the **User Interface Preferences** action to set various time intervals to control when the software takes an action or performs an operation.

- Session timeout: Sets the interval to wait for activity before ending your user interface session. The default value is 30 minutes. You can set the time to wait from 5 to 120 minutes.
- Console timeout: Sets the interval to wait for activity from the serial console of managed assets before ending the session. The default value is 120 minutes. You can set the time to wait from 5 to 120 minutes.
- Connectivity check interval: Sets the time to wait before the software checks its access to the Internet, Knowledge Base, and My Oracle Support Services. The default value is 15 minutes and the minimum value is 1 minute.
- Table refresh frequency: Sets the time to wait before refreshing the tables in the user interface. The default value is 30 seconds and the minimum value is 10 seconds.
- Job status popup duration: Sets the time to wait after a job completes to display a status message window. The default value is 5 seconds.

2

Asset Management

Asset management is the process through which you discover your assets and organize them using groups and tags.

The following features and topics are covered in this chapter:

Introduction to Asset Management

Asset management is the process through which Oracle Enterprise Manager Ops Center begins to manage and monitor your assets, which includes server hardware, chassis, racks, network equipment, operating systems, virtualization software and clustering software. Discovering and managing your assets is a prerequisite for almost every action in the software.

Roles for Asset Management

Table 2–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 2–1 Asset Management Tasks and Roles

Task	Role
View Assets	Read
Add Assets	Asset Admin
Find Assets	Asset Admin
Create Discovery Profile	Asset Admin
Update Management Credentials	Security Admin

Table 2–1 (Cont.) Asset Management Tasks and Roles

Task	Role
Edit Asset Attributes	Asset Admin
Edit Access Points	Security Admin
Delete Assets	Asset Admin
Edit Tags	Asset Management
Create Group	Asset Admin SuperCluster Systems Admin
Edit Group	Asset Admin SuperCluster Systems Admin
Move Group	Asset Admin SuperCluster Systems Admin
Add or Remove Assets From a Group	Asset Admin SuperCluster Systems Admin
Delete Group	Asset Admin SuperCluster Systems Admin

Actions for Asset Management

You can perform a variety of asset management actions, depending on the needs of your environment:

- Add Assets by Declaring Servers for OS Provisioning
- Add Assets by Declaring Servers for Service Processor Configuration
- Add Assets Using a Discovery Profile
- Find Assets
- Create a Discovery Profile
- Edit a Discovery Profile
- Copy a Discovery Profile
- Delete a Discovery Profilead
- Install Agent Controllers From the Command Line
- Update Management Credentials
- Edit Asset Attributes
- Delete Assets
- Use Access Points

Location of Asset Management Information in the User Interface

Asset Management actions and information are located in several sections of the user interface.

Table 2–2 Location of Asset Management Information in the BUI

To Display:	Select:
Asset Discovery and Management actions	Expand Assets in the Navigation pane.
Discovery Profiles	Expand Plan Management in the Navigation pane, then select Discovery in the Profiles and Policies section.

Discovering and Managing Assets

Assets that have been discovered and managed by Oracle Enterprise Manager Ops Center can be monitored and targeted with jobs. You can discover assets by using specific discovery criteria, by running a search for ins, or by specifying server information.

After assets have been discovered, they are automatically managed, giving Oracle Enterprise Manager Ops Center full access to the assets and enabling you to monitor, update, and provision them.

Oracle Enterprise Manager Ops Center requires at least one configured Proxy Controller that is accessible on a network before discovering and managing an asset. See Managing Proxy Controllers for Asset Discovery for more information.

Hardware assets are managed using a set of credentials. Operating systems and virtualization software can be managed using an Agent Controller installed on the system or using only a set of credentials. In addition to the default Agent Controllers, Oracle Enterprise Manager Ops Center uses specialized virtualization Agent Controllers called VC Agent Controllers. See Using Agent Management for Operating Systems for more information about Agent Controllers.

Some features are not available when the operating system is managed agentlessly. Table 2–3 shows the information that are available for each management type.

Table 2–3 Information Available for Agent Managed and Agentlessly Managed Assets

Tab or Feature	Agent Managed	Agentlessly Managed
Dashboard	Yes	Yes
Summary	Yes	Yes
Libraries	Yes	No
Storage	Yes	No
Utilization	Yes	Yes
Analytics	Yes	Limited
Virtualization Analytics for Oracle VM Server	Yes, if the guest is agent-managed	No
Virtualization Analytics for Oracle Solaris 10 Zones	Yes, if the global zone is agent-managed or if the non-global zone is agent-managed.	No
Networks	Yes	No
Incidents	Yes	Yes
Monitoring	Yes	Yes
Charts	Yes	Yes
Terminal	Yes	No

Table 2–3 (Cont.) Information Available for Agent Managed and Agentlessly Managed

Tab or Feature	Agent Managed	Agentlessly Managed
Jobs	Yes	Yes
Configuration	Yes	Yes

Managing Proxy Controllers for Asset Discovery

Before you can discover an asset, create at least one proxy controller, which must be associated with a network.

See the following sections in the *Oracle Enterprise Manager Ops Center Installation Guide Oracle Solaris Operating System* for more information about the configuration of new proxy controllers in Oracle Enterprise Manager Ops Center:

- *Installing and Configuring a Proxy Controller Remotely*.
- *Installing and Configuring a Proxy Controller Manually*.

You can find similar information for Linux in the *Oracle Enterprise Manager Ops Center Installation Guide for Linux Operating Systems*.

The *Oracle Enterprise Manager Ops Center Administration Guide* includes information to view and manage proxy controllers:

- To associate, enable, and disable a network in a proxy controller see *Managing Proxy Controller Networks*.
- To migrate an asset to a different proxy controller see *Migrating Assets Between Proxy Controllers*.
- To view the status of a proxy controller see *Viewing Proxy Controllers*.
- To check the status, start, stop, and set the maintenance mode of a proxy controller see *Viewing and Changing the Enterprise Controller and Proxy Controller Status.*

Declaring Servers for OS Provisioning

The Declare Server option lets you declare one or more bare metal systems in preparation for OS provisioning, even if the systems have no service processor.

You can declare a single server by entering the server information directly into the wizard, or declare multiple servers using a discovery file containing the information for all of the servers.

To declare a single server for OS Provisioning, select All Assets in the Navigation pane, then click **Add Assets** in the Actions pane. Select Manually Declare a Server to be a Target of OS provisioning, then select Declare a single server. Enter the server information, then click **Declare Asset**.

- Server Name: name of the server appearing on the UI.
- IP Address: specify an IP address to route the discovery to the correct Proxy Controller. You do not need to use a server's actual IP address. You can use an IP address that is on the same subnet as that of the server to be discovered.
- Model Categories: select the category in which the asset model appears.
- Model: the model of the asset.
- MAC Address and Port combination: used to connect to the server once it is available on the network. Click the Add or Edit icons to add or edit a MAC Address/Port combination, then select the combination.

Enter a logical port name for each network interface. One of these logical port names must be GB_0. Available logical port names are GB_0 through GB_11. You can also use mgmt as a management port. These logical port names will be mapped to network interfaces after the asset has been provisioned, according to the MAC addresses that you specify. If the server has only one network interface, use GB_0.

Enter the MAC addresses of the network interfaces in the server that you want to declare.

To declare multiple servers for OS Provisioning, perform the following steps:

Create a discovery file using the format in the example below.

1. Select **All Assets** in the Navigation pane, then click **Add Assets** in the Actions pane.
2. Select **Manually Declare a Server to be a Target of OS Provisioning**, then select **Declare All Servers**.
3. Enter the location of the discovery file, then click **Declare Asset**.

Example Discovery File

Here is an example of a discovery file.

```
<?xml version='1.0' encoding='utf-8'?>
<servers>
<server name="T5440" model="Sun SPARC Enterprise T5440 Server"
        guid="12345678"
        proxyHostname="server"
        ipAddress="10.0.0.0" >
<ethernetPort name="GB_0" mac="01:23:45:67:89:AB"/>
</server>
</servers>
```

The following variables can be used:

- **server name**: the name that the server has in the UI
- **model**: must be a model supported by Oracle Enterprise Manager Ops Center
- **guid**: a unique identifier for the server
- **proxyHostname**: the Proxy Controller to be used to connect to the server
- **ipAddress**: the IP address for the server
- **netmask**: (optional) the netmask for the server
- **gateway**: (optional) the gateway for the server
- **unconfigured**: (optional) specifies that the server is unconfigured when set to true
- **ethernetPort**: used to connect to the server once it is available on the network. You can specify multiple Ethernet ports. Includes:
  - **Ethernet port name**: the name of the Ethernet port
  - **Ethernet Port mac**: the MAC address for the Ethernet port

Declaring Servers for Service Processor Configuration

The Declare Servers for Service Processor Configuration option lets you declare one or more bare metal systems in preparation for service processor configuration.

The assets being declared do not need to be physically connected to the network at the time of the discovery, because the assets produced by an asset declaration are skeletal representations of the real assets. These assets can then be targeted with service processor configuration jobs. Once the real assets are connected to the network, provisioned and discovered, they are correlated with the declared version into complete assets.

To declare unconfigured assets for service processor configuration, perform the following steps:

1. Select **All Assets** in the Navigation pane, then click **Add Assets** in the Actions pane.
2. Select **Declare an Unconfigured Hardware Asset**, then enter data for the servers to be declared:

Figure 2–1 Declare Unconfigured Hardware Assets

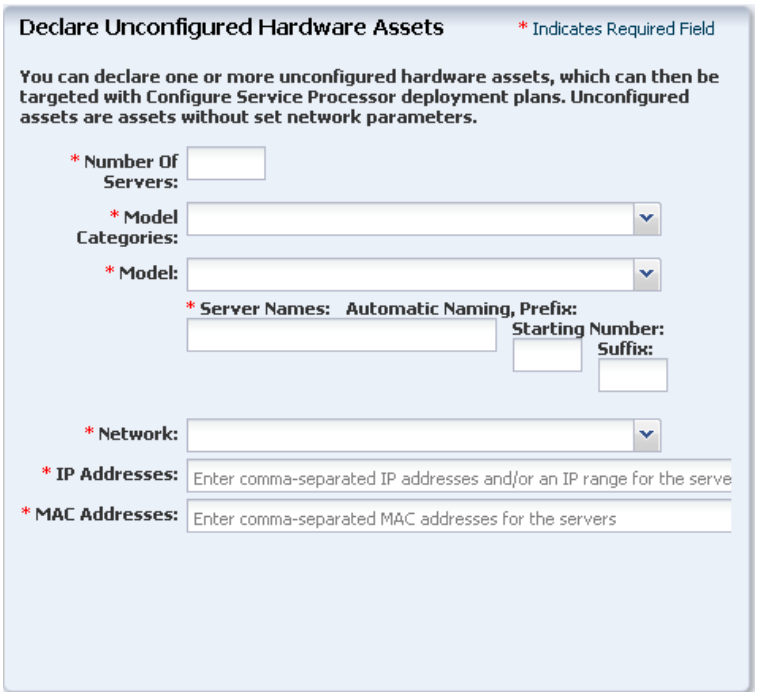

- Number of Servers: the total number of servers to be discovered
- Model Categories: the model category of the servers
- Model: the specific model of the servers
- Server Names: the names of the servers, including:
  - Prefix: a prefix that appears before each server name. This field is required.
  - Starting Number: the number of the first server. The number is increased by one for each additional server. This field is required.
  - Suffix: a suffix that appears after each server name
- Network: the network on which the server or servers is added
- IP addresses: the IP addresses to be used for the servers
- MAC Addresses: the MAC Addresses of the servers

3. Click **Declare Asset**.

Adding Assets Using a Discovery Profile

After you have created a discovery profile, you can run it to discover and manage assets. This lets you discover assets using a pre-existing set of protocols and credentials, and manage assets consistently.

To declare unconfigured assets for service processor configuration, perform the following steps:

1. In the Navigation pane, select **All Assets**.
2. In the Actions pane, click **Add Assets**.
3. Select **Add and Manage Various Types of Assets via Discovery Probes**, then select a discovery profile.
4. Add or edit the IP addresses, host names, and credentials for the targets, then click **Add Now**.

Note: Any host names must be resolvable from the Enterprise Controller to be discovered.

Finding Assets

Service tags are small XML files containing product information. Many Oracle systems come equipped with service tags. If you have hardware assets equipped with service tags, you can discover them using the Find Assets Wizard. This method lets you discover large numbers of assets quickly.

The Find Assets Wizard searches known networks for service tags, then uses credentials that you specify to manage the discovered assets.

Products without service tags cannot be discovered using this method. For example, ALOM systems do not have service tags.

Starting with Oracle Enterprise Manager Ops Center 12c Release 2 (12.2.2.0.0), the Find Assets action is disabled by default. To enable the action, see Service Tag Discovery.

To find assets, perform the following steps:

1. In the Navigation pane, select **All Assets**.
2. In the Actions pane, click **Find Assets**.
3. Select **Run Discovery Now**.
4. When the initial discovery is complete, select the assets in each category (hardware, operating systems, and engineered systems) that you want to manage and provide credentials for them, then click **Finish**.

Service Tag Discovery

Starting with Oracle Enterprise Manager Ops Center 12c Release 2 (12.2.2.0.0), the Find Assets action is disabled by default. Discovery might fail in big data centers where there are more numbers of assets to scan. If you want to discover assets using service tags, you can edit the system property file to enable the Find Assets action.

To enable the system property file, perform the following steps:

1. In the Navigation pane, click **Administration**, then click **Enterprise Controller**.

2. In the center pane, select the **Configuration** tab.

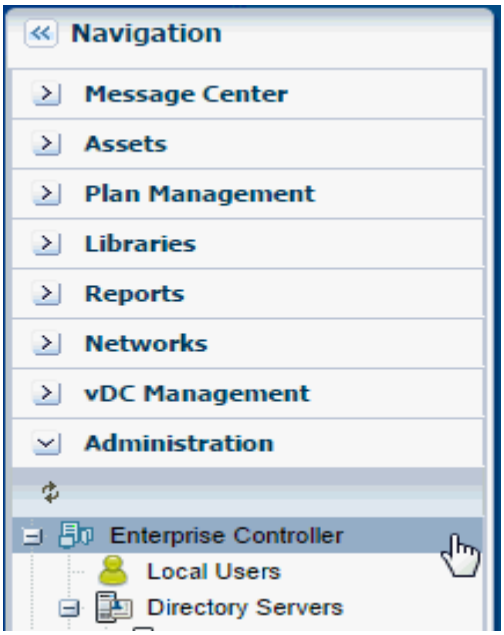

3. In the Configuration Management section, select **Discovery** from the Subsystem drop-down list.

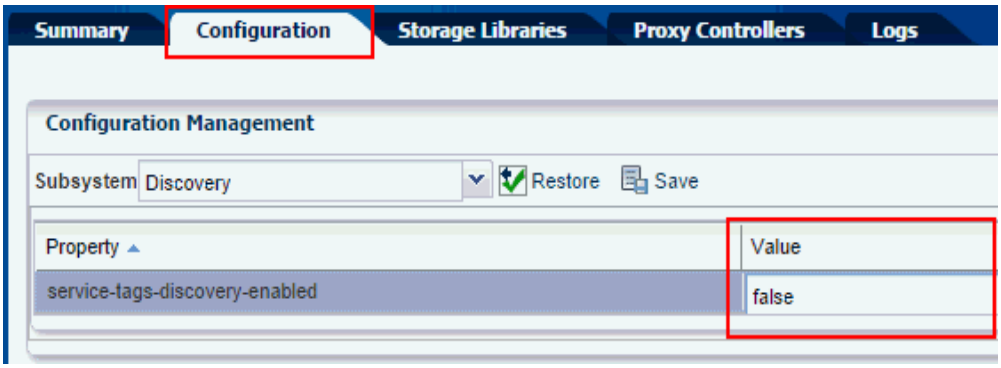

4. Set the value of the **service-tags-discovery-enabled** property file to **true**.

Creating a Discovery Profile

You can create a discovery profile and then discover the assets that comply with the profile.

To create a discovery profile, perform the following steps:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
3. In the Actions pane, click **Create Profile**.

Figure 2–2 Create Discovery Profile

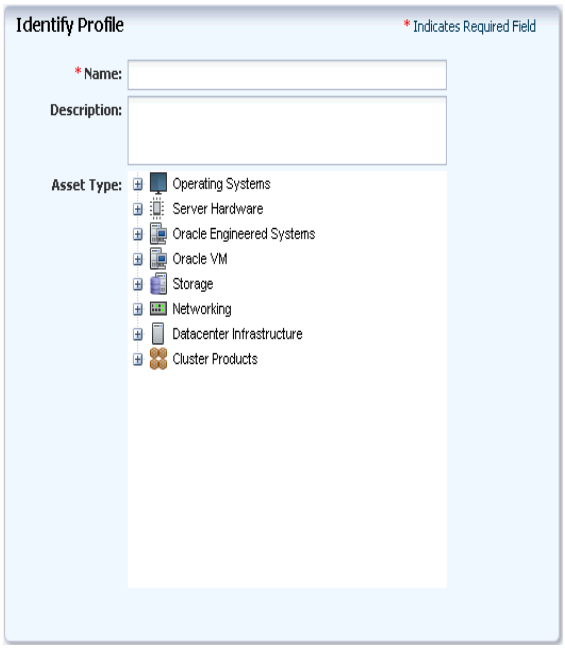

4. Enter a name for the discovery profile and select a type of asset to discover. You can also add a description. Click **Next**.
5. (Optional.) Click the add icon to add one or more tags to discovered assets, then click **Next**.
6. (Optional.) Click the add icon to add one or more IP ranges. If you do not add IP ranges when creating a discovery profile, you can do so when you run the profile. Enter the following information for each IP range:
   - Name: a name for the IP range
   - Description: a description of the IP range
   - Network: identify the managed network associated with the host names or IP addresses to route the discovery to the correct Proxy Controller. Select Automatic to route the job to the default proxy controller. The IP address of a target must resolve to only one known network for automatic routing to succeed. Select Networks to change the association among networks and proxy controllers.
   - IP ranges: an IP range or comma-separated list of IP addresses. If you want to target specific host names, you must enter them when you run the discovery profile.
7. Provide the discovery credentials information, then click **Next**.
   - Click **New** to add new discovery credentials for each protocol, or click Select to select an existing set of credentials.
   - Select **Deploy Agent Controller to deploy an Agent Controller on discovered operating systems**, or select **Manage Without Agent Controller to manage discovered operating systems without installing an Agent Controller**.
   - If the service tag parameters have been modified for the target assets, enter the Service Tag parameters:

  - Service tag passphrase: necessary if a service tag has been configured to be encrypted.
  - Service tag port: necessary if a service tag has been configured to use a port other than the default of 6481.
  - Service tag timeout: the default value is 20 seconds.

8. Review the summary and click **Finish** to create the discovery profile.

Editing a Discovery Profile

To edit a discovery profile, perform the following steps:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
3. Select the profile and click the Edit Profile icon.
4. Edit any of the information in the discovery profile, then click **Finish** to save your changes.

Copying a Discovery Profile

To copy a discovery profile, perform the following steps:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
3. Select the profile and click the Copy Profile icon.
4. Edit any of the information in the discovery profile, then click **Finish** to save the new discovery profile.

Deleting a Discovery Profile

To delete a discovery profile, perform the following steps:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
3. Select the profile and click the Delete Profile icon. The discovery profile is deleted.

Installing Agent Controllers From the Command Line

Use these procedures to install an Agent Controller and to register the target system. See Using Agent Management for Operating Systems and Virtualization Agent Controllers for more information about Agent Controllers.

Before You Begin

To use the `agentadm` command, you need the following information:

- To configure your Agent Controller software using an administrative user account on the Enterprise Controller you need:
  - User name: the user account provides authentication that supports Agent Controller registration. Use the user name of this account as the argument for the `-u` option of the agentadm command.

  - Password: use this password to populate the `/var/tmp/OC/mypasswd` file. Then use this file name as the argument for the `-p` option of the agentadm command.

- The auto-reg-token registration token from the `/var/opt/sun/xvm/persistence/scn-proxy/connection.properties` file on the appropriate Proxy Controller – If you decide not to use user credentials to configure your Agent Controller software, use this token to populate the `/var/tmp/OC/mytoken` file. Then use this file name as the argument for the `agentadm -t` option.
- IP address or host name of the Proxy Controller with which you will associate the Agent Controller – Use this IP address or host name as the argument for the `agentadm -x` option. Typically, you would associate the Agent Controller with the Proxy Controller that is connected to the same subnet as the target system.
- The IP address of the network interface that the Agent Controller will use for registration – Use this IP address as the argument for the `agentadm -a` option.

Some example `agentadm` commands in this procedure use the alternative administrative user name droot. In these examples, the droot user exists on the Enterprise Controller.

When you install an Agent Controller on a global zone, the Agent Controller installation installs, or upgrades to Java Runtime Environment (JRE) 1.6.0_51. Later versions of JRE are not affected.

Requirements for Installing an Agent Controller on Oracle Solaris 11

If you are installing an Agent Controller on a server with Oracle Solaris 11 check the following requirements:

- The Oracle Solaris 11 Update Software Library must be configured and its initial synchronization completed.
- The version of Oracle Solaris 11 installed on the target system must be available in the Oracle Solaris 11 Update Software Library.
- The Agent Controller packages must be in the Oracle Solaris 11 Update Software Library.

To Manually Install and Configure an Agent Controller Using User Credentials

This procedure creates a file that holds the password of the administrative user for your Oracle Enterprise Manager Ops Center installation.

1. On the Enterprise Controller, change to the `/var/opt/sun/xvm/images/agent/` directory, and list the files that it contains. This directory contains the Agent Controller installation archives. For example:

```
# cd /var/opt/sun/xvm/images/agent/
# ls
OpsCenterAgent.Linux.i686.12.2.0.2503.zip
OpsCenterAgent.Linux.i686.12.2.0.2503.zip.sig
OpsCenterAgent.Solaris.i386.12.2.0.2503.zip
OpsCenterAgent.Solaris.i386.12.2.0.2503.zip.sig
OpsCenterAgent.Solaris.sparc.12.2.0.2503.zip
OpsCenterAgent.Solaris.sparc.12.2.0.2503.zip.sig
OpsCenterAgent.SolarisIPS.all.12.2.0.2503.zip
OpsCenterAgent.SolarisIPS.all.12.2.0.2503.zip.sig
#
```

2. Identify the Agent Controller archive that is appropriate for the system where you intend to install the Agent Controller. See Table 2–4 for a description of the available packages.

Table 2–4 Agent Controller Packages and their related Target Operating System and Architecture

File prefix	Operating System / Architecture
OpsCenterAgent.Linux.i686	Oracle Linux/x86
OpsCenterAgent.Solaris.i386	Oracle Solaris 10/x86
OpsCenterAgent.Solaris.sparc	Oracle Solaris 10 / Oracle SPARC
OpsCenterAgent.SolarisIPS.all	Oracle Solaris 11 / x86 and Oracle SPARC

3. On the system where you want to install the Agent Controller (the target system), create a directory named `/var/tmp/OC`.

```
# mkdir /var/tmp/OC
```

4. Use scp or ftp to transfer the correct Agent Controller archive from the Enterprise Controller to the `/var/tmp/OC` directory on the target system. Respond to any authentication or confirmation prompts that are displayed. For example:

```
# scp OpsCenterAgent.Solaris.sparc.12.2.0.2503.zip root@10.0.0.0:/var/tmp/OC
Password:
OpsCenterAgent.S 100%
|*****************************************************************| 187078
KB 00:32
#
```

5. On the target system, change to the `/var/tmp/OC` directory.

```
# cd /var/tmp/OC
#
```

6. Use the `unzip` command to uncompress the Agent Controller archive. For example:

```
# unzip OpsCenterAgent.Solaris.sparc.12.2.0.2503.zip
(output omitted)
```

7. Run the `install -a` script in the `OpsCenterAgent` directory. For example:

```
# OpsCenterAgent/install -a
Installing Ops Center Agent Controller.
No need to install 120900-04.
No need to install 121133-02.
No need to install 119254-63.
No need to install 119042-09.
No need to install 121901-02.
No need to install 137321-01.
Installed SUNWjdmk-runtime.
Installed SUNWjdmk-runtime-jmx.
(output omitted)
6 patches skipped.
19 packages installed.
Installation complete.
Detailed installation log is at /var/scn/install/log.
Uninstall using /var/scn/install/uninstall.
#
```

If you are installing the Agent Controller on Oracle Solaris 11, run the `install` command with the `-p` option and the IP address. The local IPS pubs are configured for access to the Oracle Solaris 11 Update Software Library using the IP address. For example:

```
# OpsCenterAgent/install -p 10.0.0.1
#
```

If you are installing an Oracle VM Server Virtualization Controller Agent use the `-l` (or `--ldom`) option.

8. Create an empty file named `/var/tmp/OC/mypasswd`, and set its permission mode to 400. For example:

```
# touch /var/tmp/OC/mypasswd
# chmod 400 /var/tmp/OC/mypasswd
```

9. Edit the `/var/tmp/OC/mypasswd` file so that it contains the password for the administrative user that exists on the Enterprise Controller to which the Proxy Controller is connected. The following echo command appends the password to the `/var/tmp/OC/mypasswd` file. Replace password with the correct password. For example:

```
# echo 'password' > /var/tmp/OC/mypasswd
```

10. Use the `agentadm` command to associate the Agent Controller with the Proxy Controller.
    - Oracle Solaris OS: use the `/opt/SUNWxvmoc/bin/agentadm` command
    - Linux OS: use the `/opt/sun/xvmoc/bin/agentadm` command. The example commands below use the following options:
    - configure: causes an Agent Controller configuration operation to take place.
    - -u: specifies the administrative user that exists on the Enterprise Controller to which the Proxy Controller is connected. Be certain that the password that you specified in the `/var/tmp/OC/mypasswd` file is correct for the user that you specify for this option.

      Note: The example below uses ***droot*** as the administrative user.

    - -p: specifies the absolute path name of the file that contains the password for the user that you specified with the `-u` option.
    - -x: specifies the IP address or host name of the Proxy Controller to which this Agent Controller will connect.
    - -a: specifies the IP address to use during Agent Controller registration. This selects the network interface that the Agent Controller will use for registration. Accept the server's certificate when prompted. For example:

```
# /opt/SUNWxvmoc/bin/agentadm configure -u droot -p /var/tmp/OC/mypasswd -x
10.0.0.0
agentadm: Version 1.0.3 launched with args: configure -u droot -p
/var/tmp/OC/mypasswd -x 10.0.0.1
workaround configuration done.
Certificate:
Serial Number: 947973225
Version: 3
```

```
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_Agent Controller
Not valid before: Thu Jun 19 15:36:59 MDT 1969
Not valid after: Thu Apr 19 15:36:59 MDT 2029
Certificate:
Serial Number: 1176469424
Version: 3
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_ca
Not valid before: Thu Jun 19 15:36:56 MDT 1969
Not valid after: Thu Apr 19 15:36:56 MDT 2029
Accept server's certificate? (y|n)
y
Connection registered successfully.
scn-Agent Controller configuration done.
Checking if UCE Agent Controller process is still running, it may take a
couple of minutes ...
Process is no longer running
UCE Agent Controller is stopped.
UCE Agent Controller is in [online] state.
Checking if UCE Agent Controller process is up and running ...
The process is up and running.
UCE Agent Controller is started.
Added the zone configuration automation successfully.
Added the service tags recreate script successfully.
#
```

Error messages similar to *Connection cannot be registered* in the following example typically indicate problems with the user credentials that you specified in the `agentadm` command. In this example, the user `droot` was not authenticated on the Enterprise Controller. If you see this error, check that the user name that you supplied for the `agentadm -u` option, and the password in the file that you specified for the `agentadm -p` option, match an existing administrative user on the Enterprise Controller.

```
Accept server's certificate? (y|n)
y
Error with connection to CRS: com.sun.scn.connmgt.SCNRegClientException:
droot, Code: 4, Code: 4
ERROR : Connection cannot be registered.
Code--2
sc-console registration failed on [2].
sc-console : User authentication error.
Error executing step : sc_console
```

If the system where you are installing the Agent Controller has multiple active network interfaces, you can use the `-a` option to specify the IP address of the interface that you want to use for Agent Controller registration. For example:

```
# /opt/SUNWxvmoc/bin/agentadm configure -u droot -p /var/tmp/OC/mypasswd -x
10.0.0.0 -a 10.0.0.1
(output omitted)
```

11. If you encountered a *Connection cannot be registered* error message from the `agentadm` command, use `agentadm` to unconfigure the Agent Controller. For example:

```
# /opt/SUNWxvmoc/bin/agentadm unconfigure
agentadm: Version 1.0.3 launched with args: unconfigure
verified sc_console command is OK
```

```
End of validation
{output omitted}
End of configuration.
```

After the Agent Controller has been unconfigured, correct the problem that was indicated by the error message, and re-run the `agentadm configure` command.

12. Use the `sc-console` command to list the Agent Controller connection. For example:

```
# sc-console list-connections
scn-Agent Controller https://10.0.0.0:21165
urn:scn:clregid:abcdef12-6899-4bcc-9ac7-a6ebaf71c1f5:20090420171121805
#
```

To Manually Install and Configure an Agent Controller Using a Token

This procedure uses a token to configure your Agent Controller software.

1. On the Enterprise Controller, change to the `/var/opt/sun/xvm/images/agent/` directory, and list the files that it contains. This directory contains the Agent Controller installation archives. For example:

```
# cd /var/opt/sun/xvm/images/agent/
# ls
OpsCenterAgent.Linux.i686.12.1.0.zip
OpsCenterAgent.Linux.i686.12.1.0.zip.sig
OpsCenterAgent.SunOS.i386.12.1.0.zip
OpsCenterAgent.SunOS.i386.12.1.0.zip.sig
OpsCenterAgent.SunOS.sparc.12.1.0.zip
OpsCenterAgent.SunOS.sparc.12.1.0.zip.sig
#
```

2. Identify the Agent Controller archive that is appropriate for the system where you intend to install the Agent Controller. See Table 2–4 for a description of the available packages.

3. On the system where you want to install the Agent Controller (the target system), create a directory named `/var/tmp/OC`.

```
# mkdir /var/tmp/OC
```

4. Use scp or ftp to transfer the correct Agent Controller archive from the Enterprise Controller to the `/var/tmp/OC` directory on the target system. Respond to any authentication or confirmation prompts that are displayed. For example:

```
# scp OpsCenterAgent.SunOS.sparc.12.1.0.zip root@10.5.241.74:/var/tmp/OC
Password:
OpsCenterAgent.S 100%
|*****************************************************************| 34695
KB 00:32
#
```

5. On the target system, change to the `/var/tmp/OC` directory.

```
# cd /var/tmp/OC
#
```

6. Use the `unzip` command to uncompress the Agent Controller archive. For example:

```
# unzip OpsCenterAgent.SunOS.sparc.12.1.0.zip
(output omitted)
```

7. Run the `install -a` script in the `OpsCenterAgent` directory. For example:

```
# OpsCenterAgent/install -a
Installing Ops Center Agent Controller.
No need to install 120900-04.
No need to install 121133-02.
No need to install 119254-63.
No need to install 119042-09.
No need to install 121901-02.
No need to install 137321-01.
Installed SUNWjdmk-runtime.
Installed SUNWjdmk-runtime-jmx.
(output omitted)
6 patches skipped.
19 packages installed.
Installation complete.
Detailed installation log is at /var/scn/install/log.
Uninstall using /var/scn/install/uninstall.
#
```

   If you are installing the Agent Controller on Oracle Solaris 11, run the `install` command with the `-p` option and the IP address. The local IPS pubs are configured for access to the Oracle Solaris 11 Update Software Library using the IP address. For example:

```
# OpsCenterAgent/install -p 10.0.0.1
#
```

8. On the Proxy Controller that will communicate with this Agent Controller instance, examine the `/var/opt/sun/xvm/persistence/scn-proxy/connection.properties` file. The last line in this file contains the auto-reg-token that is required for Agent Controller registration. For example:

```
# cat /var/opt/sun/xvm/persistence/scn-proxy/connection.properties
#Generated by a program. Do not edit. All manual changes subject to deletion.

(output omitted)

trust-store=/var/opt/sun/xvm/security/jsse/scn-proxy/truststore
auto-reg-token=abcdef12-1700-450d-b038-ece0f9482474\:1271743200000\:T
#
```

9. On the system where you have installed the Agent Controller software, create an empty file named `/var/tmp/OC/mytoken`, and set its permission mode to 400. For example:

```
# touch /var/tmp/OC/mytoken
# chmod 400 /var/tmp/OC/mytoken
```

10. Edit the `/var/tmp/OC/mytoken` file so that it contains the auto-reg-token string from Proxy Controller with the following changes:

    - Remove the `auto-reg-token=`.
    - Remove any backslash characters from the token string. For example:

      ```
      abcdef12-1700-450d-b038-ece0f9482474:1271743200000:T
      ```

11. Use the `agentadm` command to associate the Agent Controller with a Proxy Controller.

- Oracle Solaris OS: use the `/opt/SUNWxvmoc/bin/agentadm` command
- Linux OS: use the `/opt/sun/xvmoc/bin/agentadm` command. The example commands below use the following options:
- configure: causes an Agent Controller configuration operation to take place.
- -t: specifies the absolute path name of the file that contains the registration token.
- -x: specifies the IP address or host name of the Proxy Controller to which this Agent Controller will connect.
- -a: specifies the IP address to use during Agent Controller registration. This selects the network interface that the Agent Controller will use for registration. Accept the server's certificate when prompted. For example:

```
# /opt/SUNWxvmoc/bin/agentadm configure -t /var/tmp/OC/mytoken -x 10.0.0.0
agentadm: Version 1.0.3 launched with args: configure -t
/var/tmp/OC/mytoken -x 10.0.0.0
workaround configuration done.

Certificate:
Serial Number: 947973225
Version: 3
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_Agent Controller
Not valid before: Thu Jun 19 15:36:59 MDT 1969
Not valid after: Thu Apr 19 15:36:59 MDT 2029

Certificate:
Serial Number: 1176469424
Version: 3
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_ca
Not valid before: Thu Jun 19 15:36:56 MDT 1969
Not valid after: Thu Apr 19 15:36:56 MDT 2029

Accept server's certificate? (y|n)
y
Connection registered successfully.
scn-Agent Controller configuration done.
Checking if UCE Agent Controller process is still running, it may take a
couple of minutes ...
Process is no longer running
UCE Agent Controller is stopped.
UCE Agent Controller is in [online] state.
Checking if UCE Agent Controller process is up and running ...
The process is up and running.
UCE Agent Controller is started.
Added the zone configuration automation successfully.
Added the service tags recreate script successfully.
#
```

If the system where you are installing the Agent Controller has multiple active network interfaces, you can use the `-a` option to specify the IP address of the interface that you want to use for Agent Controller registration. For example:

```
# /opt/SUNWxvmoc/bin/agentadm configure -t /var/tmp/OC/mytoken -x 10.0.0.0
-a 10.0.0.1
(output omitted)
```

12. If you encountered a *Connection cannot be registered* error message from the `agentadm` command, use `agentadm` to unconfigure the Agent Controller. For example:

```
# /opt/SUNWxvmoc/bin/agentadm unconfigure
agentadm: Version 1.0.3 launched with args: unconfigure
verified sc_console command is OK
End of validation

{output omitted}
End of configuration.
```

    After the Agent Controller has been unconfigured, correct the problem that was indicated by the error message, and re-run the `agentadm configure` command.

13. Use the `sc-console` command to list the Agent Controller connection. For example:

```
# sc-console list-connections
scn-Agent Controller https://10.0.0.0:21165
urn:scn:clregid:abcdef12-6899-4bcc-9ac7-a6ebaf71c1f5:20090420171121805
#
```

Using Management Credentials

Oracle Enterprise Manager Ops Center stores the credentials that are used to manage each asset. You can update, edit, and delete these credentials.

Upgrading Management Credentials From a Prior Version

Assets that were discovered and managed in prior versions of Oracle Enterprise Manager Ops Center might not have management credentials associated with them. You can associate new or existing sets of credentials with these assets.

If a discovered asset is blacklisted, the same can be removed by updating the management credentials.

To upgrade management credentials, perform the following steps:

1. On the Navigation pane, select **All Assets**.
2. In the Actions pane, click **Upgrade Management Credentials**.
3. Select an asset category: operating systems; servers; or chassis, m-series, and switches.
4. Select one or more assets of that category.
    - To assign an existing set of credentials, select **Assign existing set** and then select an existing set of credentials.
    - To assign a new set of credentials, select **Create and assign new set** and then enter a protocol, name, and credential information.

Updating Management Credentials

You can change the set of management credentials used by an asset or group of assets.

To update management credentials, perform the following steps:

1. On the Navigation pane, select an asset or group.
2. In the Actions pane, click **Update Management Credentials**.
3. Click **Select** to select an existing set of credentials, or click **New** to create a new set.

Creating Management Credentials

You can create a new set of management credentials. These credentials can then be used to discover and manage new assets or to manage existing assets.

To create management credentials, perform the following steps:

1. On the Navigation pane, under Administration, select **Credentials**.
2. In the Actions pane, click **Create Credentials**.
3. Select a protocol, then enter a name for the set of credentials and the information required by the protocol.
4. Click **Create** to create the management credentials.

Creating Credentials for Access to the Serial Console or SSH Tunnel

To enable a connection to a service processor or virtual machine, define the user account that Oracle Enterprise Manager Ops Center uses to open an SSH tunnel on the Enterprise Controller or to create a serial connection.

> **Note:** If you do not specify this account, Oracle Enterprise Manager Ops Center creates an account each time it accesses a serial console and deletes the account when the connection is no longer needed. This activity might not conform to your site's security policy.

The following types of assets use SSH to connect to a serial console. Create an account for each type and define the same password for each account.

- Proxy Controllers
- Global zones that use agents and require access to the consoles of non-global zones
- Control domains that use agents and require access to the consoles of logical domains

To create the account, define the `ConsoleSSHCredname` system property using the procedure in Defining the system property for console access and then define a user account for that property using either the procedure in Creating the account using Oracle Enterprise Manager Ops Center or the procedure in Creating the account using the `useradd` command.

Defining the system property for console access

1. Select the **Administration** section in the Navigation pane.
2. Select the **Configuration** tab in the center pane.
3. In the Subsystem list, select **Console Access Configuration**. The `ConsoleSSH.Credname` system property is displayed.
4. Click in the **Values** column.
5. Enter the name of the new user account. For example, `SERIALCONSOLE_CRED1`.

Figure 2–3 Configuring Console Access

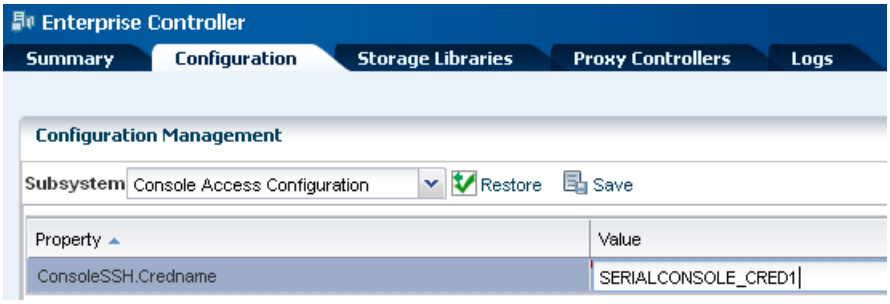

6. Click **Save**.

When the job is completed, define the account using the following procedure.

Creating the account using Oracle Enterprise Manager Ops Center

You must have the Security Admin role to perform this procedure.

After you define the user account, the account is created automatically in `/etc/passwd` the first time a job for console access is run. However, if your site's security policy requires that the operating system account must be created outside of Oracle Enterprise Manager Ops Center's control or if you prefer to create the account manually, use the procedure described in Creating the account using the `useradd` command.

1. Select the **Administration** in the Navigation pane.
2. Select **Credentials** in the Navigation pane.
3. Click **Create Credentials** in the Actions pane.
4. Select the **SERIAL_CONSOLE_SSH** protocol and enter the following details:
   - Name of the credential: Enter the value of the `ConsoleSSH.Credname` system property. In this example, `SERIALCONSOLE_CRED1`.
   - Login User: Enter a convenient or descriptive name for the user account, for example, `ConsoleAccess`.
   - Password for the user account and its confirmation.

Figure 2–4 User Account for Console Access

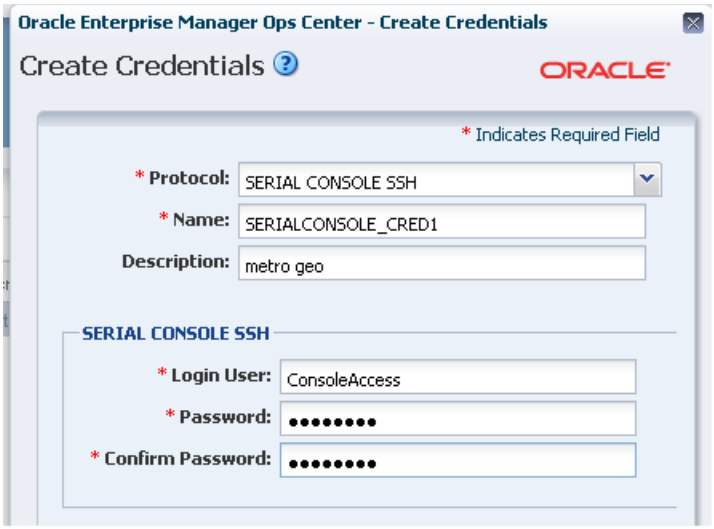

5. Click **Create** to submit the job.

Creating the account using the useradd command

1. Create the home directory for the account. In the following example, the account is named `consolex`:

   ```
   mkdir /var/tmp/consolex
   ```

2. Add the user account with its shell, `/opt/sun/n1gc/bin/serial_console`:

   ```
   useradd -s "/opt/sun/n1gc/bin/serial_console" -d /var/tmp/consolex -u uid -P
   "profile" -A "solaris.zone.manage" consolex
   ```

where *uid* is an available user ID on the Enterprise Controller's system and *profile* is either `LDoms Review` for a control domain or `Zone Management` for a global zone. The `-A` option is a feature of Oracle Solaris 11's `useradd(1m)` command that includes an authorization defined in `auth_attr(4)`.

3. Change the ownership of the home directory:

   ```
   /bin/chown consolex /var/tmp/consolex
   /bin/chmod 700 /var/tmp/consolex
   ```

4. Set and confirm the password for the account:

   ```
   passwd consolex
   ```

Using Custom SSH Keys for OS Discovery

You can use SSH keys as an option when performing the OS discovery.

You must create SSH keys on each Proxy Controller that might access the target asset and also add the SSH public key to the `~/.ssh/authorized_keys` location on the remote OS or use the hardware's web interface to upload the public keys.

1. On the Navigation pane, under Administration, click **Credentials**.
2. In the Actions pane, click **Create Credentials**.

Figure 2–5 Create Credentials

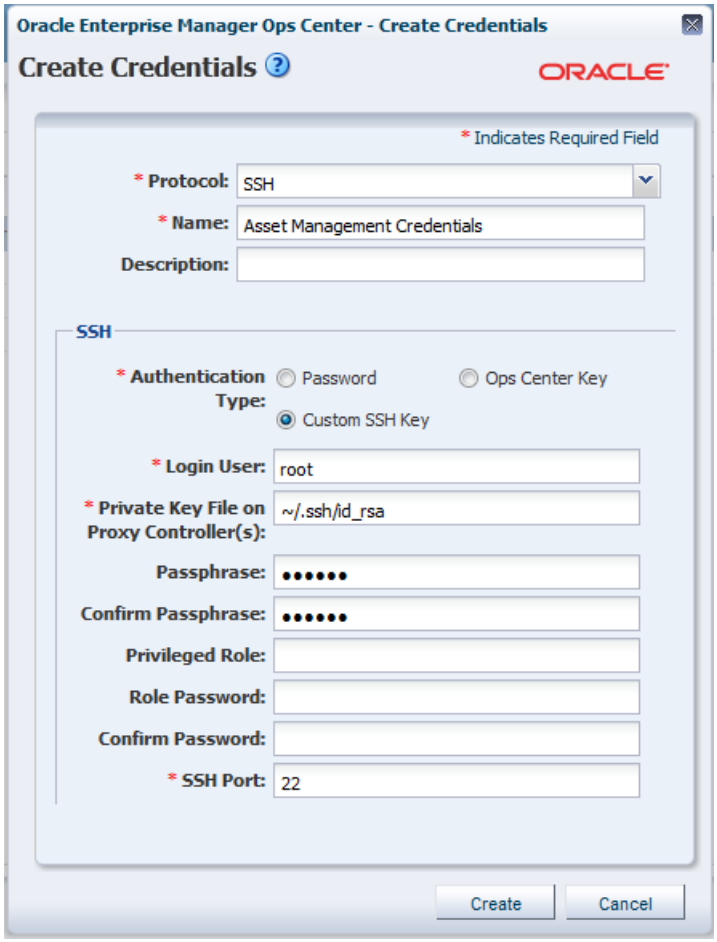

3. In the Name field, enter a name for the credential.
4. In the Description field, enter a description for the credential.
5. In the SSH section, select Custom SSH Key as Authentication Type.
6. In the Login User field, enter the login credential.
7. In the Private Key File on Proxy Controller(s) field, the location is set by default; change this value if you want to refer to other keys.
8. In the Passphrase field, enter the passphrase, if one was specified, that was set when the key was created.
9. In the Confirm Passphrase field, enter the passphrase again.
10. In the SSH Port field, enter the SSH port number. The default is 22.
11. Click **Create**.

For a nonprivileged user, enter the Privileged User role credentials in the Privileged Role and Role Password fields respectively.

Editing Management Credentials

You can edit an existing set of management credentials to reflect changes to the managed assets.

To edit management credentials, perform the following steps:

1. On the Navigation pane, under Administration, select **Credentials**.
2. In the center pane, select a set of credentials and click the Edit Credentials icon.
3. Edit the description and the information required by the protocol, then click **Update** to save the changes.

Copying Management Credentials

You can copy an existing set of management credentials to create a new set.

To copy management credentials, perform the following steps:

1. On the Navigation pane, under Administration, select **Credentials**.
2. In the center pane, select a set of credentials and click the Copy Credentials icon.
3. Edit the name, description, and the information required by the protocol, then click **Copy** to save the new set of credentials.

Deleting Management Credentials

You can delete an existing set of management credentials. Discovery profiles that use the credentials might no longer function, and Agentless assets that are managed using the credentials must be given a new set.

To delete management credentials, perform the following steps:

1. On the Navigation pane, under Administration, select **Credentials**.
2. In the center pane, select a set of credentials and click the Delete Credentials icon.
3. Click **OK** to delete the credentials.

Editing Asset Attributes

All assets have a description and a set of tags that can be edited.

The description field can be used for descriptive information about a system.

You can use tags categorize assets and simplify later searches. Each tag consists of a tag name and a value.

To edit asset attributes, perform the following steps:

1. On the Navigation pane, under Assets, select an asset.
2. In the Actions pane, click **Edit Asset**.
3. Edit the name, description, and tags, then click **Save**.

Using Access Points

An asset's access points show how Oracle Enterprise Manager Ops Center connects to the asset.

The following are possible access points:

- The discovery credentials used to discover the asset.
- The discovery credentials used to discover a related asset. For example, an access point for a service processor is the discovery credentials of its operating system.
- Agent Controller installed on the asset.
- A virtual asset's virtualization host.

To view access points, select an asset and click the Configuration tab.

To delete an access point, perform the following steps:

1. On the Navigation pane, under Assets, select an asset and click the Configuration tab.
2. Select one or more access points, then click **Delete Access Point**.

Deleting Assets

The Delete Asset option uninstalls Agent Controller software if it is present, removing the asset from Oracle Enterprise Manager Ops Center. All data for the asset is removed. You delete assets to stop managing them with Oracle Enterprise Manager Ops Center.

The operating systems that support the Enterprise Controller and Proxy Controllers cannot be deleted or removed.

Note: Deleting a global zone also deletes its non-global zones. However, if a global zone is managed with an Agent Controller, its non-global zones continue to be managed agentlessly even if their Agent Controllers are removed.

To Delete Assets

1. Click **All Assets** in the Assets section of the Navigation pane.
2. Select the asset or assets that you want to delete from the Managed Assets or Unprocessed Assets tabs.
3. Click **Delete Asset**.
4. If the assets do not have Agent Controllers, a confirmation window is displayed. Click **Delete**. A job is launched to delete the assets.
5. If the assets have Agent Controllers, the Management Credentials page is displayed. Provide management credentials using one of the following methods:
   - Click **New** to create a new set of credentials. Enter a name and the credential information, then click **OK**.
   - Click **Select** to select an existing set of credentials. Select an existing set of credentials, then click **OK**.
6. Click **Next**.

   The Summary page is displayed.
7. Review the summary, then click **Finish**.

Special Discovery and Management Procedures

Most assets can be discovered and managed using the procedures in Discovering and Managing Assets. However, some types of assets must be discovered or managed using special procedures. Use these procedures when any of the following asset types is discovered or managed.

- Oracle Engineered Systems: you can also discover an engineered system using a discovery profile. See Oracle Engineered Systems Management in the Oracle Engineered Systems chapter for more information.

- Oracle Solaris 11: you can discover and manage Oracle Solaris 11 operating systems only if the Enterprise Controller and Proxy Controller are installed on Oracle Solaris 11 and the Oracle Solaris 11 Update Library is configured.
- Microsoft Windows: you must enable Windows Management Instrumentation (WMI) on Microsoft Windows systems before discovering them. Once WMI is enabled, they can be discovered normally.
- Sun SPARC Enterprise M-Series Servers: you must ensure that user privileges and the status of each dynamic system domain are correct before discovering a Sun SPARC Enterprise M-Series server.
- Oracle SPARC M5 and M6 Servers: Oracle SPARC M5 and M6 servers have ILOM server processors. You must ensure that you have access to the server with Ops Center Admin and Ops Center Security Admin user roles enabled before discovering an Oracle SPARC M5 or M6 Server.

> **Note:** M5 and M6 servers are supported, but some features have additional limitations. For more information see the `Target Servers` section of the *Certified Systems Matrix* document in the Oracle Enterprise Manager Ops Center document library.

- Oracle ZFS Storage Appliances: you must discover both the storage appliance and its service processor, and follow special procedures to manage them.
- Oracle Solaris Clusters: you must discover and manage Oracle Solaris Clusters in a specific order so that Oracle Enterprise Manager Ops Center can manage the entire cluster.

Windows Systems

Microsoft Windows systems can be discovered and managed using normal discovery and management procedures. However, before the system can be discovered, the Windows Management Instrumentation (WMI) utility must be configured and have access through the Windows Firewall or Internet Connection Firewall.

To Enable WMI

This procedure allows the Enterprise Controller or a Proxy Controller to connect to the target system.

1. Log in to the WMI on the target host.
2. Click Administrative Tools, then click Computer Management.
3. Expand Services and Applications.
4. On WMI Control, right click Properties.
5. Click the Security tab.
6. Click the Security button.
7. Select the Administrators group.
8. Select the option to allow Remote Enable.

To Allow WMI Through the Windows Firewall

This procedure allows WMI to send data through the target system's firewall.

1. Go to the command prompt on the target system.

2. Use the `netsh` command to allow WMI to send data through the firewall.
   - On Windows Server 2008 R2, use the following command:

     ```
     netsh advfirewall firewall set rule group="remote administration" new
     enable=yes
     ```

   - On other Windows systems, use the following command:

     ```
     netsh firewall set service RemoteAdmin enable
     ```

Servers

Most servers can be discovered using the normal discovery and management procedures. However, certain servers must be properly configured before they can be discovered and managed.

Discovering a SPARC Enterprise M-Series Server

To discover, manage, provision, and update a Sun SPARC Enterprise® M3000, M4000, M5000, M8000, M9000, or Fujitsu M10 server (SPARC Enterprise M-series servers), you monitor its XSCF service processor and its dynamic system domains.

The SPARC Enterprise M-series servers have a dedicated processor for system control that is independent of the system processor. A SPARC Enterprise M3000, M4000, and M5000 server has one service processor. The SPARC Enterprise M8000 and M9000 servers, each have two service processors; however, only one service processor is active at a time. The eXtended System Control Facility (XSCF) firmware runs on the dedicated service processor. The firmware manages hardware configuration, monitors cooling system (fan units), domain status, and error status, and can power on and power off peripheral devices.

The XSCF firmware can also create dynamic system domains. Each domain is a logical unit that can function as a system. An Oracle Solaris OS can operate in each domain.

Note: Fujitsu M10 servers require XSCF Control Package (XCP) firmware version 2050 (XCP2050) to use the Automatic Service Requests (ASR) functions and to use service tags to enable server discovery with the **Find Assets** function.

To Discover a SPARC Enterprise M-Series Server

To discover a Sun SPARC Enterprise® M3000, M4000, M5000, M8000, M9000, or Fujitsu M10 server, run an Add Assets Using Discovery Profile job for the XSCF service processor. The discovery job discovers the XSCF service processor and its dynamic system domains.

Perform the following tasks to discover this type of asset:

1. In the XSCF service processor, create a user account with `platadm` privilege if one does not exist.
2. Ensure that user privileges and the status of each dynamic system domain are correct.
3. Check the status of each dynamic system domain, using the `showdomainstatus -a` command. Oracle Enterprise Manager Ops Center can only discover domains that do not have a "-" status.
4. Log in to the XSCF shell from an XSCF-LAN port or from the serial port.
5. Discover the system using the Add Assets Using a Discovery Profile procedure.

A Discovery job is launched. Dynamic system domains and the XSCF service processor are discovered at the same time.

When the job is complete, the XSCF service processor and the dynamic system domains appear in the Managed Assets tab. Additionally, the service processor and the domains appear under the M-Series Servers group in the Assets tree.

See Related Resources for Asset Management for more information.

Discovering an Oracle SPARC M5 or M6 Server

To discover a configured Oracle SPARC M5 or M6 server, run an **Add Assets** job for an ILOM service processor discovery profile. The job discovers the ILOM service processor and its dynamic system domains.

Perform the following tasks to discover this type of asset:

1. Create a Discovery Profile choosing an ILOM Service Processor for the Asset Type.
2. Select existing or create new SSH and IPMI credentials.
3. Run an **Add and manage various types of assets via discovery probes** job.
4. Choose the Discovery Profile for the server and include the correct Hostnames or IP addresses.
5. After the discovery job completes successfully find the server in the Navigation pane under the M-Series Server tree.

See *Oracle Enterprise Manager Ops Center Discovering and Managing Oracle SPARC M5 and M6 Servers Guide* and Related Resources for Asset Management for more information.

Oracle ZFS Storage Appliance

The Oracle ZFS Storage Appliance family of products provides rich and efficient data services for file and block storage formats. Each appliance has the Analytics feature for observing the condition and behavior of the appliance in real time and the ZFS Hybrid Storage Pool that uses Flash-memory devices, high-capacity disks, and DRAM memory within a data hierarchy to provide solid-state response time with spinning disk capacity.

When you manage Oracle ZFS Storage Appliances within the Oracle Enterprise Manager Ops Center environment, you not only manage the appliance as one of the assets in the data center, but you can also make use of the storage provided by these appliances as a backing storage for storage libraries and software libraries.

Because the storage appliance contains a service processor, it is possible to discover the service processor but not the appliance, discover the appliance but not the service processor, or to discover both. The recommended procedure is to discover both aspects of the Oracle ZFS Storage Appliance.

- Use **Find Assets**. Each storage appliance is discovered as two assets: a service processor and a storage appliance. Select both the appliance and its service processor and enter the user credentials. Oracle Enterprise Manager Ops Center displays the asset in the Storage section.
- Use **Add Assets** and a discovery profile to discover the storage appliance first and then discover its service processor. When you discover the storage appliance, Oracle Enterprise Manager Ops Center displays the device in the Storage section of the Assets tree. When you discover the service processor, its information is mapped to the discovered appliance.

Note: If you change the order of discovery, the result is the same. However, Oracle Enterprise Manager Ops Center displays a generic asset in the Server section of the Asset tree. When the appliance discovery succeeds, Oracle Enterprise Manager Ops Center removes the generic server from the Servers section and displays a new asset in the Storage section.

Oracle ZFS Storage 7320 and 7420 Appliances provide a two-node cluster configuration. To discover the storage appliance, the administrative interfaces of both nodes must be private so that each node has a different static IP address. To verify that the appliance's nodes are using private administrative interfaces, you must use the appliance's user interface. Use this procedure for storage appliances with a cluster configuration.

1. Log into the Oracle ZFS Storage Appliance.
2. Select **Configuration** and then **Clusters**.

Figure 2–6 Configuration View of Oracle ZFS Storage Appliance

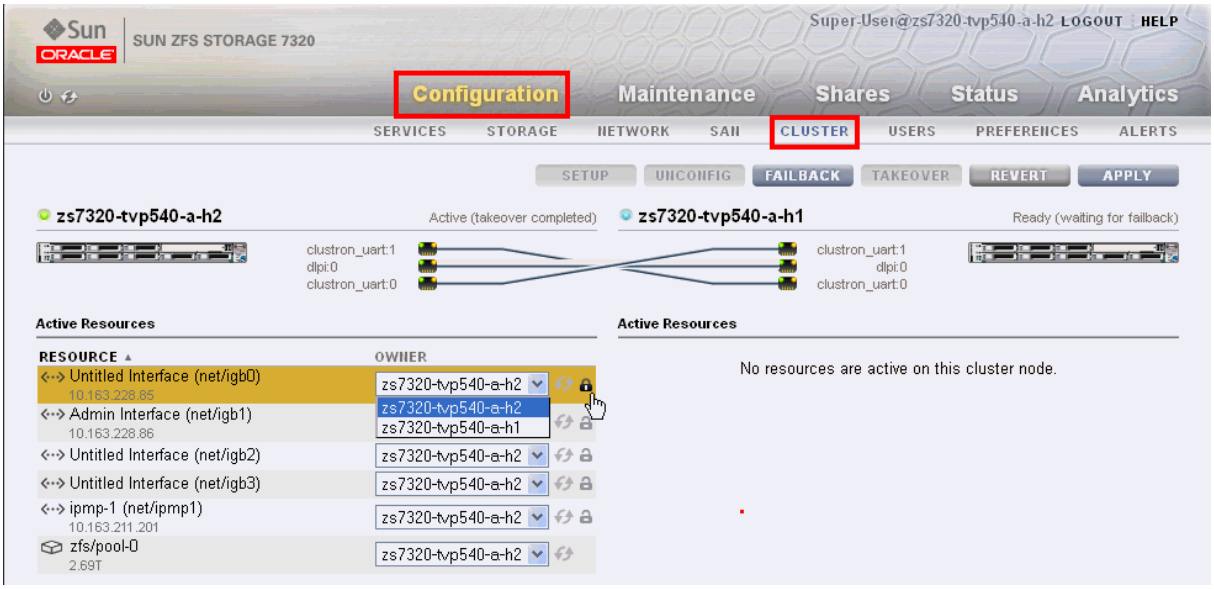

The Resource table includes the Administrative Interfaces. A private resource has an icon of a closed lock in bold.

3. To set the administrative interface as private, click on the open lock icon and then click **Apply**.

For information about configuring and discovering an Oracle ZFS Storage Appliance see the Deploy How To library: http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm.

Oracle Solaris Cluster

Discovering and managing an Oracle Solaris Cluster has some specific requirements. The discovery and management must be performed in the correct order.

Note: Do not configure a new zone cluster or add a new global node during the discovery and management process.

To Discover an Oracle Solaris Cluster

1. Verify that all global nodes in the cluster are in cluster mode.
2. Discover all of the cluster's global nodes using a discovery profile. You must provide both ssh credentials and create JMX credentials for each global node. The credentials authenticate the Oracle Solaris Cluster's agent.

Figure 2–7 Discovery Profile for Oracle Solaris Cluster

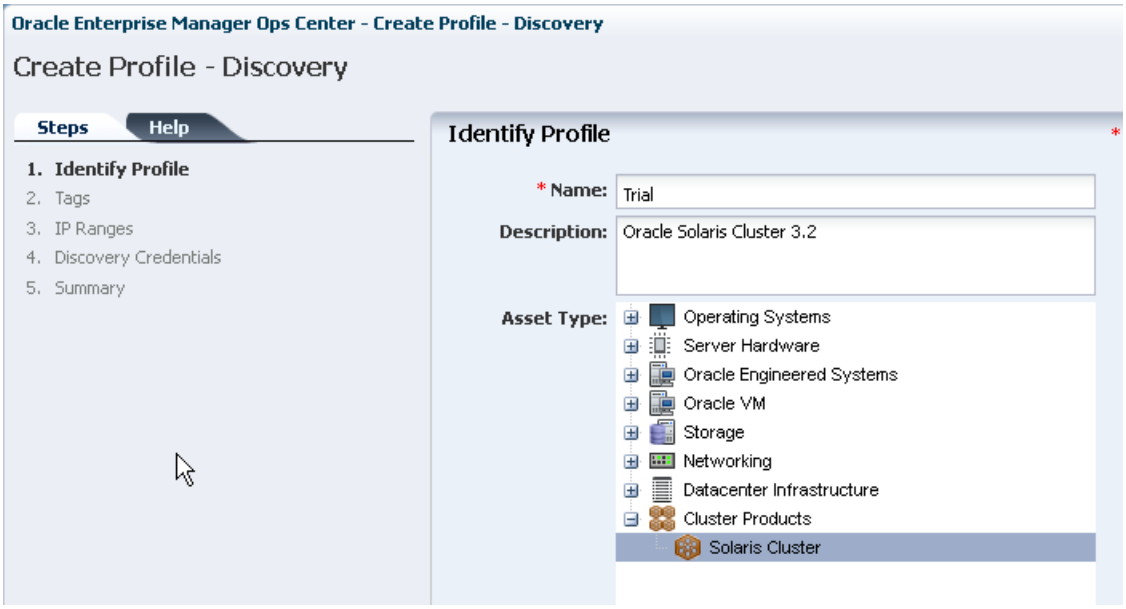

3. Discover the cluster using the discovery profile.

Starting in this version, you can launch the Oracle Solaris Cluster Manager:

1. In the Navigation pane, select **Solaris Cluster** and then select a cluster.
2. In the Actions pane, click **Launch Web Console**.
3. Log into the Oracle Solaris Cluster Manager.

Using Tags

A tag is a set of information attached to an asset. Each tag consists of a tag name, which is drawn from a list of values appropriate for each asset, and a tag value, which can be any text string. For example, an asset could have a tag with a tag name of oracle.cloud.resource.creation.time and a value of 12 June.

You can use tags to associate information with assets, and to group assets based on their tags or tag values. You can add tags to assets during or after discovery.

Adding Tags

To add tags to one or more assets, perform the following steps:

1. On the Navigation pane, select an asset or group.

2. In the Actions pane, click **Edit Tags**.

Figure 2–8 Add Tags

3. Click the add icon and select a tag name from the drop-down list. Enter a tag value, then click **Save**.

Viewing Tags

Tags are displayed in the dashboard tab for managed assets. Select an asset, then mouseover the tags icon in the center pane to display the tags.

Grouping Assets Using Tags

To add assets to a group using tags, create a new group or edit an existing group using the procedures described in the Using Groups section, and create one or more rules that add assets based on tags.

Deleting Tags

To delete tags, perform the following steps:

1. On the Navigation pane, select an asset or group.
2. In the Actions pane, click **Edit Tags**.
3. Select a tag, then click the Delete icon.
4. Click **Save**. The tag is deleted.

Using Groups

Assets are automatically placed into system groups based on asset type. You can also create your own groups and add assets to them, either by manually adding assets or by creating rules that add assets automatically based on asset characteristics.

Types of Groups

Two types of groups are available in Oracle Enterprise Manager Ops Center System groups, and User-defined groups. System groups are automatically generated for each category of asset. User-defined groups can contain any assets, and can be configured with rules to add assets that meet the rule criteria.

You can take any action on a group that is applicable to the assets in the group. For example, you can update Automated Service Request (ASR) contacts for all assets in a group, or target a group containing operating systems with an OS update job.

System Groups

Oracle Enterprise Manager Ops Center automatically creates groups of major asset types called system groups. Within these system groups, a subgroup is created for each type of asset that Oracle Enterprise Manager Ops Center is managing. You can

use System Groups to locate and view assets of a specific type. You can also act on System Groups, such as changing monitoring thresholds and updating management credentials.

You can view system groups by selecting them from the drop-down list at the top of the Assets section in the Navigation pane.

Oracle Enterprise Manager Ops Center creates these system groups:

- All Assets: contains all discovered and managed assets. This is the default view of the assets section.
- Engineered Systems: contains all Oracle Engineered Systems.
- Operating Systems: contains all operating systems with subgroups for each type of operating system such as Oracle Solaris OS, Oracle Linux, and SUSE Linux. The subgroups are further organized by version, such as Oracle Solaris 9 and Oracle Solaris 10 software.
- Servers: contains all hardware that can receive OS provisioning.
- Chassis: contains all hardware that can receive firmware updates but not OS provisioning.
- Network Switches: contains all network switches.
- Racks: contains all racks.
- Server Pools: contains all server pools.
- Storage: contains all storage systems.
- Solaris Clusters: contains all Oracle Solaris Clusters.

User-Defined Groups

User-defined groups can contain any type of asset and can be organized by any criteria. You can configure rules for user-defined groups that automatically add assets with specific characteristics to the group.

You can specify the following characteristics for a user-defined group:

- **Group Name**
- **Description**
- **Group Location**
- **Group Rules**: group rules add any assets to the group that match the attributes and rules.
- **Subgroups:** groups can be organized hierarchically.

Viewing Group Data

Each group can list all the assets in the group and can display data about their assets. Chassis and hardware groups display power usage information, and operating system groups display CPU, network, memory, and system load information.

You can select a group to see a dashboard page with information about the group, including:

- A group summary that shows the group's name, description, tags, location, and number of members.

- A membership graph showing the group's assets, any child groups, and any parent groups.
- A status summary showing the problems of the assets within the group.
- An asset summary showing basic data about the assets within the group.

Creating a Group

You can organize your assets into groups to aid in management and inventory.

To create a group, click Create Group in the Actions pane, then enter the group information:

- **Group Name**: this name is displayed in the User-Defined Groups of the Navigation pane.
- **Description (optional):** this is a description of the group that is displayed in the group's dashboard.
- **Group Location**: a group's position within the tree. You can create a group at the top level (root) or as a child of an existing user-defined group.
- **Configure Group Rules**: use this option to specify the membership rules. Any assets that match the rules are added to the group.
- **Configure Subgroups:** use this option to specify any existing user-defined group as a child of this group.
- **Preview Group Before Creation**: if you are configuring group rules or subgroups, select this option to see the assets and subgroups that will be added to the new group before it is created.

If you are creating a rules-based group, provide the rule information:

- **Matching Policy:** a rule can contain one or more attributes. For each rule, specify whether an asset must match all of the rule's attributes or any of the rule's attributes.
- **Asset Type:** select the asset type that uses the rule.
- **Attribute:** select an attribute to be considered by the rule. For example, if you select Name, the rule compares the asset's name to a given value. The available attributes vary depending on the selected asset type. You can add additional attributes to a rule using the Add Attribute icon.
- **Condition**: select a condition for the rule. For example, if you select Contains for the Name attribute, the rule includes assets if their name contains a given value.
- **Value**: enter a value to be used by the rule. For example, if you select Name as an Attribute, Contains as a condition, and Pod3 as a value, any asset of the specified type with Pod3 in its name is added to the group.

If you are adding subgroups, drag and drop one or more groups from the available groups list to the selected groups list.

If you chose to preview the group, view the assets and subgroups that will be added to the group. If necessary, go back and change the group rules or the list of subgroups.

See *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for a list of the asset attributes that can be used in group rules.

Editing a Group

You can change the attributes of an existing group, including its name, description, rules, subgroups, and parent group. If you change or remove the rules of an existing group, assets added by those rules are removed if they do not match the new rules. However, assets that were added manually can only be removed manually.

To edit a group, perform the following steps:

1. On the Navigation pane, select All User Defined Groups from the Assets drop-down list.

Figure 2–9 All User Defined Groups

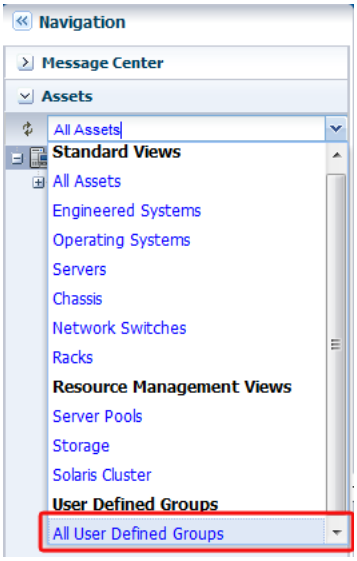

2. In the Actions pane, click Edit Group.

> **Note:** You can edit any of the group's characteristics. If you edit a group's rules, assets that were added by group rules will be removed if they do not meet the new rules, but assets that were manually added will not be affected.

See *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for a list of the asset attributes that can be used in group rules.

Adding Assets to a Group

You can place assets into user-defined groups to organize them. When an asset is added to a group, it continues to be displayed in the All Assets section. Assets can be added to any number of groups.

To add assets to a user-defined group, select the assets in either the managed assets tab of the center pane or the Assets section of the Navigation pane. Click **Add Asset to Group**, select the destination group, and click **Add Assets to Group**.

Removing Assets from a Group

You can remove assets from any user-defined group. The assets remain in any other groups.

Note: If an asset was added to a group by the group's rules, it cannot be manually removed.

To remove an asset from a group, select the group. Select the asset and click **Remove Asset from Group**.

Moving Assets to a Group

When you move assets to a new user-defined group, the assets are removed from the current group and added to the new group.

Note: If the current group has rules that match the asset you want to move, the asset is not removed from the original group.

To move assets to a new group, perform the following steps:

1. On the Navigation pane, under Assets, select **All User Defined Groups**.
2. Select the group that currently contains the assets.
3. Select one or more assets and click **Move Asset to Group**.
4. Select the destination group and click **Move Assets to Group/Subgroup**.

Moving a Group

You can move a user-defined group within the hierarchy of groups, making it a child of a different user-defined group or placing it at the top level.

To move a group, select the group and click **Move Group**. Select a destination group and click **Move Group**.

Deleting a Group

Deleting a user-defined group removes the group and all of its subgroups from the user-defined group hierarchy without removing any of the assets.

To delete a group, select a group and click **Delete Group** in the Actions pane.

Related Resources for Asset Management

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- See Chapter 4, "Monitoring Rules and Policies" for information about the monitoring rules and policies to generate alerts and incidents in the user interface.
- See Chapter 23, "Oracle Engineered Systems" for information about Oracle Engineered Systems management.
- See Oracle SuperCluster for more information about discovering the Oracle SPARC SuperCluster T4-4, T5-8, and M6-32 systems.
- For end-to-end examples, see the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm`. The following documents relate to asset management in Oracle Enterprise Manager Ops Center:
  - *Deploy Hardware Workflow*

  - *Discovering and Managing Hardware Guide*
  - *Discovering and Managing Oracle SPARC M5 and M6 Servers Guide*
  - *Discovering and Managing a Fujitsu M10 Server Guide*
  - *Discovering and Managing an Oracle SPARC T5 Server Guide*
  - *Adding InfiniBand Switches Guide*
  - *Discovering and Managing a Power Distribution Unit Guide*
- For information about Oracle SPARC servers, including SPARC T5, SPARC M5-32, and SPARC M6-32 servers, see SPARC Systems at http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html.
- See Systems Management and Diagnostics at http://www.oracle.com/technetwork/documentation/sys-mgmt-networking-190072.html for information about ILOM configurations.

3

Jobs

This chapter describes in detail about jobs, the different roles for managing jobs, viewing jobs, and the different actions that can be performed on jobs.

The following information is included:

Introduction to Jobs

Any action performed by Oracle Enterprise Manager Ops Center creates a job. This job is either run on the Enterprise Controller system or picked up by a Proxy Controller to be run on a Proxy Controller or a managed asset. The progress of the job is tracked and displayed in the user interface. Each job is made up of one or more tasks, and has one or more targets; information on the job's progress on any task or target can be viewed.

The Jobs pane is located at the bottom of the user interface and displays a list of all current and historical jobs. From the Jobs pane, you can monitor the progress of current jobs and can also review historical jobs. You can view the status of all jobs, view detailed information about specific jobs, and take actions on jobs.

To view the details of the job, double-click the job in the Job Summary table. The Job Details view displays the targets of the job, and also the following information:

- Job status:
  - Running – The job is in progress.
  - Waiting for User Input – The job has started, but needs information from a user before it can be completed.
  - Failed – The job was not successful.
  - Partially Successful – Some tasks were completed successfully. This could be a job with multiple tasks, in which some tasks completed but others failed, or a

job with multiple targets, in which the job was successful for only some targets.

- Stopped – The job was stopped by the user.
- Scheduled – The job has been scheduled to run at a specific time. It might be a one-time job or a recurring job.
- Successful – The job has completed. All of its tasks were completed successfully.

- Job ID – A unique identification number for the job.
- Type of job – For example, the discovery custom type identifies a job as a result of a custom discovery action.
- Name of job – A name for the job.
- Run ID – If a job has been run multiple times, each run of the job has a separate run ID.
- Mode of job – Simulated or Actual Run. Some jobs can be simulated. Simulated jobs verify that the necessary permissions, images, space, and other job requirements are satisfied. To perform an actual run of a simulated job, you must create a new job.
- Owner of job – The user who launched the job.
- Start / Scheduled Date – The date and time when the job was started or is scheduled to start.
- Elapsed time – The amount of time the job has been running, if the job is running. The amount of time it took for the job to complete, if the job has completed.
- Description – Description of the job.
- Failure policy – The failure policy of the job, can be one of the following:
  - Continue on Failure – if a task fails, continue to run other tasks.
  - Abort on Failure – if a task fails, all remaining tasks are aborted.
  - Rollback on Failure – if a task fails, abort all remaining tasks, and run the rollback method on the failed task and all previous successful tasks.
- Next Scheduled Time (For recurring scheduled jobs) – This option is available only in the job summary table. It displays the time when the job started. If it is scheduled to run in the future, then it displays the next run time.
- Next Scheduled Date – The next date that the job will be run, or null if the job will not run again.
- Start Date – The date/time that the run of this job started.
- Creation Date – The date/time that the job was first created.
- Task Execution Order:
  - Sequential – the tasks are run one after another.
  - Parallel – all the tasks run at the same time.
  - Sequential_on_task – this applies to tasks that target more than one target. The task must complete on all targets before the next task begins.
- List of tasks – This option is available on the job target details view. It displays the following:

  - the task target
  - the last log entry for the task, or the result of the task if the task is completed.
  - the elapsed time – the amount of time that this task has been running.
- Task execution order – The execution order of the task.
- Task progress/results – Progress or result of the task.
- Event logs – Event logs are generated by tasks as they progress.

You can perform the following actions on jobs:

- View details about the job, including the status of all tasks.
- Respond to prompts during a job.
- Interrupt or stop a running job. You can re-run the job later.
- Remove a job from the history.
- Re-run a job that had partial success, failed, or was stopped.
- Create a new job with a new Job ID based on a previous job; for example, to repeat a successful complex multi-step job on another target.

Roles for Job Management

A cloud user can only view or act on jobs that were created by the user. The cloud user cannot view jobs that were created by another user.

If a job has multiple targets, you can see only the targets for which you have the correct role. If you do not have the role for any of the targets, the job is not displayed.

The following table lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 3–1 Job Management Tasks and Roles

Task	Role
Viewing Job Status	The same as the role required to launch the job.
Viewing Job Details	The same as the role required to launch the job.
Monitoring Jobs for an Asset	The same as the role required to launch the job.
Answering Questions	The same as the role required to launch the job.
Stopping a Job	The same as the role required to launch the job.
Re-running a Job	The same as the role required to launch the job.
Copying a Job	The same as the role required to launch the job.
Deleting a Job	The same as the role required to launch the job or Job Management.
Changing the Maximum Time for a Job	Root access on Enterprise Controller system.

Job Management

The job management service is responsible for managing the execution of jobs and monitoring them. The job manager communicates instructions in the task to one or more agents in the order specified by the task. The job manager monitors the state of execution, logs each successful or unsuccessful operation, and maintains the information state that enables the job manager to report success or failure.

Within job management, you can perform the following actions, depending on what jobs have been launched and their status:

- View all jobs
- View all jobs with a specific status
- View details for a job
- Monitor jobs for an asset
- Search for jobs
- Answer questions for a job
- Stop a job
- Re-run a job
- Re-run a job on failed targets
- Copy a job
- Delete a job
- Debug a job using the OCDoctor
- Change the job properties

Viewing Jobs

You can view jobs in different ways. The Jobs pane displays all current and past jobs that have not been deleted. You can also view the details of a specific job, and view all jobs with a specific status or for a specific asset.

Viewing All Jobs and Jobs With Specific Status

You can view a list of all jobs run by Oracle Enterprise Manager Ops Center.

To view all jobs, perform the following steps:

1. Click and expand the Jobs pane at the bottom of the UI.

   The Jobs pane is displayed as shown in the following figure.

Figure 3–1 Jobs pane

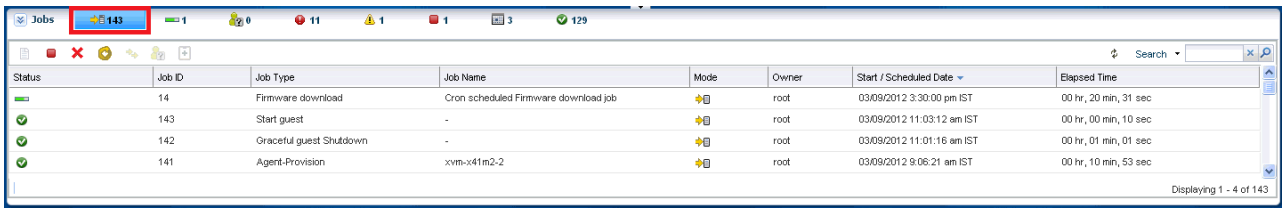

2. In the Jobs pane, click the **All Jobs** icon, highlighted in Figure 3–1.

All the jobs are displayed.

3. To view jobs with a specific status (such as jobs in progress, jobs waiting for user input, failed jobs, partially successful jobs, stopped jobs, scheduled jobs, and successful jobs), click the respective job status icon. The jobs with the selected status are displayed.

Viewing Job Details

You can view detailed information about a job, including the status of the tasks that make up the job.

To view job details, perform the following steps:

1. Select a job in the Jobs pane.

Figure 3–2 View Job Details

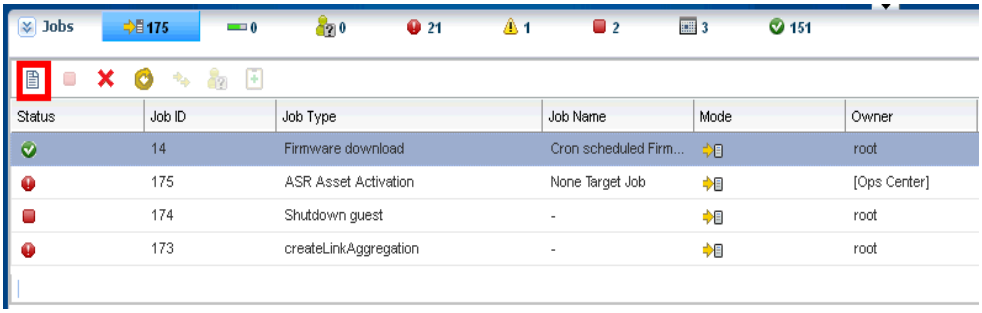

2. Click the **View Job Details** icon, highlighted in Figure 3–2 or double-click the job. The Job Details window is displayed.

Figure 3–3 Job Details Window

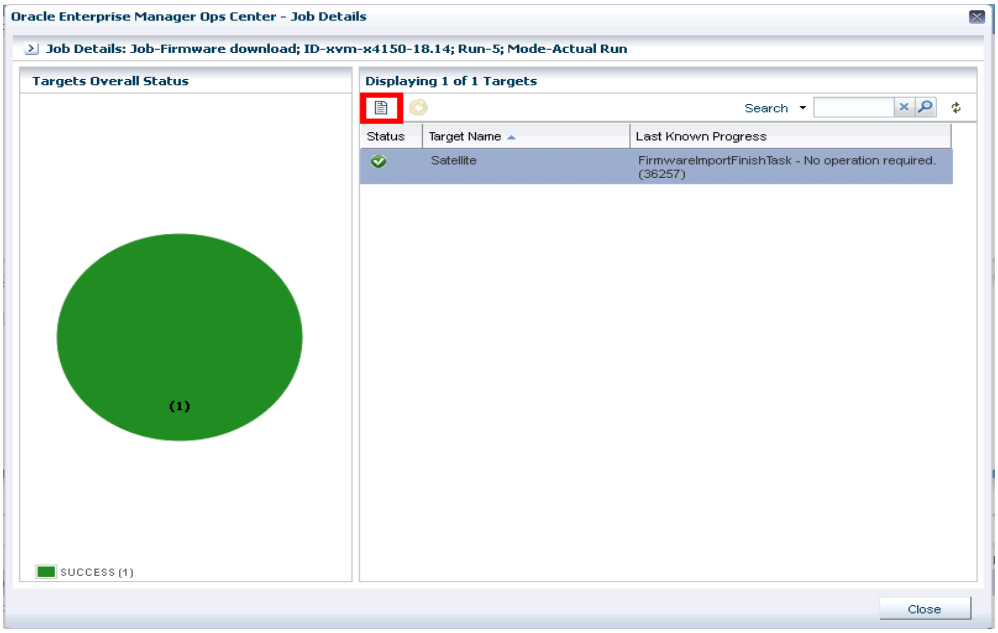

The overall status of the target is displayed on the left. On the right, a table lists all the targets.

3. Select a target and click **Display Selected Target Details** icon, highlighted in Figure 3–3. You can also double-click the target to display the target details.

   The Job Target Details window is displayed.

Figure 3–4 Job Target Details Window

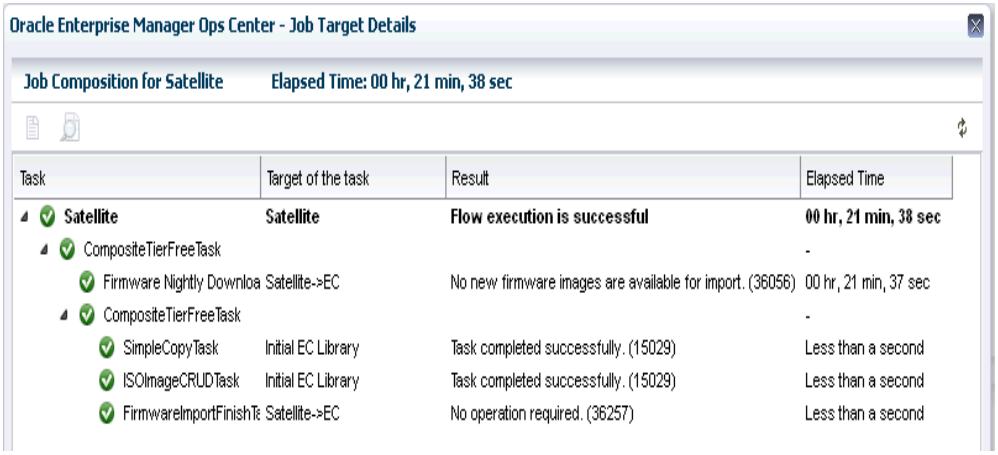

4. Select a target and click **Display Selected Task Details** icon. You can also double-click a task.

   The Task Logs window is displayed with the Event Logs tab active, as show in Figure 3–5.

Figure 3–5 Task Logs Window

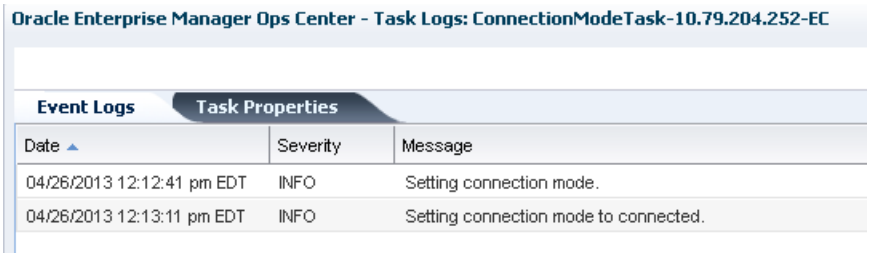

   You can click the **Task Properties** tab to see additional information about the task, such as the Task ID and the permissions used to complete the task.

5. Click **Export Logs** on either tab to save the event log as a text file. The View Target Logs window is displayed with two options:
   - **Selected Task Target Log** exports the current task.
   - **Full Job Log** exports all tasks.

6. Click **Save** to export the event log to a file and location of your choice.

You can also view and export the event log for a task or job by selecting the **View Target Logs** icon on the Jobs Target Details window, as show in Figure 3–4.

Monitoring Jobs for an Asset

You can view jobs for a particular asset, both running jobs and jobs that have been completed.

To view only those jobs for a particular asset, perform the following steps:

1. In the Navigation pane, under Assets, select an asset.
2. In the center pane, click the **Jobs tab**.

Figure 3–6 Jobs pane in Assets View

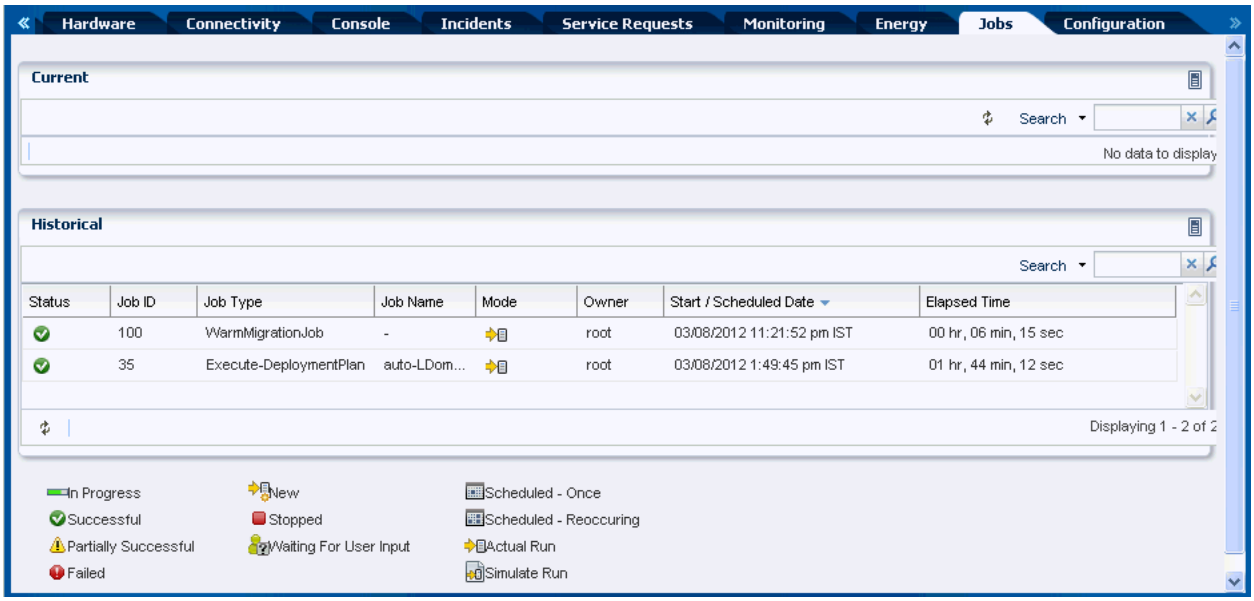

Figure 3–6 shows the Jobs pane for a selected asset. The Current and Historical tables show the jobs for the asset.

3. Select a job and double click to view the details of the job.

Search for Jobs

You can search for specific jobs in any state. To search for a particular job, perform the following steps:

1. Click **Search** in the upper right corner of the Jobs pane.

Figure 3–7 Jobs Search

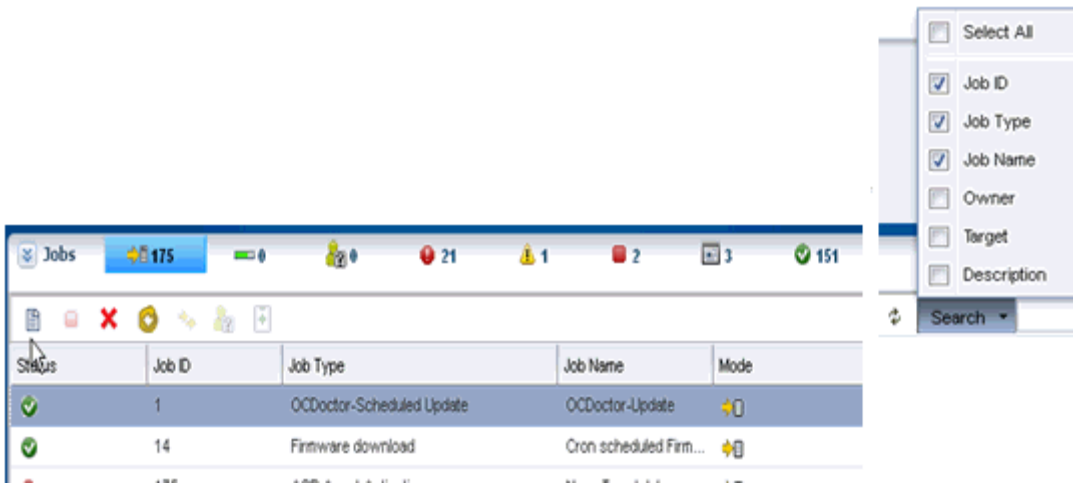

2. Select one or more job characteristics to search, then enter a search term and click the **Search** icon.

   Jobs for which the selected characteristics match or include the search term are displayed.

Taking Actions on a Job

You can take actions on jobs that have been launched. These actions vary depending on the current status of the job. If a job has stopped and is waiting for user input, you can answer questions and resume the job. If a job is running, you can stop it. If a job has completed, you can re-run it entirely, re-run it only on failed targets, copy it to create a new job, or delete it.

Answering Questions for a Job

Some jobs require user input to complete. A job that requires a response has the status of Waiting for User Input. To provide input, perform the following steps:

1. Select a job with the Waiting for User Input status.
2. Click the **Answer Question**s icon.
3. Select **Yes** or **No** for each question, or click **Yes to All** or **No to All**.
4. When you have answered each question, click **Submit** to resubmit the job using the same Job ID and Run ID.

Stopping a Job

You can stop a job that is running. All tasks in progress are interrupted and tasks that have not yet started are aborted. The completion status of the job depends on the number of tasks that have completed, been interrupted, and not started.

To stop a job, perform the following steps:

1. Select a job that is running.
2. Click **Stop Selected Job**s.
3. Click **Stop Job** to confirm.

Re-running a Completed Job

If a job is completed, you can repeat the job with a new Run ID. The re-run option is disabled for some jobs.

To re-run a job, perform the following steps:

1. Click and expand the Jobs pane.
2. Select a job.
3. Click **Re-Run Selected Jobs**.
4. Click **Run Job**.

The job is re-run with a new Run ID.

Re-running a Job on Failed Targets

If a job is partially completed, failed, or stopped, you can repeat the job on failed or incomplete targets. To re-run a job on failed targets, perform the following steps:

1. Select the job and view the job details.
2. Select one or more failed targets from the list of targets and click **Re-Run Selected Failed Targets**, then click **OK**.

The job is re-run with the same Run ID on the failed targets.

Copying a Job

You can copy an OS or firmware update job, using a completed job as a template for a new job. To copy a job, perform the following steps:

1. Select the job that you want to copy and click **Copy Job**.

   A new job wizard is displayed, using the information from the existing job.
2. Modify the job information, select a schedule, and click **Submit Job**.

The new job is submitted with a new Job ID.

Deleting a Job

Deleting a job removes it from the queue entirely. It cannot be re-run or resumed, and its job details are not available.

> **Note:** You cannot delete jobs that are running. You must first stop the job and then delete it.

To delete a job, perform the following steps:

1. Select the job that you want to delete.
2. Click **Delete Selected Jobs**.
3. Click **Delete Jobs** to confirm.

Debugging a Job Using the OCDoctor

You can debug a job that has been run on a managed asset. This action runs a self-diagnosis using the OCDoctor's `--troubleshoot` option on the asset.

To debug a job, perform the following steps:

1. Select the job and click the **Debug Job Using OCDoctor** icon.
2. Select either **Run New Self Diagnosis** to run a new diagnosis or **Work on Previous Self Diagnosis** to view the results of a prior diagnosis. If you select Work on Previous Self Diagnosis, the results of the prior self-diagnosis are displayed, and you can re-run the self diagnosis.
3. Attempt to fix issues by selecting **Attempt to Fix Issues**, or select **Collect Logs** to collect log files from the target.
4. Review the summary information and click **Finish** to launch the self-diagnosis job.
5. Review the job details or rerun the wizard and select **Work on Previous Self Diagnosis** to view the data gathered by the self diagnosis.

Changing the Job Properties

The Enterprise Controller sets several variables for job management, including the time-out values for jobs and tasks. You can modify the values to improve performance. The following are the time-out values for jobs and tasks:

- jobtypeweight.FirmwareProvisioningJob – This value sets the relative weight of firmware provisioning jobs. Do not modify this value unless directed by Oracle Support.
- taskTimeOut – This value sets the time, in minutes, after which a task times out and fails.
- tierTimeOut – This value sets the time, in minutes, after which a job that has not been picked up by a Proxy Controller fails.
- totalAllowedWeight – Do not modify this value unless directed by Oracle Support.
- totalDispatchWeight – This value sets the total weight of jobs that can be performed in parallel. Increase this value if your Enterprise Controller system can handle additional jobs.

To modify the Job Manager values, perform the following steps:

1. In the Navigation pane, select **Administration**, then select the Enterprise Controller and click the **Configuration** tab.
2. In the drop-down menu, select **Job Manager properties**.
3. Edit the values, then click **Save**.

Events for Job Management

To follow the progress of a job, view Job Details. If a job does not complete successfully, you can examine the log file in the following location:

On Oracle Solaris system: `/var/cacao/instances/oem-ec/audits`

On Linux System: `/var/opt/sun/cacao2/instances/oem-oc/audits`

Viewing Audit Logs

Oracle Enterprise Manager Ops Center logs events to create a record of the following operations:

- Adding and deleting a user account
- Changing the roles for a user account
- Connecting to the Enterprise Controller
- Starting and ending jobs

The audit log contains information about:

- Date and time of event
- User ID executing a job
- Type of event
- End of the job
- Details of the login connection: type of connection, port, and IP address

Note: You must have a `root` access to view the log files.

Logs cannot be edited or modified by any user.

Related Resources for Job Management

For instructions on performing actions or to learn more about the role of this feature, refer to the following:

- *Oracle Enterprise Manager Ops Center Administration Guide*
- See the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm`.
- See the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm`.
- See Performance and Security in the Appendix B, "Logs and Directories" for more information on audit log files.

Part II

Configure

Part II contains the following chapters:

4

Monitoring Rules and Policies

This chapter discusses the types of monitoring rules and how the software uses monitoring rules and policies to generate alerts and incidents in the user interface.

The following information is included:

Introduction to Monitoring Rules and Policies

Monitoring detects components or attributes of a managed resource that are not operating within specified parameters. Resource is a generic term for an asset (such as hardware or operating system), a group, a network, or a library that is managed by Oracle Enterprise Manager Ops Center.

Monitoring rules and policies define the monitoring parameters and are the main components of a complete monitoring configuration. They are defined as follows:

- Monitoring Rules: Define alerting conditions. You can apply one or more rules to an asset to monitor the asset and raise an alert when the an asset is operating outside the defined parameter.
- Monitoring Policies: A set of monitoring rules targeted to a specific asset type. Default monitoring policies contain a set of rules that are automatically applied to monitored resources. You can create your own policies, define the rules for the policy, and make those policies the default policies for new assets.

A monitor gathers information about the condition of a resource. When you activate a monitoring rule on an asset, a monitor is activated on the asset's management access point (usually the Agent Controller or the Proxy Controller.) When an operating system is an agentless managed system, the software monitors the operating system remotely.

Roles for Monitoring Rules and Policies

Table 4–1 lists the roles required to complete monitoring tasks. You can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 4–1 Monitoring Tasks and Roles

Task	Role
View a monitoring rule	Read
	Plan/Profile Admin
View a monitoring policy	Read
	Plan/Profile Admin
View the association of an asset and a monitoring policy	Read
	Asset Admin
	Plan/Profile Admin
View the historical data of a threshold rule	Read
	Asset Admin
Create, edit, or delete a monitoring rule	Fault Admin
Create, copy, extract, edit, and apply a monitoring policy	Plan/Profile Admin
Modify the monitoring configuration of an asset	Fault Admin
Delete a monitoring policy	Plan/Profile Admin
Group assets by a monitoring policy	Asset Admin
Apply a monitoring policy to an asset	Fault Admin
Apply a monitoring policy to a group	Asset Admin and Fault Admin

Actions for Monitoring Rules and Policies

You can perform the following actions, depending on the type of rules, the requirements, and your roles:

- Create, edit or delete a monitoring rule.
- Create, copy, extract, edit, apply, or delete a monitoring policy.
- Modify the monitoring configuration for an asset.
- Group assets by a monitoring policy.
- Apply a monitoring policy to an asset or a group.

Location of Monitoring Rules and Policies in the User Interface

Table 4–2 Location of Monitoring Rules and Policies in the UI

Object	Location
Monitoring policies	Expand **Plan Management** in the Navigation pane, then click **Monitoring Policies**.
Monitoring rules	Expand **Plan Management** in the Navigation pane, then click **Monitoring Policies**. Open a policy to see the rules.
Asset-specific monitoring rules	Select an asset from the **Asset** view, then click the **Monitoring** tab.

Monitoring Rules

Monitoring rules state the values and boundaries for an asset's activity. A monitoring policy is a set of rules. When you apply monitoring policy to all the assets, it enforces consistency in monitoring. Each monitoring policy contains rules for threshold levels. Default policies for monitoring hardware, operating systems, and Oracle Solaris Clusters are included in the software. You can use the default policies, but you cannot edit them. To edit or add monitoring rules to a monitoring policy, you must make a copy and then set that policy to be the default.

Monitoring rules define the alerting conditions. Rules are associated with, and determined by, the type of managed resource. You can apply a generic monitoring rule to many different attributes, but other monitoring rules are attribute-specific, hard-coded into drivers and cannot be relocated or reconfigured.

When you set a threshold, the UI displays the existing historical data for an attribute. The software might propose some default threshold values based on the analysis of historical data, and display existing thresholds as a bar on that historical data. You can enter any threshold value.

Each managed resource has a Monitoring tab. You can add, edit, enable, disable, and remove resource monitoring rules. You can tune the rules for a specific managed resource.

> **Note:** Tuning rules for a specific managed resource detaches the resource from the monitoring policy, keeping only a copy of the rules. When you modify rules in the policy, the change is not made to the rules that are associated with the resource.

Rules have Info, Warning, and Critical severity levels. Default values and severity levels are provided at installation, but you can edit the rules for your organization. For user-defined rules, you can define the time between when the alerting condition occurs and when the software generates an alert or incident. You can configure the software to send an e-mail or pager message when it identifies a Warning or Critical incident.

You can view rules for a specific asset or a specific policy from different places in the UI, as follows:

- **Asset View**: Rules for a specific asset are located in the Monitoring tab for the asset. To display the tab, click the asset in the Assets navigation tree, then scroll over to the Monitoring tab in the center pane. The name of the monitoring policy applied to the asset appears at the top of the monitoring rule grid, next to the number of rules.

- **Policy View**: Rules for a specific policy are located under Monitoring Policies in the Operational Plans section of the Plan Management tree.

The following categories of monitoring rules are available:

- **System-defined rules**: These are attribute- specific and are hard-coded into drivers. You can disable a system-defined rule, but you cannot edit, relocate, or reconfigure these types of rules.
- **User-defined rules**: These are associated with, and determined by, the type of managed resource. You can apply a user-defined rule to many different attributes.

User-Defined Rule Parameters

The following types of editable user-defined rule parameters, also known as rule types, are available:

- **Threshold**: Sets an upper or lower monitoring threshold for the monitored attribute.
- **Boolean Control**: Sets a logical operator of true or false for the monitored attribute.
- **Enumerated Control**: A series of values that defines a subset of specific values among the possible values of the monitored attribute. An alert occurs when the attribute matches one of those specific values.
- **Expression**: Defines the variables, literals, and operators for an attribute. An expression is an instruction to execute something that returns a value.

You cannot modify all rules, but most rules include some parameters that you can tune, or edit, to meet your organization's requirements.

The following are some examples of editable parameters:

- Severity level of the alert: You can define the parameters for informational, warning, and critical alerts.
- Raising and clearing values: These are threshold settings that determine when an alert is raised and cleared. These two values are always the same. For example, you can configure the software to raise an alert when a value reaches 90% and clear the alert when the value falls below 90%.
- Monitor for alert limits at specific time: Defines when you want monitoring to occur, or to not occur. You might use this parameter when a daily maintenance procedure causes an attribute to operate outside of the normal monitoring threshold, but you do not want to raise an alert. You can define a period of time when the monitors are disabled and you can perform maintenance.
- Generate alert after: Defines how long an issue occurs before an alert is generated. The number defines the time between when a threshold is exceeded and when an alert is generated. The alert is not triggered immediately. An alert is generated when the monitored attribute value is outside the threshold after the specified delay. You might use this parameter to limit false positive alerts due to a temporary condition.

Enabled and Active Rules

Monitoring rules have two types of states:

- **Enabled or Disabled**: Disabling a rule removes that attribute from monitoring. You can disable and enable rules on a per asset or group basis.

- **Active or Inactive**: Reflects the system's state and indicates whether the software is monitoring the asset or group. When a rule is not enabled, monitoring is not active.

By default, all monitoring rules are enabled. The status appears on the Alert Monitoring Rules page, which you can access from the Monitoring tab. When *Yes* is in the Enabled and Active fields, the rule is enabled and active. When *No* appears in the corresponding field, a rule is disabled or inactive. Figure 4–1 shows the Service Alert Monitor status as enabled, but inactive.

Figure 4–1 Enabled and Active Monitoring Rules

An **Enabled** field appears in the list of monitoring rules for an asset. You can disable one or more rules for a specific asset. When *No* appears in the Enabled column, the rule is disabled.

When a rule is enabled, the active state reflects the system's actual state and indicates whether the software is using the rule. The following are some reasons that an enabled rule might be inactive:

- When a specific attribute is not hard-coded into the driver, monitoring is not possible for that attribute.
- The software cannot reach the resource or the attribute cannot be refreshed.
- Some type of misconfiguration, such as a missing mandatory parameter or an illegal value for a parameter.
- An internal error specific to the monitor, particularly for driver-specific monitors.

You can view the rules for a specific asset or you can view the rules for a policy that is associated with one or more assets.

Editing Monitoring Rules

Monitoring rules have pre-defined parameters. You can change the parameters, including the threshold values and the monitoring level, to meet your data center guidelines. You can define the parameters for an individual asset or for a group of assets.

You can create separate monitoring groups to consistently and efficiently define the parameters for all systems in each group. For example, you can create a group for a set of high priority systems. You edit one set of specific threshold values for all members

of that group and apply a monitoring policy to the group. All members of the group are now monitored in the same way for the same values.

You can edit a monitoring rule for a specific asset or the monitoring rule for a monitoring policy. Editing the monitoring rule parameters for an individual system might be useful when a particular system is on a critical path. For example, you might consider the systems that you have the Enterprise Controller and Proxy Controllers on critical path systems, you can monitor the system continuously and create more stringent monitoring thresholds for those systems.

> **Note:** When you update a monitoring rule in a policy, the monitoring configuration of all assets that are associated with the policy are updated to reflect the revised rule.

You can perform the following tasks:

- Change the values for the Warning and Critical thresholds
- Change file system thresholds
- Change the thresholds by system or by group
- Set specific threshold values for different operating systems

For example, you can create a threshold on the Enterprise Controller system that sends a warning when the file system use exceeds 90%. This alerts you when the Enterprise Controller file system is almost full.

For a threshold alert, you can change how often and for how long the software monitors the resource. You can change the threshold values.

- **Alert window**: Enables you to specify a period of the day when the monitoring rule is enabled. For example, when a daily maintenance operation causes a monitored attribute to exceed a threshold, you can exclude monitoring for that time to disable monitoring for that maintenance window.
- **Generate alert after**: Enables you to configure monitoring to ignore a monitored attribute that is outside the defined monitoring parameters for a short period of time. Specifying a delay means that the software generates an alert only when the value remains above the specified limit for a given duration. The software does not generate an alert when the value goes above the limit once and then immediately goes back to normal.

To Edit a Monitoring Rule

Perform the following steps to edit a monitoring rule:

1. Determine whether to edit a rule for a policy or for an individual asset:
   - For a policy: From the Plan Management section of the Navigation pane, click **Monitoring Policies**, then click a policy.
   - For an asset: From the Assets section of the Navigation pane, select an asset, then click the Monitoring tab.
2. Click the **Edit Alert Monitoring Rule Parameters** icon.
3. Click the entry in the Value column and type the new value. For a threshold alert, you can change how often and for how long the value is monitored, and you can change the threshold values.
4. Click **Apply** to submit the changes.

See the *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for more about the attributes that you can use in monitoring rules.

Using Historical Data to Determine Threshold Limits for a Specific Asset

The software maintains a history of each monitored asset's performance against its assigned threshold rules. When you add a new threshold-type monitoring rule or modify a threshold configuration for a specific asset, it does not impact the asset's attribute history.

The asset's attribute history, which is maintained by the software, is not specifically tied to the threshold rules. Oracle Enterprise Manager Ops Center only records statistical values of the asset attributes over time. However, the threshold wizard can leverage the history to suggest meaningful threshold values. When there is historical data for the monitored attribute, a graphical representation of the historical data appears along with a proposed default threshold value based on the analysis of the data.

When modifying a threshold configuration for a specific asset, you can choose the time frame from a list of options, from one day up to six months, to display a graphical representation of the historical data. You can use this information to tune the threshold rules.

Note: When you edit the threshold limit for a specific asset, the software disassociates the asset from the default monitoring policy and creates a new monitoring policy for the asset.

Adding Monitoring Rules

Each monitoring policy contains a default set of rules. See Monitoring Policies for information on policies. The rule set and default parameters depend on the managed asset subtype. See Monitoring Rules for a definition of type.

When specifying an Expression monitoring rule, you use the Oracle Enterprise Manager Ops Center query language to write a logical expression that defines the alerting condition for one or more resource attributes. The logical expression includes attribute names, operators, and literal values. You can use a dotted notation to reference attribute subfields.

When adding a Threshold, Enumerated, or Boolean monitoring rule, you must define the monitored attribute. When adding an Expression rule type, specify an expression that references one or more attributes to monitor. The following are some examples of monitored attributes:

- `CpuUsage.usagePercentage`
- `ProcessUsage.topMemoryProcesses.pid=*.physicalMemoryUsage`
- `DiskUsageSet.name=*.busyPercentage.`

Monitored attributes are available in the Javadoc that is in the Oracle Enterprise Manager Ops Center Software Developer's Kit (SDK). Go to the `dvd/`*`platform`*`/Product/components/packages` directory and install the SUNWxvmoc-sdk.pkg package.

Browse the available attributes and names for the monitoring framework. Attributes always start with an upper case letter, such as `SystemUpTime`, and fields always start with a lower-case letter.

For example, when you want to list the valid monitored attributes for an operating system, go to the *com.sun.hss.type.os.OperatingSystem Javadoc* page. This page displays all of the attributes of an OperatingSystem.

Each of these attributes is either a simple type, a structure or struct-like type, or a collection type. The following are examples of the different types of attributes:

- Simple: You can use the name, such as `SystemUpTime`
- Struct-like: You can drill-down into a field of the structure type. The fields always start with a lower-case letter, such as `SystemLoad.average1Minute`
- Collection: You can drill-down into a member of the collection. For Maps you do this by specifying the appropriate 'key'. When you set the key, specify the value for the 'name' field, to get the value of a single member. For example, use the following to check the 'enabled' value of the interface named *eth0*: `InterfaceInfos.name=eth0.enabled`

Note: Structures are sometimes nested. For example, a struct-like attribute can contain another struct-like field, or a collection. Collections typically contain struct-like values. To drill down, continue to append the appropriate field names.

You can perform a query which scans across all members of a collection by specifying the '*' wildcard value for the key or name. When you perform a query, you must use one of the following operators: *max*, *min*, or *like*. For the query syntax, see the *DomainQuery Java class* Javadoc. See the *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for details on the Javadocs.

See the *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for more information on the expression query language, grammar, lexical elements, and method details.

Monitoring Policies

Monitoring policies contain the information needed to monitor a managed asset with alert configurations, including defined thresholds and alert monitors. You can revise many of the threshold and monitor settings, or use one of several methods to create new policies.

The following information is in this section:

System and User Defined Monitoring Policies

A monitoring policy is a collection of rules that are associated with each type of monitored asset. The policy defines the resources monitored and the rules for that asset type. You can add and edit policies and select the default policy for a specific asset type.

A set of default monitoring policies is included with the software. The policies are based on the type of resource being monitored, such as operating systems, servers, power distribution units, and SAN server arrays, and Oracle VM Server. Each policy contains a default set of rules. The rule set and parameters depend on the managed asset subtype. Alert monitors watch the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

You can use the default rules and policies or you can use the following methods to create a new policy:

- Copy an existing policy and then edit it, as needed, to modify the rules.
- Extract an existing policy from an asset and modify it.
- Create a new policy, then edit the policy to add rules.

When you discover and add an asset, the software applies the default monitoring policy for the asset's type immediately. Some monitoring policies install monitors or agents on managed resources, while other policies are designed to invoke arbitrary actions or scripts against the managed resource.

Go to the Plan Management section of the UI for a list of available monitoring policies. Click Monitoring Policies to display a list of all policies, the default status of the policy, and the intended asset or target type. The following types of monitoring policies are available:

- **System Defined Policies**: Typically use asset specific rules. The monitored details are determined by what is implemented on the specific asset. An example of a system-defined policy is the MSeriesChassis policy that monitors the Xsb Mode of a SPARC M-series chassis. A system-defined policy is read-only, you cannot disable or modify the rules defined in the policy. You can turn the rule off and on.
- **User-Defined Policies**: Typically use generic rules and contain monitoring thresholds. An Operating System policy is an example of a user-defined policy that monitors the following generic operating system parameters: CPU usage, disk I/O queue length and utilization percentage, file system usage, memory usage, network bandwidth utilization, and swap usage, and system load.

The Navigation pane and center pane both display user-defined policies followed by system-defined policies. The icons and naming convention help to identify the type of policy in the Navigation pane. The names of all system-defined policy use the prefix OC and have a green check mark in the lower. For example, in Figure 4–2 *OC - Chassis* is a system-defined policy and *Copy of OC - Chassis* and *OS Profile - Dev* are user-defined policy.

Figure 4–2 List of Monitoring Policies

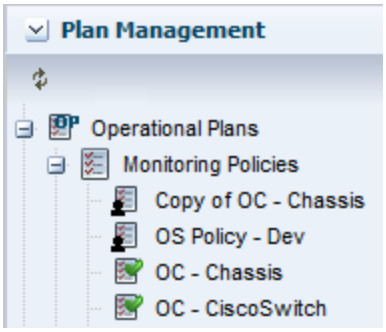

User-defined policies have a silhouette of a person in the lower left corner. A green check mark in the icon indicates that this is a default monitoring policy. In Figure 4–2, OC - Chassis is a system defined policy that is the default policy. In Figure 4–3 Copy of OC - Chassis is a user defined policy and is also the default policy.

Figure 4–3 Default and User Defined Policies

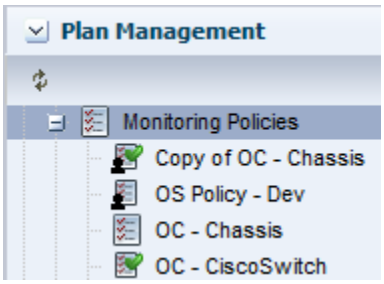

To view a detailed list of user defined and system defined policies, click **Monitoring Policies** in the Navigation pane, as shown in Figure 4–4.

Figure 4–4 List of Policies, by Type

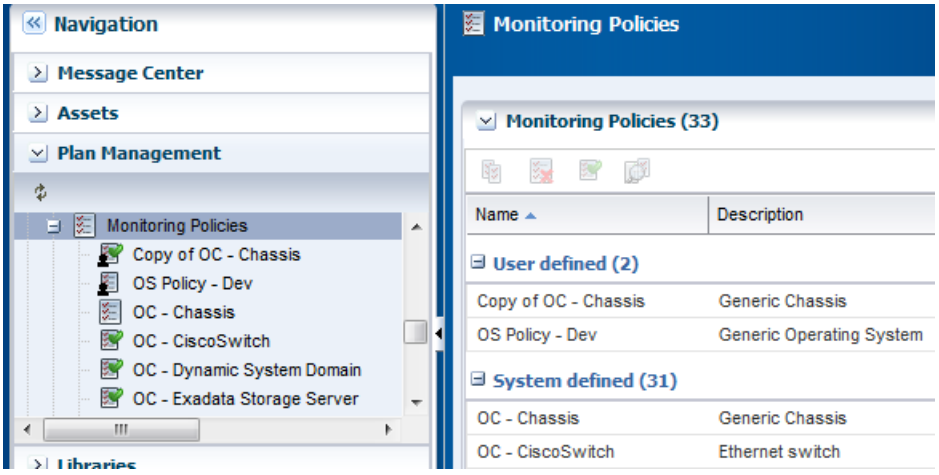

You can create your own user-defined policies by creating, copying, or extracting a monitoring policy. You perform the copy policy action from the policy view and the extract policy action from the asset view. By extracting a policy from the asset view, you ensure that the subtypes are valid for that asset type. You can change the target

subtype to a more specific or a more generic eligible target type. To be eligible, the policy must be a member of a more general policy for the specified target.

For example, you can highlight an Oracle Solaris operating system and extract a monitoring policy for an operating system. You can specify one of the following as the OS subtype for the new policy:

- Oracle Solaris 10: Any supported Oracle Solaris 10 operating system.
- Oracle Solaris 10 Operating System: Any supported Oracle Solaris operating system, beginning with Oracle Solaris10 8/07 (update 4). You might use this subtype when you use the Oracle Solaris Zones policy.
- Oracle Solaris: Any supported Oracle Solaris operating system release.
- Operating System: Any supported Oracle Solaris, Linux, or Windows operating system.

Monitoring Policy Details

Each asset has a Monitoring tab. The tab displays monitoring configuration details, such as the associated monitoring policy.

When an asset is managed, the software automatically assigns a default monitoring policy to the asset. The name of the associated monitoring policy appears above the Monitoring Rules table. The field is empty when an asset is not associated with a policy. If you associate the asset with a user-defined monitoring policy and then later delete that policy, the asset might not have an associated monitoring policy.

Monitoring policy details vary, depending on the asset and associated resources being monitored. Figure 4–5 is an example of a system-defined operating system monitoring policy, including the Policy Details. The Details section of the page contains information about the policy, including the name, description, and the type of policy. The policy details also describes the applicable types of assets or targets, whether the policy is a sub-type of another policy, when the policy was last modified, and whether the policy is the default policy for the target types.

Figure 4–5 System-Defined Operating System Monitoring Policy

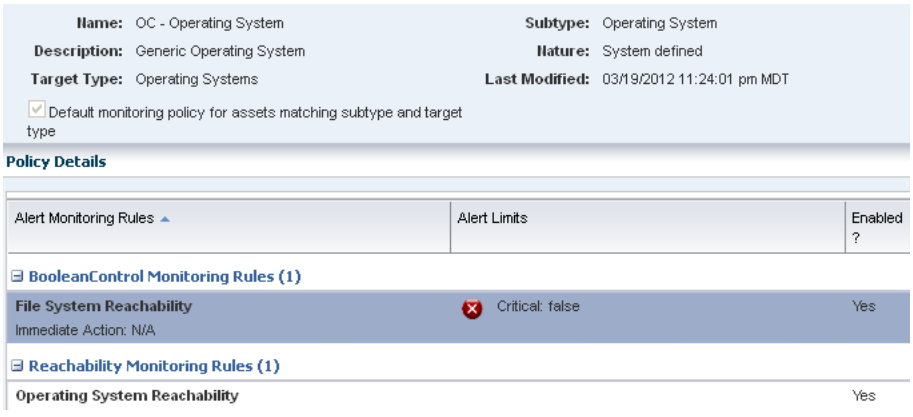

The following system-defined and generic monitoring policies are available:

- OC – Chassis: Monitors the chassis fan and power supply
- OC – CiscoSwitch: Monitors an Ethernet switch's power status, switch port status, and switch status.

- OC – Dynamic System Domain: Monitors the state and status of the dynamic system domains on eligible servers. Server reachability is monitored. Informational alerts are available for the power state, either on or off, and the operating system state, either running or not. The Server Port Status generates a Warning alert when the server port is disabled or down.
- OC – File Server: Monitors the file server reachability, backing devices usage percentage, storage allocation percentage, and storage usage percentage of generic file servers.
- OC – Global Zone: Monitors the DHCP status, appliance health, CPU usage, disk I/O, file system usage, memory usage, network bandwidth, Swap usage, and system load of a global zone.
- OC – ISCSI Storage Array: Monitors the iSCSI storage array reachability, storage allocation percentage, storage usage percentage, volume group allocated space percentage, and volume group used space percentage.
- OC – Local Library: Monitors the storage library usage percentage for a local storage library.
- OC – Logical Domain: Monitors Oracle VM Server for SPARC guest status, including the migration status, running or not, and whether the guest is powered on.
- OC – M-Series: Monitors the Xsb Mode of a SPARC M-series chassis
- OC – NAS Library: Monitors the storage library status and storage library usage percentage of network attached storage (NAS) libraries.
- OC – Non-global Zone: Monitors CPU usage, disk I/O queue length and utilization percentage, memory usage, network bandwidth utilization, and swap usage for non-global zones.
- OC – Operating System: Monitors the following generic operating system parameters: CPU usage, disk I/O queue length and utilization percentage, file system usage, memory usage, network bandwidth utilization, and swap usage, and system load.
- OC – Oracle VM Server for SPARC: Monitors Oracle VM Server for SPARC status, including the DHCP client, appliance health, and free virtual CPU (VCPU) usage.
- OC – Oracle VM Server for x86: Monitors Oracle VM Server for x86 status, including the DHCP client, appliance health, and free virtual CPU (VCPU) usage.
- OC - Power Distribution Unit: Monitors Ampere levels of rack Power Distribution Units.
- OC – Remote Oracle Engineered System: monitors remote Oracle engineered systems.
- OC – SAN Library: Monitors the status of the storage area network (SAN) libraries.
- OC – SAN Storage Array: Monitors the SAN storage array reachability and usage. The usage includes the storage allocation percentage, storage usage percentage, volume group allocated percentage, and volume group used space percentage.
- OC – Server: Monitors the following parameters on a generic server: Power status, server port status, CPU, NIC, fan, fan tray, memory, and power supply.
- OC – Server Pool: Monitors the total CPU allocation percentage, total cryptographic units allocation percentage, and the total memory usage percentage.

- OC – Solaris Cluster: Monitors the Oracle Solaris cluster install mode, reachability, and monitor.
- OC – Solaris Cluster Node: Monitors the online status, scalert.node reachability, and scalert.node alert monitor.
- OC – Solaris Cluster Zone Cluster Group: Monitors the scalert.zone cluster alert monitor.
- OC – Solaris Cluster Zone Cluster Node: Monitors the scalert.zone Cluster Node Alert Monitor
- OC – Storage: Monitors the following parameters for a generic storage device: operating system status, power status, server port status, and storage alert.
- OC – Switch: Monitors an Ethernet switch's power status, switch port status, and switch status.
- OC – Virtual Machine: Monitors the reachability of the virtual machine, if the operating system is running, and the migration status of a virtual machine.

Copying or Creating a Monitoring Policy

You can add a new monitoring policy or copy an existing monitoring policy, or you can extract a policy. See Extracting a Monitoring Policy for more details on extracting a policy for a specific type of asset.

To Copy or Create a Monitoring Policy

1. Expand **Plan Management** in the Navigation pane, then click **Monitoring Policies**.
2. Click **Copy Policy** or **Create Policy** in the Action pane.
3. Enter a name and description for the new policy, then select the subtype and target type.
4. Click **Next**, then click **Finish** to add the policy to the list of available policies.
5. To see the new policy, expand **Plan Management** and click **Monitoring Policies**. Click the policy to display details.
   - To add or remove rules or change monitoring parameters, double-click the policy in the center content pane.
   - To view all assets that are monitored with this policy, click **View Associated Assets**.
   - To make this policy the default monitoring policy, click the **Set as Default Policy** icon.

Extracting a Monitoring Policy

Extracting a monitoring policy is similar to copying a policy. Copying a policy is performed from the policy view, while extracting a policy is performed from the asset view, as shown in Figure 4–6.

Figure 4–6 Extract a Monitoring Policy Action

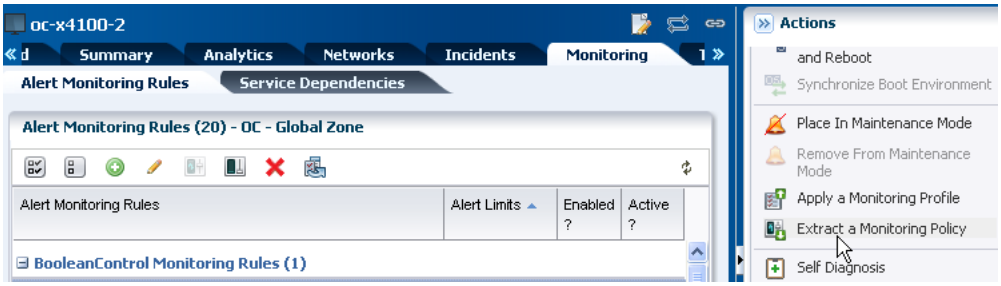

By extracting a policy from the asset view, you filter the possible subtypes to only those that are valid. You can change the target subtype to a more specific or a more generic eligible target type. To be eligible, the policy must be a member of a more general policy for the specified target. If a list of valid target types does not display, then the target type cannot be changed.

You can specify one of the following as the OS subtype for the new policy:

- Oracle Solaris 10: Any Oracle Solaris 10 operating system.
- Oracle Solaris 10 Operating System: Any Oracle Solaris 10 8/07 (update 4) or higher operating system. You might use this subtype if you want to use the policy for zones.
- Oracle Solaris: Any supported Oracle Solaris operating system release.
- Operating System: Any supported Oracle Solaris, Linux, or Windows operating system.

To Extract a Monitoring Policy

1. Expand **Assets** in the Navigation pane, then click an asset type, such as operating system.
2. Click the **Monitoring** tab, then click **Extract Monitoring Policy** in the Action pane.
3. Enter a name and description for the new policy, then select the subtype and target type. Click **Next**.

Figure 4–7 Extract Monitoring Policy

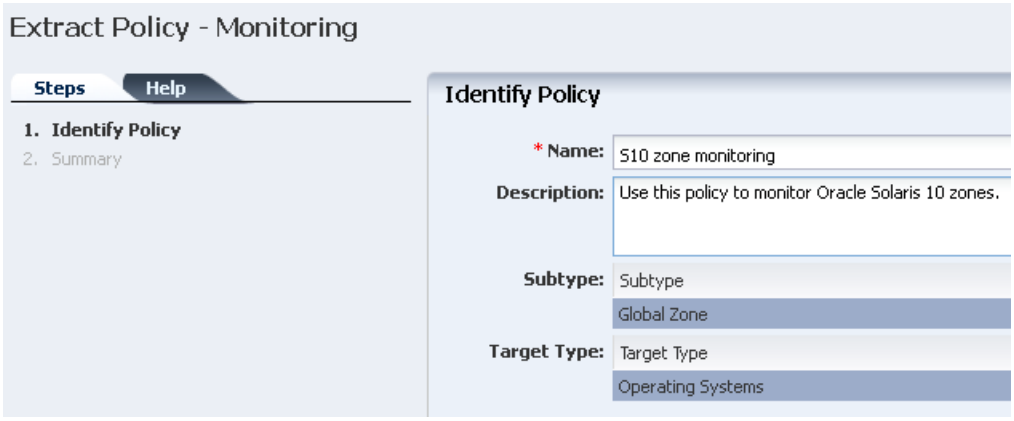

4. Click **Finish** to add the policy to the list of available policies.

5. To see the new policy, expand Plan Management and click **Monitoring Policies**. Click the policy to display details.
   - To add or remove rules or change monitoring parameters, double-click the policy in the center pane.
   - To view all assets that are monitored with this policy, click **View Associated Assets**.
   - To make this policy the default monitoring policy, click the **Set as Default Policy** icon.

Editing the Monitoring Configuration of an Asset Bound to a Monitoring Policy

Monitoring configurations contain monitoring policies and are associated with an asset type, such as an operating system. You can edit specific rules and parameters, but you cannot edit general monitoring policy properties, such as name and description after you create the policy. Editing a policy changes the monitoring configuration for all associated assets.

You can modify the monitoring configuration of an individual asset. When you create a monitoring configuration, the asset is no longer associated with the monitoring configuration and policies for that asset type. Instead, it has its own independent monitoring configuration.

When you modify a monitoring configuration for an asset, you create a policy and the asset is associated with the new created policy.

Applying a Monitoring Policy to a Group

You can apply, or associate, a monitoring policy with a user-defined group. When you apply a monitoring policy to a group, all applicable members of the group are associated with the policy. The policy associated with an asset appears next to the asset in the wizard. When you apply a policy to all assets in a group, the software disassociates the asset from the assigned policy and associates it with the policy that is assigned to the group.

When you remove the asset from the user-defined group, the asset is associated with the group's monitoring policy until you relocate it to another user-defined group that has a different policy or you manually associate the asset with a different monitoring policy.

Creating a Group of Assets According to Monitoring Policy

When you create a user-defined group, you can specify a rule that filters out the assets based on the name of the associated monitoring policy. You can use this to group all the assets associated with a given monitoring policy or to find assets that are not associated with a specific policy.

When viewing the monitoring configuration of an asset, you can navigate to the definition of the monitoring policy bound to the asset.

Deleting a Monitoring Policy

You can delete user-defined monitoring policies in the Plan Management section. Resources are associated with a monitoring policy. When you delete a user-defined monitoring policy, any asset or group associated with that policy is automatically detached from the policy. However, the asset does retain the monitoring configuration that was defined by the policy.

You cannot delete system-defined policies.

Disabling and Enabling Monitoring Policies

Monitoring policies are enabled by default. You can disable all of the monitoring policies for your data center. When you disable monitoring policies, the monitors are no longer deployed on the assets and incidents and alerts are not generated.

Note: Disabling monitoring policies disables the evaluation of monitoring rule conditions against collected data and prevents the deployment of monitors across your data center, it does not disable the collection of data on managed assets.

When you enable monitoring that you previously disabled, the software applies the monitoring rules that are defined by the default monitoring policies to all of the assets.

You must use the command line interface to disable or reenable monitoring policies. See *Oracle Enterprise Manager Ops Center Command Line Interface Guide* for more information.

To temporarily disable incidents from generating on a single asset, or a group of assets, you can place an asset in Maintenance Mode. This is useful when you need to perform maintenance on an asset and you want to temporarily disable the incidents that might generate during system maintenance. See Using Maintenance Mode for information about temporarily disabling incidents.

Related Resources for Monitoring Rules and Policies

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources in the *Oracle Enterprise Manager Ops Center Feature Reference Guide*:

- See Chapter 11, "Hardware" for information about hardware monitoring.
- See Chapter 12, "Operating System Management" for information about monitoring operating systems.
- See Chapter 9, "Incidents" for information about managing incidents that result from monitoring.

See the *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide* for details about the expression query language, grammar, lexical elements, and method details.

The following How To documents in the Operate How To Library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm` relate to monitoring:

- *Oracle Enterprise Manager Ops Center Managing Incidents*
- *Oracle Enterprise Manager Ops Center Tuning Monitoring Rules and Policies*
- *Oracle Enterprise Manager Ops Center Understanding OS Performance and Capacity*
- *Oracle Enterprise Manager Ops Center Using Service Requests*

5

Software Libraries

The following information is included:

Introduction to Software Libraries and Repositories

Oracle Enterprise Manager Ops Center uses libraries to store and manage cached data, images, packages, and metadata. A library that stores images for provisioning operations is a Software Library. At least one software library always exists on the Enterprise Controller.

A Software Library accepts the following types of images:

- OS images that install an operating system
- Branded images that install a specialized version of an operating system
- Firmware images and the supporting metadata to update existing firmware on service processors, RAID controllers, and disks

Another type of library is a storage library. See Chapter 16, "Storage Libraries for Virtualization" for information about this type of library.

Roles for Software Libraries

Table 5–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See

the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 5–1 Software Libraries Tasks and Roles

Task	Role
Set Enterprise Controller Storage Library	Ops Center Admin
Create Library	Storage Admin
Delete Library	Storage Admin
Associate Library	Storage Admin
Import image	Storage Admin
Upload image	Storage Admin
View details of an image	Storage Admin
Moving an image	Storage Admin
Edit Attributes	Storage Admin
Associate Library to Server Pool	Cloud Admin

Actions for Software Libraries

Perform the following actions, depending on the requirements:

- Designate a default software library
- View details of an image
- Create Library
- Delete Library
- Associate Library
- Import Image
- Upload Image
- Download an OS image
- Moving an image
- Edit Attributes
- Create a library of OS images and manage content

Location of Software Library Information in the User Interface

Table 5–2 shows where to find information.

Table 5–2 Location of Library Information in the BUI

Object	Location
To see the software libraries	Expand Libraries in the Navigation pane, then expand Software Libraries.
To see firmware and OS provisioning profiles	Expand Plan Management in the Navigation pane, then click Profiles and Policies.

Knowledge Base and Parent Repository

By default, Oracle Enterprise Manager Ops Center operates in Connected mode. In this mode, the Enterprise Controller uses the Internet to download images and metadata from the Oracle Knowledge Base and the Oracle Solaris 11 Package Repository.

- The Oracle Knowledge Base (KB) contains metadata for Oracle Solaris 10, 9, and 8 and Oracle Linux OS components.
- The Oracle Solaris 11 Package Repository includes packages of images in IPS format for Oracle Solaris 11.

To use the Oracle Enterprise Manager Ops Center software without Internet access, your site can maintain a local version of the Knowledge Base or Repository on your site's network. In this case, the Enterprise Controller connects to the local network location to get the latest information.

EC Library

The Enterprise Controller must have at least one Software Library to store the new versions of images that are downloaded from the Oracle Knowledge Base. Oracle Enterprise Manager Ops Center updates this library each week, by default.

In addition to the routine download operations, you can create jobs to update assets. When you submit an update job for specific target assets, the Agent Controllers on the targets send a request to the KB through the Enterprise Controller to download the latest information.

The product installation procedure creates the Initial EC Library. At any time, you can specify a different software library to accept the automatic download operations from the Knowledge Base. A badge identifies the current active library. Figure 5–1 shows the badge, a white down arrow on a green background.

Figure 5–1 Badge Identifying the Current Default Library

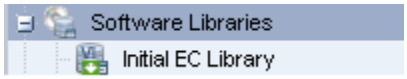

To Change EC Software Libraries

To specify a different software library as the default software library:

1. Create a software library, as described in Creating a Software Library.
2. Expand **Administration** in the Navigation pane. An alternative is to expand Libraries in the Navigation pane.
3. Click **Set Enterprise Controller Storage Library** in the Action pane. The window lists all libraries with the current library highlighted.
4. Click the new software library.
5. Click **Apply**.

When the job is completed, the Library section of the Navigation pane shows the software library you selected as the default library.

When the product software is updated, the schedule of download is interrupted. Use the **Set Enterprise Controller Storage Library** to restore the schedule.

Publishers and Parent of the Oracle Solaris 11 Repository

The Oracle Solaris 11 Package repository is similar to the Knowledge Base but is only for Oracle Solaris 11 images and updates. The Oracle Solaris 11 Package Repository resides at the Oracle site: `http://pkg.oracle.com/solaris/release`.

The Support Repository Updates (SRUs) for Oracle Solaris 11 contains bug fixes or minor feature enhancements and is released monthly. This repository, `https://pkg.oracle.com/solaris/support`, is available to users with an Oracle support agreement.

Both repositories have the role of Publisher and Parent for your local repository.

Other parent repositories, such as Oracle Solaris Cluster, are available. For a list of available Oracle repositories and to download the key-certificate pair, see the `https://pkg-register.oracle.com` site.

Using Software Libraries

In addition to the EC Library, you can create more Software Libraries and organize their content, according to your site's purposes. You can use a file system on the Enterprise Controller's system or a shared file system on an NFS server that the Enterprise Controller mounts. The file system on the Enterprise Controller is called a local software library. The file system on the NFS server is called a NAS software library. See Types of Storage for Libraries.

When Oracle Enterprise Manager Ops Center provisions target systems with an operating system or firmware, it copies the images files from the designated Software Library to the Proxy Controller that manages the target. The Proxy Controllers handle the provisioning operations. When Oracle Enterprise Manager Ops Center provisions target systems with an update to an operating system, it uses the software library named Linux and Oracle Solaris 8-10 Software Update Library or the Oracle Solaris 11 Software Update Library.

Viewing the Contents of a Software Library

You can display the contents of the software library, its associations, and details about the disks in the software library. You can also see how Oracle Enterprise Manager Ops Center monitors the library and any problems.

To View the Contents of a Software Library

1. Expand **Libraries** in the Navigation pane.
2. Click a software library.

   The details of the selected library are displayed in the center pane in a set of tabs.

   The Summary tab displays information about the entire software library:

   - URL – File for a local library and NFS for a NFS share
   - Size – Total storage capacity of the library
   - Used Space – Percentage of used space as compared to available space
   - State – Status of the library
   - Access – Read-Write

   The Library Contents table lists all the images in the library, organized by type, and includes the size and the date the image was modified. The other tables on the

Summary tab describe the types of images: Service Processor Firmware, Component Firmware, and BIOS Configuration Snapshots

3. To see the results of monitoring the software library, click the **Incidents** tab.
4. To see the attributes and values that are being monitored, click the **Monitoring** tab.

Creating a Software Library

You can create a software library that uses space on a file system on the Enterprise Controller's system, which is a Local Software Library. You can also create a software library that uses space on a shared file system on an NFS server. This type of library is a NAS Software Library.

Note: When you use both local and NAS software libraries, do not use the same name for the library.

Creating a Local Software Library

1. Expand **Libraries** in the Navigation pane.
2. Click **New Local Software Library** in the Action pane.
3. Enter a unique name and description.
4. In the URL field, enter the location of the file system.
5. Click **Create**.

Creating a NAS Software Library

1. Expand **Libraries** in the Navigation pane.
2. Click **New NAS Software Library** in the Action pane.
3. Enter a unique name and a description.
4. Select one or more server pools to use this storage library.
5. Choose the type of service:
   - To use a storage device, select the storage device and specify the exported share to use.
   - To use an NFS server, enter the host name, port, and path.
6. Click **Create**.

Libraries for Oracle Solaris 11

The Oracle Solaris 11 Image Packaging System (IPS) contains the packages that you need to install, provision, and update your Oracle Solaris 11 operating system. Each IPS package has an associated manifest that describes how the package is assembled. The package manifest provides basic metadata about the package (such as name, description, version, and category), what files and directories are included, and the package dependencies. Packages might specify the services to restart to refresh some configuration on the system, specify the aliases to update for a given hardware driver, or the users and groups to create as part of the package installation process. A package repository holds all software packages and systems must connect to the repository to install software updates.

Oracle Solaris 11 Software Update Library

The Enterprise Controller can maintain a repository for the Oracle Solaris 11 Image Packaging System (IPS). The repository, called the Oracle Solaris 11 Software Update Library, provides the images for provisioning assets with the Oracle Solaris 11 operating system.

> **Note:** The host system for the Enterprise Controller must use the Oracle Solaris 11 operating system. Do not attempt to initialize an Oracle Solaris 11 Software Update Library on a different Oracle Solaris operating system.

Options for Configuring the Oracle Solaris 11 Software Update Library

You have several options for when and how you configure the Oracle Solaris 11 Software Update Library:

- If your site already maintains an Oracle Solaris 11 IPS Repository, direct Oracle Enterprise Manager Ops Center to use it, instead of initializing a new one. Specify the location of the existing repository during the product installation or after the product is installed. See Using an Alternate IPS Repository.
- At any time, you can create an Oracle Solaris 11 IPS repository at your site and use it as the Oracle Solaris 11 Software Update Library. See the *Copying and Creating Oracle Solaris 11 Package Repositories* at `http://docs.oracle.com/cd/E23824_01/html/E21803/toc.html`
- During installation of the product software, initialize the library. This library downloads content from the `https://pkg.oracle.com/solaris/support` repository and continues to sync with the repository routinely. However, the initialization of the library can take many hours, depending on your site's access.
- After installation, initialize the library using the **Initialize Oracle Solaris 11 Software Update Library** action. This action performs the same operation as the installation option with the same time requirements.

To connect to `https://pkg.oracle.com/solaris/support`, either at installation or at a later time, you must provide a key file and a certificate file to authenticate the connection. If these files are missing or have expired, provide a new key and certificate, using the procedure in the `Oracle Enterprise Manager Ops Center Administration Guide`.

Library States

The Oracle Solaris 11 Software Update Library has the following states:

- Unconfigured – Oracle Solaris 11 Software Update Library was not created at installation. The **Initialize Oracle Solaris 11 Software Update Library** action is available in the Actions pane.
- Configuring – The Oracle Solaris 11 Software Update Library is in the process of being configured. The process of initializing the library takes many hours. When complete, the new library appears in the Libraries section.
- Syncing – The Oracle Solaris 11 Software Update Library is in the process of updating its contents with the parent IPS repository or is in the process of being reconfigured. The Oracle Solaris 11 Software Update Library is locked and unavailable for use when in this state.
- OK – The Oracle Solaris 11 Software Update Library is ready to use.

Summary of Oracle Solaris 11 Software Update Library

When the Oracle Solaris 11 Software Update Library is configured, its Summary tab, shown in Figure 5–2, provides an overview of the health, status, and last update. The URL listed on the Summary is the location of the Oracle Solaris 11 Software Update Library. The size is the amount of space allocated to the file system, and the used space shows the amount of space used by the packages that are located in the library.

The Publishers table shows each Parent Repository and its credentials. The URL is the location of the Parent Repository for Oracle Solaris 11 packages and content located in the Image Packaging System (IPS). Oracle Enterprise Manager Ops Center uses these URLs to synchronize the information in the Oracle Solaris 11 Software Update Library with the Image Packaging System. The Parents list is created when you use the **Configure Parent Repositories** Wizard. If your site has its own local repository for Oracle Solaris 11 packages, update Oracle Enterprise Manager Ops Center's library from that location.

Figure 5–2 Oracle Solaris 11 Software Update Library Summary

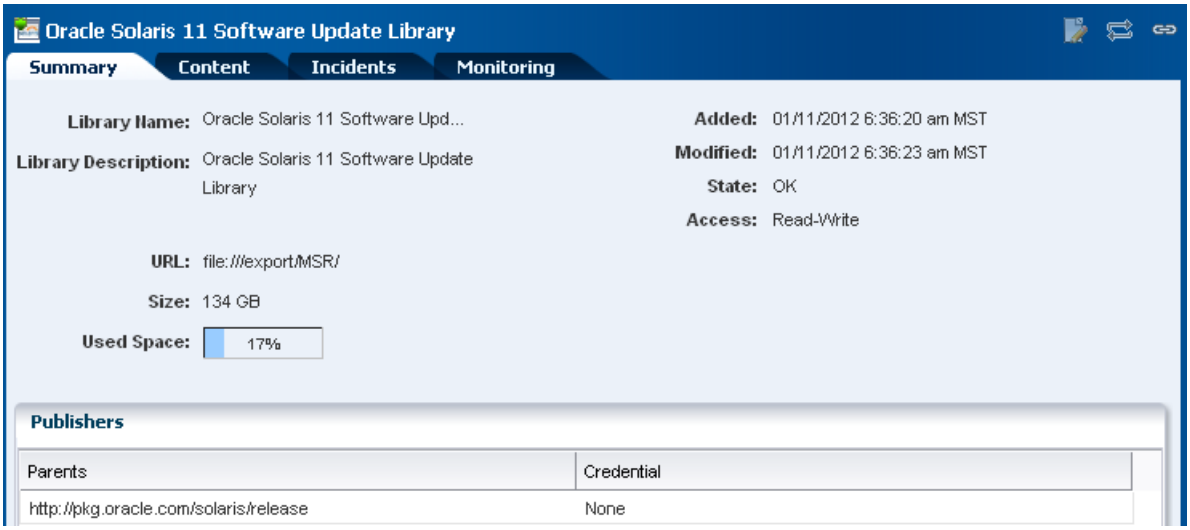

Content of the Oracle Solaris 11 Software Update Library

The Content tab, shown in Figure 5–3, displays a list of the packages in the Oracle Solaris 11 Software Update Library with a brief description of each package. You can filter the list of packages by selecting a category or version. The Support Repository Update check box is a quick way to filter the list of packages in the Oracle Solaris 11 SRU releases. You can also search for a specific package.

Figure 5–3 Oracle Solaris 11 Software Update Library Content

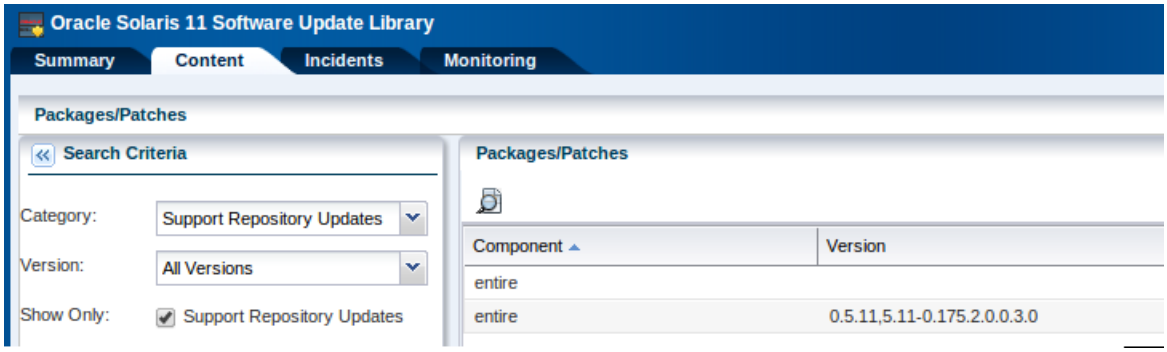

Configuring Parent Repositories to Synchronize

Use the **Configure Parent Repositories** action to manage publishers for the parent repository. After the library has been configured and the content has been downloaded the first time, create a recurring schedule to synchronize the Oracle Solaris 11 Software Update Library with its parent repository.

In the Configure Parent Repositories Wizard, create a list of parent repositories with their credentials that you use to update the Oracle Solaris 11 Software Update Library. During the synchronization operation, all the content in the parent repository is compared to the existing content in the Oracle Solaris 11 Update Software Library and new or updated content is downloaded.

Using an Alternate IPS Repository

Oracle Enterprise Manager Ops Center synchronizes with a parent repository to maintain the Oracle Solaris 11 Software Update Library, that is, the Oracle Solaris 11 Software Update Library is a child of the parent repository. If your site already has an Oracle Solaris 11 IPS repository for other purposes, you have the option of declaring that repository as the Oracle Solaris 11 Software Update Library. Operations that use IPS content retrieve the content from your IPS repository. Because it is not a parent-child relationship, no synchronization is performed.

When you install Oracle Enterprise Manager Ops Center or when you initialize the Oracle Solaris 11 Update Software Library, identify your IPS repository as the Oracle Solaris 11 Update Software Library. All operations that use this software library will retrieve content from the IPS repository.

Note: Oracle Enterprise Manager Ops Center does not maintain the content of the IPS repository. You must maintain this repository manually.

To Use an IPS Repository as the Oracle Solaris 11 Software Update Library

1. Locate your site's Oracle Solaris 11 IPS repository or create one.
2. Expand **Libraries** in the Navigation pane.
3. Click **Software Libraries**.
4. Click **Initialize Solaris 11 Software Update Library** in the Action pane.
5. In the **Specify Library Location** field, enter the location of your IPS repository.

6. In the **URL** field for the parent repository, do not enter any text. Leave this field blank.
7. Click **OK**.

Adding Content

Use the **Add Content** action to specify a parent repository and add ISO image files to the Oracle Solaris 11 Software Update Library manually. The parent repository URL and credentials are not saved when you use this action. To save a list of repositories, use the **Configure Parent Repository** action.

For each repository, enter its complete URL in HTTP protocol. If credentials are required, specify the system's credentials, **SystemDefinedSupportCredential**, or create new key and certificate identifiers to access the repository. When you select **Define New,** show in Figure 5–4, the Create Credential Wizard is displayed.

Figure 5–4 Add Content

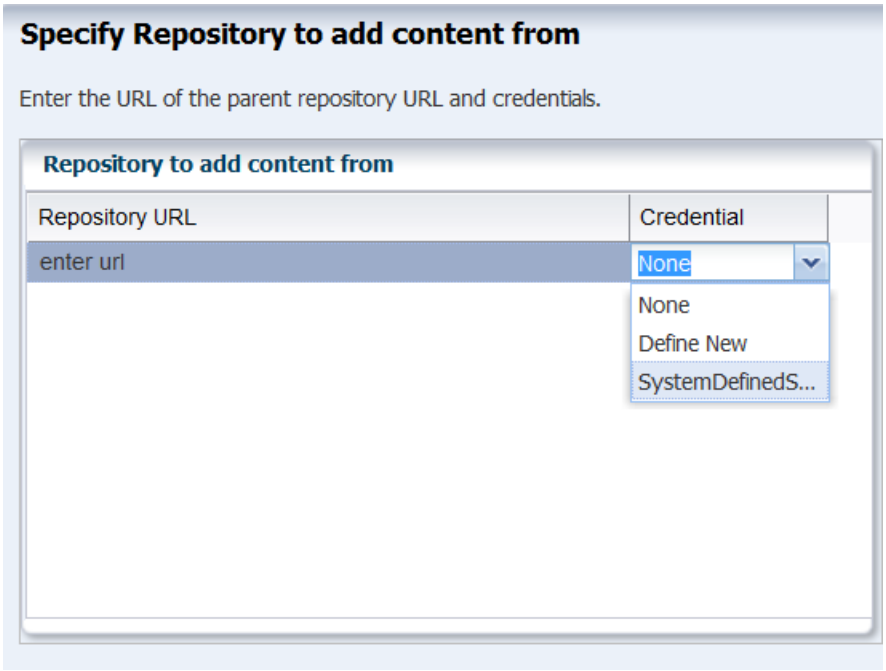

Deleting Oracle Solaris 11 Software Update Library

You can delete the library and keep its packages and content, or you can delete the library and all of its contents.

Libraries for Oracle Solaris 10, 9, 8 and Linux

Update Profiles and any profiles you create that provision an OS or firmware image rely on the contents of a software library. For Oracle Solaris 10, 9, and 8 and for Oracle Linux, a dedicated software library called Linux, Solaris 8-10 Software Update Library contains the packages, updates, and site-specific scripts and configuration files for these operating systems.

You create the Linux, Solaris 8-10 Software Update Library when you install Oracle Enterprise Manager Ops Center, or use the **Create Update Library** action to create the

library after installation. When the **Create Update Library** action is not available, the library has been created.

For each OS, this software library organizes its contents in the following categories:

- Clusters (groups of Oracle Solaris packages)
- Configuration files
- Hardware
- Local
- Local Packages
- Local RPMs
- Notifications
- Packages
- Patches
- Post-actions
- Pre-actions
- Probes
- Recommended Software Configurations
- Oracle Solaris Baselines

You can change the display of the contents of the library according to the OS distribution, the category, the type of view, and the version.

See Images for information about managing the content in the library.

See Local Content for information about managing the content of the local categories.

Images

In most cases, all the images you need are downloaded from the Oracle Knowledge Base or the Oracle Solaris 11 Package Repository. You can also create images within Oracle Enterprise Manager Ops Center or obtain them from a location external to Oracle Enterprise Manager Ops Center and then import or upload them.

An operating system can be in ISO format or FLAR format. Firmware can have various formats, depending on the vendor.

- ISO image – The image, also called a disk image, contains uncompressed directories and files of any type: application or data or both. This type of image can reside on removable media.
- FLAR image – The image is a flash archive of an Oracle Solaris 10, 9, or 8 operating system and other software. You can use a FLAR to install or restore a system to a specific configuration. Differential FLAR images are not supported. FLAR files must have the `.flar` file name extension.

> **Note:** Another type of image is a virtualization image. This type of image is stored in a Storage Library and contains the configuration information for a guest, its operating system, and the applications that the guest uses.

Images for OS Provisioning

An OS image contains an entire operating system in either ISO format or FLAR format. A subset of OS images are branded images, which install an operating system that is optimized for a specific purpose and can also include applications. OS images provision the operating system on both hardware servers and virtualization hosts. Table 5–3 shows the location of OS images.

Table 5–3 Location of Images for Each Operating System

Operating System	Location of Images
Oracle Solaris 11	Oracle Solaris 11 Software Update Library
Oracle Solaris 10, 9, 8	Linux and Oracle Solaris 8-10 Software Update Library
Oracle Linux	Linux and Oracle Solaris 8-10 Software Update Library

In addition to being grouped as packages and updates, Oracle Solaris OS images are also grouped into baselines.

Requirements for OS Images

- An OS image must be in a single image file. For example, on an Oracle Solaris system, the following command collects all OS component files on the auto-mounted file system into an ISO file.

  ```
  # mkisofs -o <name_of_OS.iso> -J -R /cdrom/<name_of_OS>
  ```

- The Oracle Enterprise Manager Ops Center software uploads or imports one ISO file per operation. If you are loading an ISO file from physical media and the file spans more than one CD, combine the content on one DVD.
- An ISO file cannot be made from Oracle Solaris installation CDs.

Images for Firmware Updates

Firmware images provision hardware assets. A firmware image is a copy of the vendor's firmware file and metadata for the firmware, such as the platform it is used on and any software dependencies. You obtain the images by downloading them from vendor websites or uploading them from their product media. Firmware images are stored in Software Libraries. The maximum size of a firmware image is 20 MB.

The following firmware types are supported:

- Service Processor firmware
- Chassis firmware
- Power distribution unit firmware
- Storage Component firmware updates firmware on RAID Controllers, Expanders and Disks.

When you import a firmware image, you might be required to provide metadata to complete the image file. You can usually find the information in the image's README file. You must provide the firmware type, the systems that the firmware supports, the version of the firmware, and any other firmware images that this firmware image depends on.

Example 5–1 Example of Firmware Metadata

The following is an example of a README file for ALOM-CMT firmware, where a single binary is deployed to the Service Processor.

- To determine the type and version of the firmware update:

```
Latest Sun System Firmware(6.1.2):
                ---------------------------------
                System Firmware 6.1.2 Sun Fire[TM] T2000 2006/01/20 18:19
                ALOM-CMT v1.1.2  Jan 20 2006  18:06:10
                VBSC 1.1.1  Jan 20 2006  17:56:19
                Reset V1.0.0
                Hypervisor 1.1.0 2005/12/15 11:10
                OBP 4.20.0 2005/12/15 16:48
                Sun Fire[TM] T2000 POST 4.20.0 2005/12/15 17:19
```

- To determine the models supported:

```
This README is intended for users who wish to
upgrade the firmware in their Sun Fire T2000.
```

- To determine if the system must be powered off before updating the firmware:

```
a)To update the Sun System Firmware, the system must be powered off (i.e. in
standby mode).
```

From this README file, identify the following metadata:

- Available platforms – Sun Fire T2000
- Type – VBSC
- Version – 1.1.1
- Require power off – Yes

For this example, the VBSC firmware subcomponent/type with version 1.1.1 was used. You can use any of the other types such as ALOM-CMT:1.1.2 or OBP:4.20.0. However, you must ensure that the version specified is always the firmware subcomponent/type.

Uploading or Importing Images

To provision firmware or an OS, you use a deployment plan to direct Oracle Enterprise Manager Ops Center to retrieve the images from the appropriate software library and install them on the targeted assets.

Starting in Release 12.2.2.0.0, you must create a profile for each image and then include the profile in a deployment plan. If you prefer to create profiles by default, you must change Enterprise Controller's configuration by editing the property file. See Default Profiles and Plans.

To use an image in a deployment plan, the image must be in one of the Oracle Enterprise Manager Ops Center's software libraries. You upload an image or import an image, depending on where the image resides. In both cases, you are moving the image from a location external to Oracle Enterprise Manager Ops Center's management into one of its libraries.

- If the image resides on the Enterprise Controller's system, import the image.
- If the image does not reside on the Enterprise Controller's system, upload the image. This operation relies on the browser to transfer the image file. The maximum size of an image that you can transfer using browser operations is 2 GB.

If the image is larger than 2 GB, move the file manually to the Enterprise Controller's system and then import it.

Importing an Image

1. Expand **Libraries** in the Navigation pane.
2. Click **Software Libraries** or **Storage Libraries** to expand.
3. Click the library.
4. Click **Import Image** in the Actions pane.

   The Import Image window is displayed.

Figure 5–5 Import Image

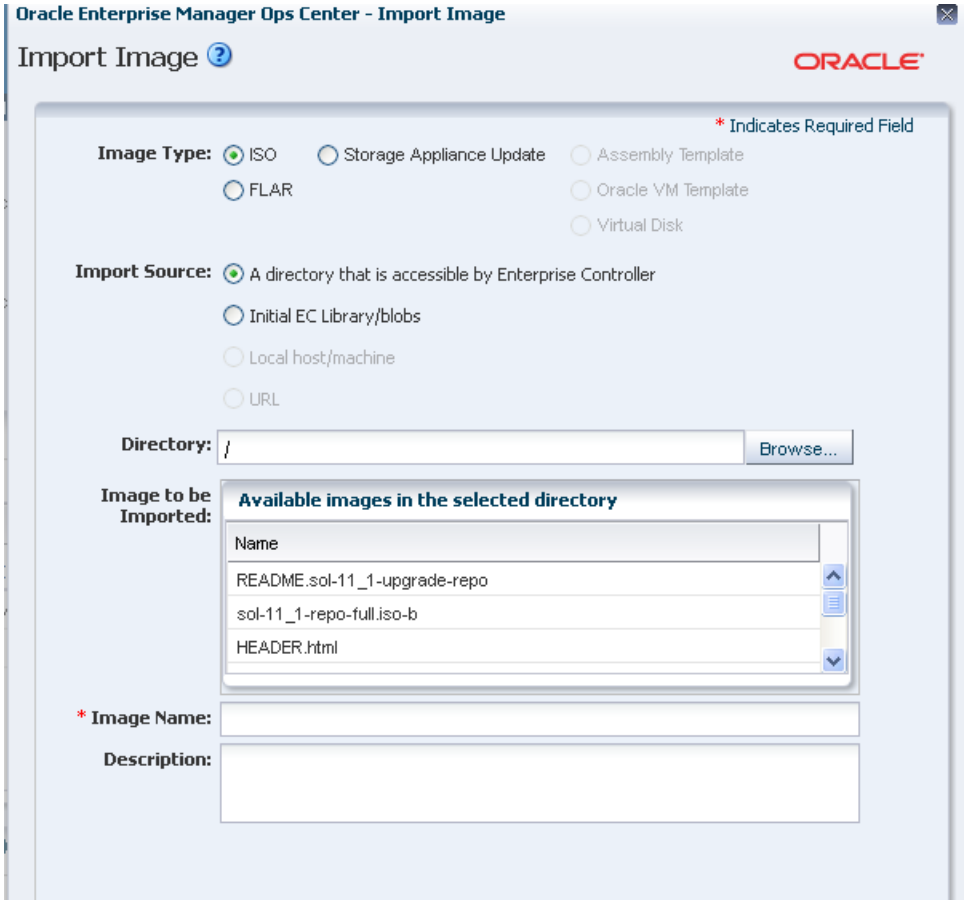

5. Identify the type of format for the image.
6. Identify the current location of the image file. For an ISO image file, the available source locations are the Initial EC Library or another location that is accessible to the Enterprise Controller.

   If the image file is located in an accessible location, enter the name of the directory or use the Browse button to navigate to the location.
7. Select the image you want to import.
8. Enter a name for the image and a description. Image names must be unique, can consist of up to 100 characters, and can include numbers, letters, and some special

symbols. The following special symbols are prohibited: comma, asterisk, single quote, double quote, parenthesis, question mark, equal sign, and newline.

9. Click **Import Image** to copy the image to the library.

Uploading an Image

Although the action name is Upload ISO Image, you can use the action to upload a FLAR file. However, the original OS image, the image that the FLAR is based on, must also be in the same library.

1. Expand **Libraries** from the Navigation pane.
2. In Software Libraries, select the library in which you want to store the image.
3. Click **Upload ISO Image** in the Actions pane.

   The Upload ISO Image window is displayed.

Figure 5–6 Upload ISO Image

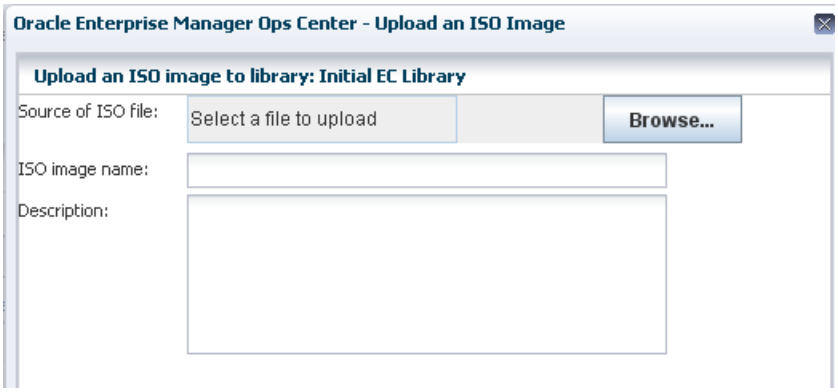

4. Click the name of the file in the Source of ISO File field or click Browse to navigate to the image.
5. Enter the name and description of the image.
6. Click **Upload Image**.

The progress of the upload to the Enterprise Controller is displayed. When the job is completed, the image is in the software library.

Uploading Firmware Images

The Upload Firmware action can be directed to retrieve firmware images from either the Enterprise Controller's file system or a local file system. Use the procedures in *Keeping Your Firmware Up-to-Date* at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm to upload firmware.

Working with Firmware for Power Distribution Units

The firmware for a PDU consists of two files, one for the firmware and one for the management software of the PDU. Each firmware image must be imported or uploaded and then a profile created that includes both images. The firmware images must have metadata, which is not always included in the image. In that case, you specify the metadata during the import or upload operation.

Uploading Firmware With Metadata

Use the procedures in *Keeping Your Firmware Up-to-Date* at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm to upload firmware. At the step for Select Firmware Components, select both firmware components:

- MKAPP_V*firmware_version*.DL
- HTML_V*firmware_version*.DL

Uploading Firmware Without Metadata

Use the procedures in *Keeping Your Firmware Up-to-Date* at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm to upload firmware. At the step for Select Firmware Components, select the firmware component, MKAPP_V*firmware_version*.DL, and click **Upload**. To define the metadata for the firmware:

1. For the target type click **PDU**.
2. For the platform, click **Sun Rack II PDU**.
3. Enter the version of the firmware.
4. Choose **Depends On None** and accept the remaining default options.

After the job is completed, repeat the procedure with the other firmware image, HTML_V*firmware_version*.DL, with one difference: At Step4, specify that this firmware image depends on the MKAPP_V*firmware_version*.DL image.

Creating a Firmware Profile for PDU Firmware Updates

1. Expand **Plan Management** in the Navigation pane.
2. Click **Create Firmware Profile** in the Actions pane.
3. Enter a name and description for the profile.
4. For subtype, click Power Distribution Units.
5. For target type, click Power Distribution Units. Click **Next**.
6. For Firmware Configuration, accept the default options and values. Click **Next**.
7. Select both images: `MKAPP_V<firmware_version>.DL` and `HTML_V<firmware_version>.DL`. Click **Next**.
8. Review the summary of the new profile and click **Finish** to submit the job.

Local Content

> **Note:** This section does not apply to Oracle Solaris 11. To add custom content for Oracle Solaris 11, use `pkg` commands to add the content to a local IPS repository and then import the repository into Oracle Enterprise Manager Ops Center. To use custom scripts, use an operational profile with either the Execute Operation action or an operational plan. See Operational Plans and Profiles.

The Linux and Oracle Solaris 8-10 Software Update Library also stores and gives access to site-specific configuration files and scripts used in deployment plans. This local content can also include data files, executable files, or binary files. For example, you might develop a script to test servers before running a provisioning job.

You can view details such as distribution, version, release, group, size URL, when the file was added or edited, any summary and description information, host, RPM, and vendor.

Local Categories

The Local categories of the Updates Library have no connection to the Knowledge Base. You upload local content to the software library, directed into one of the categories, and then maintain the files throughout their life cycle.

Local content is organized into the following default categories: local RPMs or PKGs, configuration files, macros, pre-actions, post-actions, and probes. You can create subcategories to further organize your local content. The type of local content allowed in a subcategory depends on its parent category.

Adding a Local Category

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library**.
3. Click **Add Local Category** in the Actions pane.
4. Enter a name for the new subcategory.
5. Enter a brief description for the new subcategory such as its purpose.
6. Click **Distribution** to assign to the subcategory.
7. Click **Parent Category** to select one of the system-defined categories for the subcategory.
8. Click **Apply**. The new subcategory is created under the selected default category.

You can now upload software packages and files into the new subcategory.

Uploading a Local Action

An action is a script, binary file, or executable file that makes changes to the managed host. To use the uploaded script or file, create an Update profile and include the profile in the deployment plan that installs or upgrades the OS. The following actions are available:

- Pre-Actions – Script that runs on a managed host before the provisioning step starts. When you create the deployment plan, you select the Execute Pre-Install step and then select the profile that includes the script.
- Post-Actions – Script that runs on a managed host after a job is completed. When you create the deployment plan, you select the Execute Post-Install step and then select the profile that includes the script.
- Probes – Script that runs on a managed host to verify that a job can be performed.
- Macros – Script that modifies a generic configuration file to make it specific for a managed host. Use macros to apply a single configuration file across multiple hosts by customizing the configuration file for each host's environment. The script outputs a single line that replaces a macro sign in a configuration file.

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library.**
3. Click **Upload Local Action** in the Actions pane.

The Upload Local Action window is displayed in Figure 5–7.

Figure 5–7 Local Upload Action

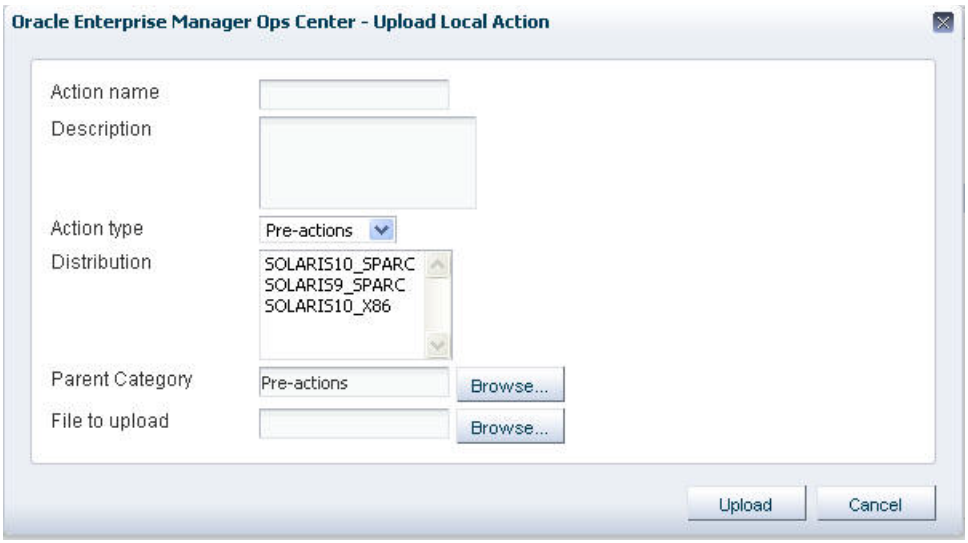

4. Enter a name for the action.
5. Enter text to describe the purpose of the script or executable file.
6. Select the type of action such as Pre-Action, Post-action, Macros, or Probes.
7. Click the name of the distribution that is appropriate for the script or executable file.
8. The Parent Category field shows the category you specified as the type of Action. If your site uses subcategories, click the **Browse** button to navigate to the specific subcategory.
9. Click **Browse** to locate and select the script or executable file.
10. Click **Upload**. The file is uploaded to the Linux, Solaris 8-10 Software Update Library in the category you specified.

You can now create a profile for the script and then use the profile in a deployment plan.

Uploading a Local Software Package

You can upload software in the following formats:

- `pkg`
- `rpm`
- `tar`
- `zip`
- `gzip`
- `compress`

If the file is in compressed format, the file is uncompressed after it is uploaded.

1. Expand **Libraries** in the Navigation pane.
2. Click the operating system's Software Update Library.

3. Click **Upload Local Software Packages** in the Actions pane.
4. Select **Yes** if the package is a security fix for a previous version of the software. Otherwise, select **No**.
5. Click the name of the distribution to which you want to add this package.
6. In the Parent Category section, click **Local PKGs** or click **Browse** to locate a subcategory.
7. In the Files section, click **Add** to see the list of files. Select at least one software package.
8. Click **Upload**. The file is uploaded to the Linux, Solaris 8-10 Software Update Library in the category you specified.

You can now include the package or RPM in an provisioning profile.

Uploading a Local Configuration File

A configuration file is a text file, binary file, or non-RPM application that contains the settings and values for an asset type. To use the uploaded file, create an Update profile and include the profile in the deployment plan that installs or upgrades the OS.

Uploading a Local Configuration File

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library.**
3. Click **Upload Local Configuration File** in the Actions pane.

   The Upload Local Configuration File window is displayed in Figure 5–8.

Figure 5–8 Upload Local Configuration File

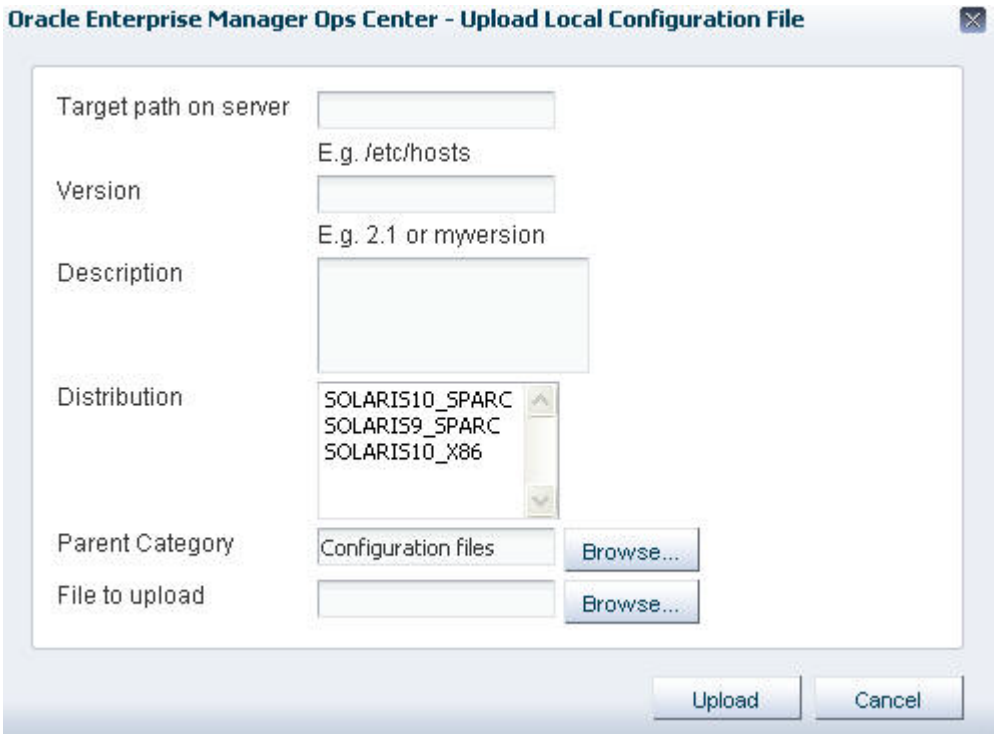

4. In Target path on server, type the full path to the configuration file.

5. In Version, type a character string to identify this version of the file. The string is appended to the file name when it is displayed in a Components list.
6. Enter a brief description of the file.
7. Select the Distribution to which this file is applied. You can choose multiple distributions.
8. In Parent Category, accept the Configuration Files category or click **Browse** to locate a subcategory.
9. Click **Browse** to locate and select the configuration file.
10. Click **Upload**. The file is uploaded to the Linux, Solaris 8-10 Software Update Library in the category you specified.

Uploading Software in Bulk

You can upload multiple files or an entire directory in one operation. For example, you can upload the contents of a DVD. All components in the directory and subdirectories are uploaded.

The files must be in the following formats:

- `pkg`
- `rpm` (for Linux RPMs)
- `tar`
- `zip`
- `gzip`
- `compress`

If files are compressed, the software extracts the files after it uploads them.

Before You Begin

- Verify that the files have the supported file types.
- Verify that the file size does not exceed 2 GB. If the file is larger than 2 GB, copy the file manually to a file system on the Enterprise Controller's system.
- If you are uploading from removable media, insert the media.

Uploading Local Software in Bulk

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library**.
3. Click **Bulk Upload Packages and Patches** in the Actions pane.

   The Upload Packages, Patches, and RPMS window is displayed in Figure 5–9.

Figure 5–9 Upload Local Software

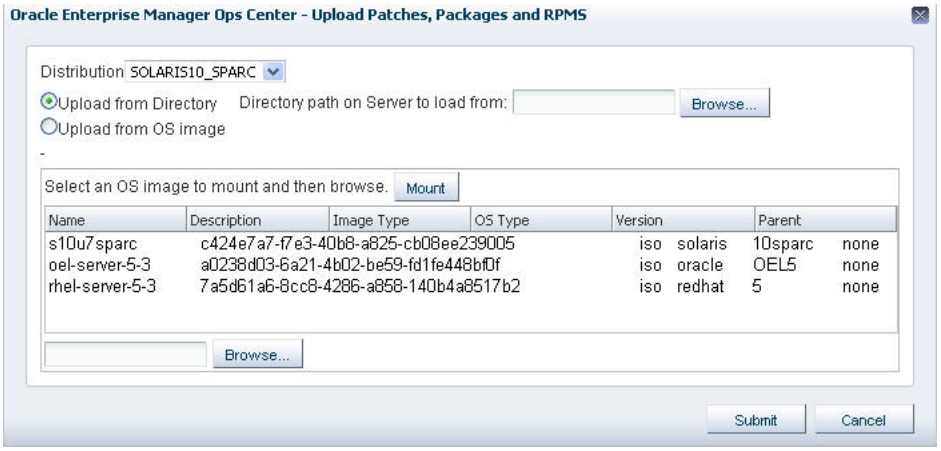

4. Click **Distribution** to select the distribution that applies to these files.
5. Select either **Upload from OS Image** or **Upload from Directory**.
6. Specify the path to the OS image or directory or click **Browse** to locate and select it. If you specify a directory, all files in the directory and its subdirectories are uploaded. If you specify an OS image, you must mount the OS image and select the files.
   a. Click one of the OS images and click **Mount**.
   b. Click **Browse** to locate and select the files.
7. Click **Submit**. The upload job is created.

To view the status of the upload job, select **Bulk Upload Results**.

To view the certified packages in the software update library, click the Content tab in the center pane and select Patches in the Category list. To view non-certified packages, click Local PKGs or Local RPMs in the Category list.

Viewing Results of a Bulk Upload Operation

You can view a detailed history of all the local components that were uploaded in bulk.

To View Bulk Upload Results

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library**.
3. Click **Bulk Upload Results** in the Actions pane. The uploaded components list displays the name, description, status, and date for each component.
4. Select a component and click **View Results**. The details of the uploaded components are displayed.

Using Local Content

The local content that you uploaded to the software library is used in profiles and deployment plans to guide and complete provisioning operations. For example, a local configuration file can be modified by the output of a macro so that each target gets the appropriate configuration file.

- Create the macro named `mymacro` and upload it to the software library.

```
# cat /var/tmp/runme.sh
#!/bin/bash
hostname
```

- Create the local configuration file and upload it.

```
### This configuration file changes
### the following line to the output of "mymacro"
### Include this configuration file in the profile.
<^AM^>mymacro<^AM^>
```

- Create an Update profile and deployment plan to include the local configuration file.
- Deploy the plan. The resulting configuration file on a particular server, for example `mymachine1`, has the following content:

```
### This configuration file changes
### the following line to the output of "mymacro"
### Include this configuration file in the profile.
mymachine1
```

Editing Local Content

You can edit files in the Local Categories. For example, if you uploaded a file that contained IP addresses and determined that there was an incorrect IP address in the file, you can edit the file to correct the IP address. You can also use this procedure to replace the file with a corrected file.

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library**.
3. Click **Edit Local Component File** in the Actions pane.
4. To specify the file, type its name or click the **Browse** button to navigate to the file. If the file is not found, click **Distribution** to select the correct distribution. Only files in the selected distribution are displayed.
5. Select either **Edit existing file** or **Replace existing file**.
   - If you choose to edit the file, make changes to the file and click **Save**.
   - If you choose to replace the file, browse for the replacement file and click **Upload**.

Deleting Local Content

You can remove your site's local content or added categories. You cannot remove the default categories.

> **Note:** Deleting content does not require confirmation and cannot be undone. Verify you are deleting the correct local component.

1. Expand **Libraries** in the Navigation pane.
2. Click **Linux and Oracle Solaris 8-10 Software Update Library.**
3. Click **Delete Local Component** in the Actions pane.

4. Expand the category to display the component you want to delete. To change the distribution, click **Distribution**.
5. Select the file or a subcategory to delete.
6. Click **Delete**.

To remove a subcategory and its components, do not attempt to remove each component and then remove the subcategory. When there are no components in a subcategory, the subcategory creates a placeholder component, which you cannot delete. Repeat the procedure and select the subcategory itself to delete. The placeholder component is also removed.

Backing Up Images and Local Content

The `ecadm backup` command does not back up the software libraries. As a good practice, create the software library for OS images on networked storage (NAS) and include the network storage device in your site's backup plan. As an alternative, back up the Enterprise Controller's directory manually.

1. Move the archive to another server, file-share facility, or a location outside of the `/var/opt/sun` directory, according to your site's disaster recovery plan.
2. If it is necessary to rebuild the Enterprise Controller, restore the Enterprise Controller and then restore the `/var/opt/sun/xvm/images/os` hierarchy.

Related Resources for Software Libraries

For instructions in performing actions or to learn more about the role of this feature, see the following resources:

- In the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm:
  - *Keeping Your Firmware Up-to-Date*
  - *Update Oracle Solaris 10 OS Workflow*
  - *Update Oracle Solaris 11 Workflow*
- *Deploy Software Libraries Workflow* in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm

6

Storage

The following information is included:

Introduction to Storage

Oracle Enterprise Manager Ops Center discover, manages, and monitors storage servers and appliances, discovers and provisions storage capacity through Storage Connect plug-in software, and makes storage resources available to virtual assets through storage libraries.

Storage libraries are the storage resources for Oracle Solaris Zones, Oracle VM Servers for SPARC, Oracle VM Servers for x86, their server pools, and virtual datacenters. These virtualization hosts and server pools store their metadata and their operational data in storage libraries. The metadata and data for Oracle Solaris zones are stored in a storage library associated with the global zone or server pool. A storage library that supports an Oracle VM Server for x86 is an Oracle VM Storage Repository.

Roles for Storage

Table 6–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 6–1 Storage Tasks and Roles

Task	Role
Create a new storage library	Storage Admin
Remove a storage library	Storage Admin
Edit attributes of a storage library	Storage Admin
Add storage capacity	Storage Admin
Update a storage appliance	Update Admin

Actions for Storage

You can perform the following actions, depending on the requirements.

- Create new storage libraries, either as a file system or as block storage.
- Modify existing storage libraries. You can change the library's attributes or add capacity.
- Update storage appliances.

The set of available actions depends on what you have selected:

- When you select a physical asset, you can launch the asset's user interface and view information about the hardware's state and configuration.
- When you choose an asset in the File Server group, the Filesystems tab in the center pane lists all of the file systems with the **Add a Backing Device** icon, the **Edit** icon, and the **Delete** icon.
- When you choose an asset in the Storage Array Group, the Logical Units tab gives you access to the **Create Logical Unit** icon, the **Resize LUN** icon, the **Delete LUN** icon, and the **Clone LUN** icon.

Location of Storage Information in the User Interface

Table 6–2 shows where to find information.

Table 6–2 Location of Library Information in the BUI

Object	Location
To see the storage libraries	Expand Libraries in the Navigation pane, then expand Storage Libraries.
To see storage hardware information	Expand All Assets in the Navigation pane, then scroll to the Storage section. Select a physical asset to see information about it in the center pane.
To see groups of storage arrays or file servers	Expand All Assets in the Navigation pane, then Storage in the Resource Management Views section.
To see storage profiles	Expand Plan Management in the Navigation pane, then click Profiles and Policies. The Discovery, RAID Controller, and Update Storage Appliance categories contain profiles.

Table 6–2 (Cont.) Location of Library Information in the BUI

Object	Location
To see the virtual host that is using a LUN	Expand Libraries in the Navigation pane, then expand Storage Libraries. Select the storage library to view the LUNs table in the center pane. The Allocated To field displays the virtual host's identifier.
To see the virtualization host that is using a storage library	Expand Libraries in the Navigation pane, then expand Storage Libraries. Select the storage library to view the Summary tab in the center pane. The IP address of the virtualization host is in the Allocated To field.
To see incidents for a storage library	Expand the Message Center in the Navigation pane. For more information about recovering from incidents, see the *Oracle Enterprise Manager Feature Reference Appendix Guide*.

Storage Libraries

- Network Attached Storage (NAS) Storage Libraries
- Storage Libraries for Server Pools
- Storage Libraries for a Virtual Datacenter
- Storage Libraries for Oracle Solaris Zones
- Storage Libraries for Oracle VM Server for SPARC
- Storage Libraries and Repositories for Oracle VM Server for x86

A storage library stores metadata for each virtualization host in the server pool that is associated with the storage library. Metadata is a virtualization host's image or identity: the configuration for its operating system, CPU, memory, and network. The virtualization host's data, which results from its use, can reside in the same storage library or in a different storage library. A storage library can be a local, that is, a file system on the virtualization host's server, or it can be accessed through an NFS server or SAN network.

When you create a virtual host, you assign it to one of the storage libraries associated with its virtualization host so that its metadata can be stored.

- For the metadata of its local virtual hosts, virtualization hosts (Oracle VM Server for SPARC, Oracle VM Server for x86, and global zones) can use a local storage library. However, storing metadata in a local storage library limits the management of the virtual host because this virtual host cannot be migrated to another server and cannot be recovered on another server if it fails.
- For the metadata of all virtual hosts, virtualization hosts must use a Network Attached Storage (NAS) storage library.
- For data, Oracle VM Server for SPARC and Oracle VM Server for x86 can use either NAS shares or SAN LUNs for itself and for its virtual hosts. A global zone can use SAN LUNs for itself and for its non-global zones. The storage library must be associated with the virtualization host.

The diagram in Figure 6–1 shows how the NAS and SAN storage libraries and virtualization hosts interact with LUNs.

- SAN storage libraries expose data as virtual disks (LUNs), identified by their LUN GUIDs. A control domain makes raw partitions available to its logical domains using Fibre Channel or iSCSI, as described in Block Storage

- NAS storage libraries expose data as raw files and files in VMDK format. All types of virtualization hosts store metadata using NFS services, as described in Network Attached Storage (NAS) Storage Libraries

Figure 6–1 SAN and NAS Storage Libraries

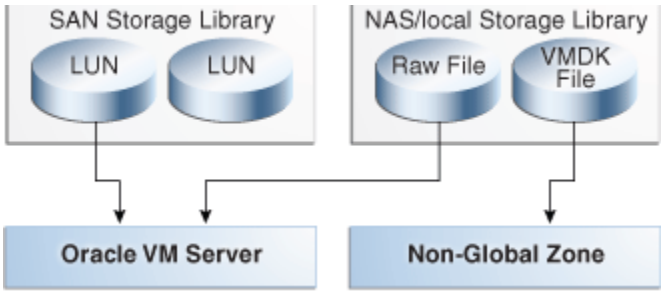

The LUNs, raw files, and raw volumes store data and metadata for the virtual hosts. Figure 6–1 expands Figure 6–1 to show that the metadata for NAS virtual disks are stored in the NAS storage library. Metadata for SAN virtual disks is persisted in the SAN Storage Library.

Figure 6–2 Storage Library Metadata

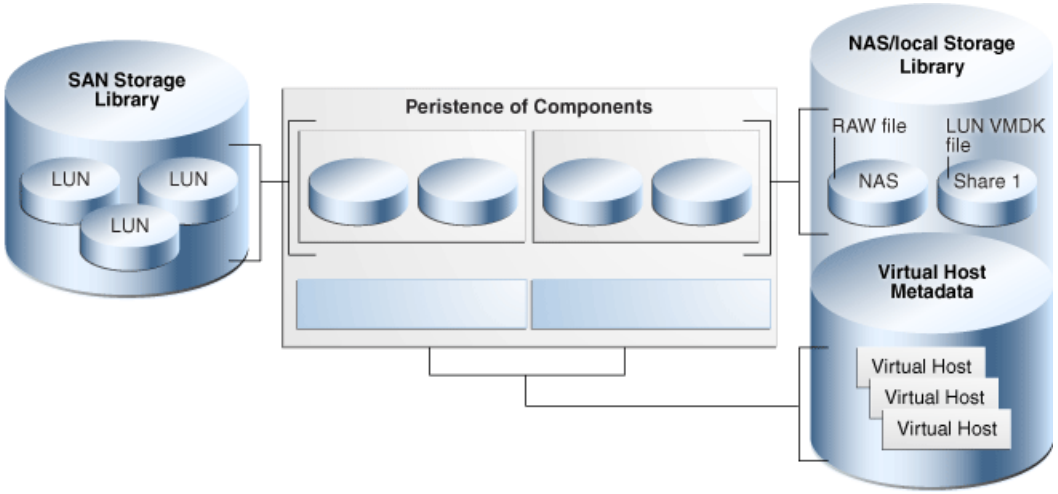

Network Attached Storage (NAS) Storage Libraries

Network-attached-storage (NAS) libraries are storage libraries for NFS storage device mount points. The virtualization hosts use NFS services to attach to the storage libraries and get access to the data and metadata.

You can store metadata for all virtual hosts in one NAS storage library or you can create separate storage libraries for each virtual host. Use separate storage libraries to increase ease of access, to increase capacity, and to increase performance.

If a NAS storage library becomes unavailable, the virtual hosts associated with the library are affected in the following ways:

- If the storage library is used for the virtual hosts' metadata, a virtual host continues to function but Oracle Enterprise Manager Ops Center can no longer manage the virtual host. Because Oracle Enterprise Manager Ops Center relies on

its interaction with the metadata in the storage library, jobs that must read or modify the metadata fail. You must manage the virtual host manually.

- If the storage library is used for NFS large files that support virtual disks, the virtual host does not function.
- If the boot disk is on the NFS share, the virtual host cannot be rebooted.
- The virtual host cannot be migrated.

Storage Libraries for Server Pools

You group virtualization hosts to create a server pool. The virtualization hosts share all the storage and networks associated with the server pool. When you add a virtualization host to a server pool, the libraries associated with that virtualization host become available to all the other hosts in the server pool.

To delete a storage library or the storage hardware asset that supports the storage library, you must remove its associations. Use the **Disassociate Library** icon in the Libraries tab to disassociate the library from the server pool.

A server pool must use a NAS storage library. Figure 6–3 shows how virtualization hosts in a server pool get access to storage resources through the storage libraries.

Figure 6–3 External Storage

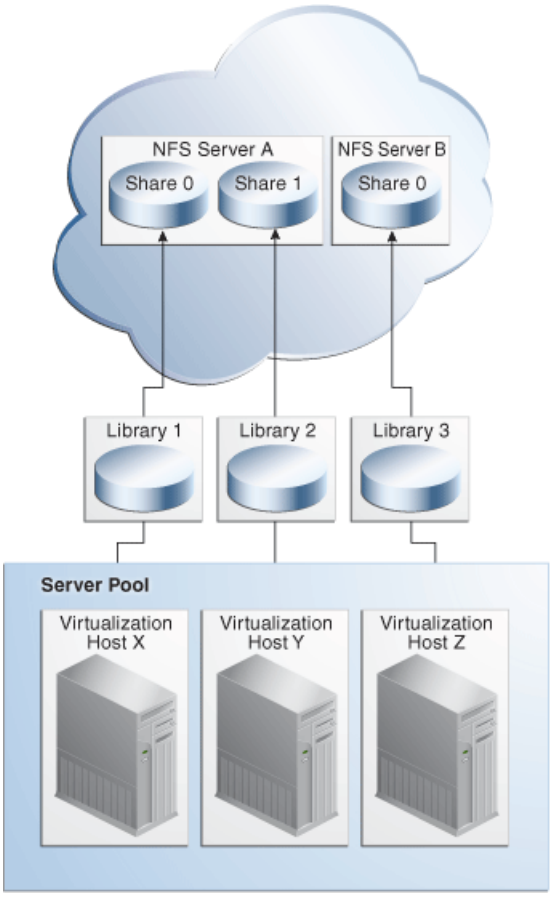

Storage Libraries for a Virtual Datacenter

The vDC inherits the storage resources allocated for the server pool. See Setting Up Storage Resources in Chapter 22.

For each account, the amount of storage used by the guests is shown on the Account's Dashboard tab. In the Resource Allocation section of the dashboard, the Storage graph shows the number of guests and the used space in gigabytes. The scale for the x-axis of the graph is adjusted according to amount of space allocated to guests. In Figure 6–4, one guest is allocated 14 GB of space so Oracle Enterprise Manager Ops Center uses 14 as the midpoint of the x-axis.

Figure 6–4 Graph of Guest Storage Resource Allocation

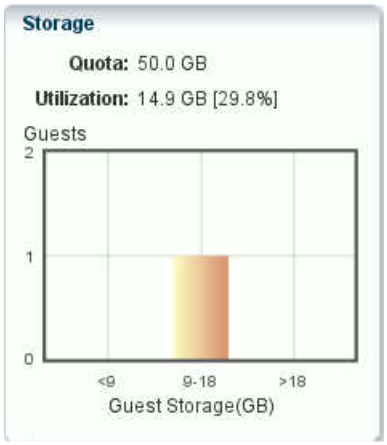

Storage Libraries for Oracle Solaris Zones

A global zone provides storage resources to its non-global zones. The Oracle Solaris Zone must be associated with a storage library. See Associating a Storage Library with a Global Zone in Chapter 18.

Storage Libraries for Oracle VM Server for SPARC

The control domain provides storage resources to its logical domain. The control domain must be associated with a storage library. See Associating Storage Library with the Domains in Chapter 19.

Storage Libraries and Repositories for Oracle VM Server for x86

Oracle Enterprise Manager Ops Center provides option to create Oracle VM Storage Repositories. This type of storage library stores virtual machine metadata, templates, assemblies, ISO images, and virtual disks for the Oracle VM Server for x86.

When you create the storage repository on a LUN, it is a block-based repository. When you create the storage repository on a NFS file server, it is a NFS-based storage repository.

Creating an Oracle VM Storage Repository

1. Expand **Libraries** in the Navigation pane.
2. Select **Storage Libraries**.
3. Click **New Oracle VM Storage Repository**.

4. Enter a name and description for the repository and choose either the NFS or OCFS protocol. Click **Next.**
5. From the drop-down lists, choose the Oracle VM Manager and the Oracle VM Server for x86 that use the storage repository, and the NFS File Server that supports the storage repository.
6. Choose the file system on the NFS file server for the storage repository and specify the share name. Click **Next**.
7. Choose the server pool to associate with the storage repository. Click **Next**.
8. Review the configuration and then click **Finish** to submit the job.

Uploading Templates and Assemblies to an Oracle VM Storage Repository

1. Expand **Libraries** in the Navigation pane.
2. Select **Storage Libraries**.
3. Select the Oracle VM Storage Repository.
4. Click **Import Image** in the Actions pane.

   The Import Image window is displayed.

Figure 6–5 Import Image

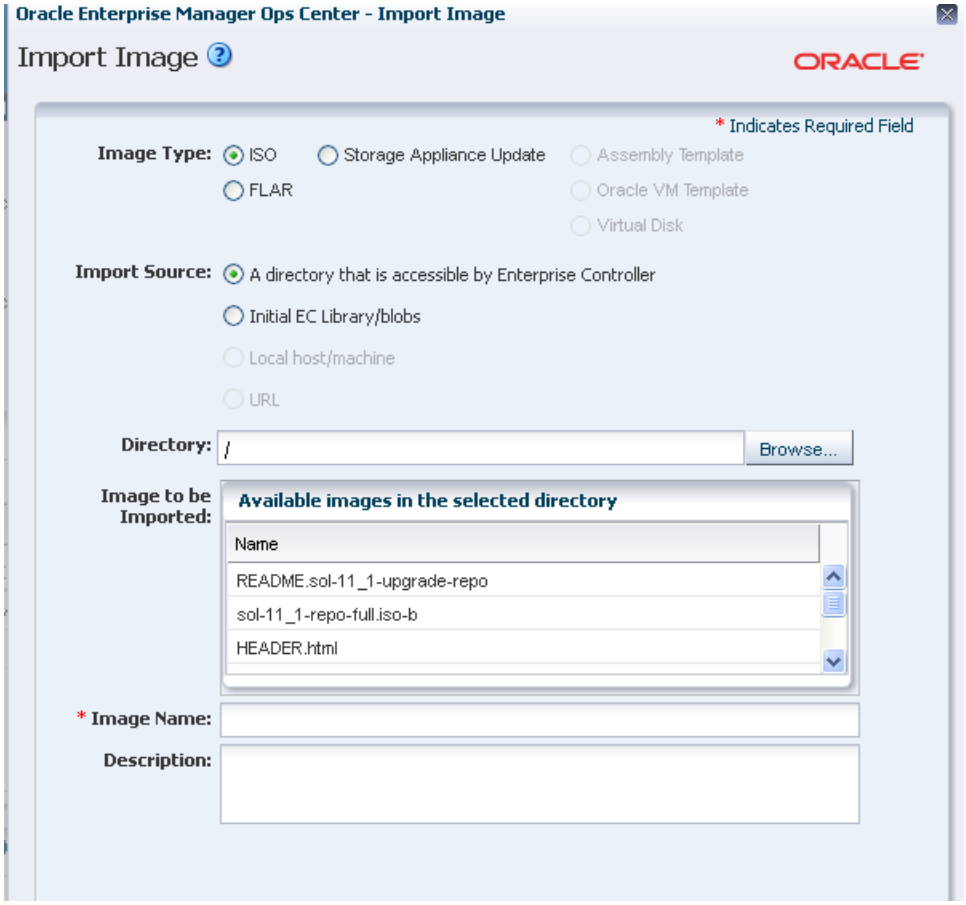

5. Identify the type of format for the image, either Assembly Template or Oracle VM Template.

6. Identify the current location of the image file. If the image file is located in an accessible location, enter the name of the directory or use the **Browse** button to navigate to the location.
7. Select the image you want to import.
8. Enter a name for the image and a description. Image names must be unique, can consist of up to 100 characters, and can include numbers, letters, and some special symbols. The following special symbols are prohibited: comma, asterisk, single quote, double quote, parenthesis, question mark, equal sign, and newline.
9. Click **Import Image** to copy the image to the library.

Types of Storage for Libraries

File Systems Libraries

Each virtualization host has a default local library named `/guests` where data and metadata for each virtual host is stored. For the purposes of storage efficiency and your site's organization, you can create and maintain other local libraries.

If the storage library becomes unavailable, the local library remains available. However, any guest with metadata in a local library cannot be migrated.

Viewing Local Libraries

Use this procedure to see the local libraries for a virtual host and the contents of a library. You can also see details of the local disks that support the local libraries.

1. Expand **Assets** in the Navigation pane.
2. Select the virtual host.
3. Click the **Libraries** tab in the center pane. The Associated Libraries table's Type column identifies the libraries of the Local type.
4. Select a library of the Local type. The Usage table shows all the guests that use that local library.
5. In the Usage table, select a guest.
6. Click the **Contents** tab to see the Library Contents table with all of the images, sorted by type.
7. To see details of the local disks, return to the Associated Libraries table and click **Local Devices**. Then select the local device library.

When you add new disks, use the **Refresh** icon to include them in the table of disks.

Editing the Attributes of a Local Library

You can rename a local library and you can change its description. You cannot change the file system defined for the local library.

1. Expand **Assets** in the Navigation pane.
2. Select the asset.
3. Click the **Libraries** tab in the center pane. The asset's associated libraries and the guests that are stored in the libraries are listed.

4. Click the **Edit Local Library** icon.
5. In the Edit Local Library pane, enter the new name or description for the library.
6. Click the **Update** button. When the job is completed, the edited local library is listed in the Associated Libraries table.

Creating a Local Library

Each virtual host has a local library, located at `file: ///guests`. In addition to the default local library, you can create other local libraries to use your storage resources efficiently or organize your images.

> **Note:** To use the Boot Environment feature of Oracle Solaris 11, the local library must be located in its own ZFS file system.

1. Create a file system with read/write permissions for only the root user.
2. Expand **Assets** in the Navigation pane.
3. Select the asset.
4. Click the **Libraries** tab in the center pane. The asset's associated libraries and the guests that are stored in the libraries are listed.
5. Click the **New Local Library** icon.
6. In the Create Local Library pane, type a name and description for the library.
7. In the URL field, enter the directory name for the location where you want to store images and metadata.
8. Click **Create Local Library**. When the job is completed, the new local library is listed in the Associated Libraries table.

Deleting a Local Library

You can delete a local library that was added to a virtual host. After the deletion, the virtual host does not have any access to either the directory defined for the local library or any of its contents. The default local library, `/guests`, cannot be deleted. You have the option of deleting the library and all of its contents or to delete the library but keep the contents in the directory of the storage resource. In either case, the directory is not deleted.

When the job is completed, the local library is removed from the Associated Libraries table. If you chose to keep the local library's content, you can create a local library with the same content by specifying the same URL for the new library.

Block Storage

Block storage libraries are used in SAN networks and define storage by Logical Units or LUNs, which are backed by either Fibre Channel disks or iSCSI disks. You can associate block storage libraries with server pools, Oracle VM Servers, or global zones to store their data.

A LUN (Logical Unit Number) is a slice of a storage volume, as defined by the following terms:

- Disk: Physical storage media. A set of disks is a disk array.
- Volume: An aggregation of storage space provided by several disks.

- Slice: A partition of a volume that is exposed to the servers connected to the disk array.
- LUN (Logical Unit Number): The representation of a slice.
- GUID: The Global Unique Identifier for a LUN.

If the storage array is a managed asset, the LUNs can provide dynamic storage. Table 6–3 compares static block storage libraries and dynamic block storage libraries.

Dynamic Block Storage

When Oracle Enterprise Manager Ops Center can create, resize, delete, and clone LUNs, the block storage library is dynamic. When you add more virtual hosts to a server pool, you create LUNs in the storage library in the same action. When you discover a storage device with LUNs backed by iSCSI disks, a dynamic block storage library is created.

Static Block Storage

When Oracle Enterprise Manager Ops Center manages a storage device with existing LUNs, it can use the LUNs but cannot change or delete them, or create new LUNs. These operations must be done through the storage device's user interface. A static block storage library supports storage devices that were discovered and configured in previous product software versions. Each LUN is for the exclusive use of its assigned virtual host. Because a LUN has a fixed size and cannot be shared, you must plan how to optimize the available storage in the storage library and assign a LUN of the appropriate size.

Table 6–3 Comparison of Static and Dynamic Libraries

	Dynamic Block Storage Library	Static Block Storage Library
Protocol	iSCSI	iSCSI Fibre Channel
To create a storage library	The storage library is created implicitly when you discover storage hardware and its target groups.	You create LUNs and then create the storage library. You assign existing LUNs to the library.
To manage a storage library	When you add a virtual host, you can create a LUN for it. When a virtual host needs more storage, you can increase the size of the LUN. Any changes made through the storage server's user interface are reflected in the Oracle Enterprise Manager Ops Center's user interface.	The number of LUNs in the storage library determines the number of virtual hosts that can use the storage library. Changes made through the storage server's user interface are not updated in the Oracle Enterprise Manager Ops Center's user interface. Some information, such as a LUN's GUID can be obtained only from the storage server's user interface.
Add LUN icon	Create a new LUN.	Select an existing LUN or enter the GUID of an existing LUN.
Edit LUN Details icon	Increase the size of the LUN.	Renames the LUN in Oracle Enterprise Manager Ops Center
Delete LUN icon	Deletes the LUN and deletes the data. You are deleting the LUN on the storage server.	Deletes the LUN from the library but does not delete data.

Selecting LUNs For the Block Storage Library

Use this procedure in the following situations:

- You are creating a new SAN storage library and have associated the new library with a virtualization host or server pool immediately. In this case, you have selected the **New SAN Storage Library** action, you have selected the option to associate the library with virtualization hosts, and you have selected one or more virtualization hosts.
- You are increasing the storage capacity of an existing SAN storage library and chose to select LUNs from the available LUNs.

The Available LUNs table, shown in Figure 6–6 lists all of the LUNs that the selected virtualization hosts can access. For each LUN, the following information is displayed to help you identify the ones you want to add to the library.

- **LUN GUID** – The unique 32-digit identifier for the LUN.
- Host information for the LUN:
  - **Hostname** – Name or IP address of the host that can access the LUN.
  - **Controller Number** – The host's identifier for the HBA port, which is the physical interface to the Fibre Channel disk array.
  - **LUN Number** – The host's identifier for the LUN.
- **Size (GB)** – Size of each LUN in Gigabytes.

Figure 6–6 List of Available LUNs

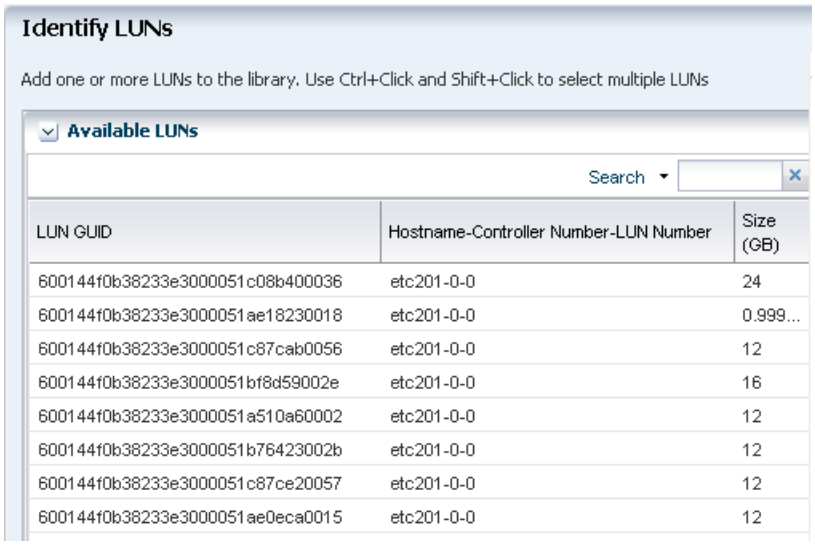

Adding Capacity to Dynamic Block Storage Libraries

If the block storage library is not associated with a server pool or virtualization host, you add storage capacity by specifying new LUNs by name. If the storage library is associated, you can also select new LUNs from a list of the LUNs available to the server pool or virtual host.

1. Expand **Libraries** in the Navigation pane.
2. Click **Block Storage** in Storage Libraries.
3. Click one of the available storage libraries.
4. Click **Add LUN** in the Actions pane.
5. Choose the method for adding LUNs:

- To add LUNs to the library manually, accept the default option: **Manually enter the GUID/WWN of the LUNs to be added**. See To Add LUNs to a Block Storage Library Manually to complete the procedure.
- To select LUNs from the available LUNs, click the **Select from available LUNs** option. See Selecting LUNs For the Block Storage Library to complete the procedure.

To Add LUNs to a Block Storage Library Manually

You use this procedure in the following situations:

- You are creating a new SAN storage library and have accepted the default action of adding LUNs later. You must add at least one LUN to create the storage library.
- You are adding a LUN to an existing SAN storage library. You selected the library and then the **Add LUN** action. The default option is to specify each new LUN by name.

The table shown in Figure 6–7 is displayed:

1. Click in the GUID/WWN field and type the GUID or WWN for the LUN. The GUID is the Global Unique Identifier associated with each LUN, which is a hexadecimal number of 32 digits. If your site uses SCSI initiators and targets, you can enter the WWN for the LUN.

Figure 6–7 Specifying a LUN

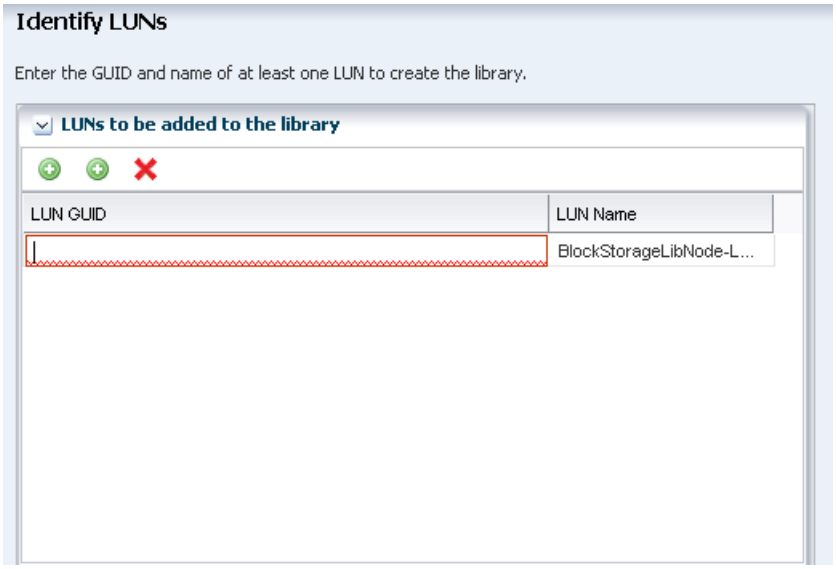

2. Click the first **Add** icon to include additional LUNs.
3. If you do not have the GUID or WWN, select a LUN from the list of available LUNs. Click the second **Add** icon to see a list of managed storage assets as shown in Figure 6–8.

Figure 6–8 Specifying an Exported LUN

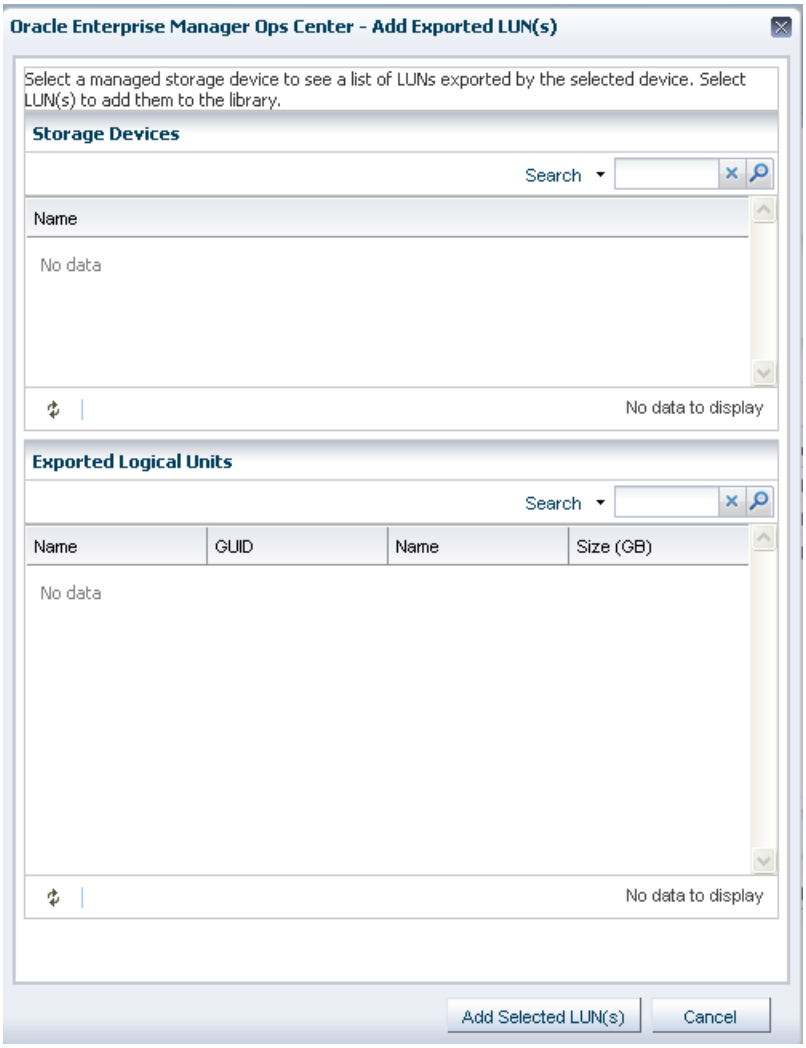

4. Select one of the storage assets to populate the list of exported LUNs on that device.
5. Select the LUNs to use in the library. Use the Search box to locate a specific LUN.
6. Click **Add Selected LUNs**.

   If you are adding a LUN to an existing storage library, a new job starts. If you are creating a SAN storage library, review the details of the LUNs you have configured in the Summary pane and click **Finish.**

Adding Storage to a Logical Domain

When a logical domain is created, it is assigned storage resources for its metadata and operational data.You can add storage for the use of the logical domain by selecting additional virtual disks from the storage library.

1. In the Asset section of the Navigation pane, select the logical domain.
2. In the center pane, click the **Storage** tab.
3. In the table of virtual disks, click the **Add Storage** icon.

The Add Storage wizard opens. The list of libraries includes all storage libraries that are available to the logical domain.

Figure 6–9 Adding LUNs

4. Select the library.
5. Select the LUN or virtual disk.
6. You have the recommended option of enabling multipathing for the storage so that an alternate path to the storage is available. Click in the Multipathing Group field and enter the name of the LUN's group.
7. Click **OK.**

Creating a LUN

In a dynamic block storage library, you can create a new LUN.

1. In the Asset section of the Navigation pane, select an asset in the Storage Array Group.
2. In the center pane, click the **Logical Units** tab.
3. Click the **Create Logical Unit** icon.

Figure 6–10 Creating a LUN

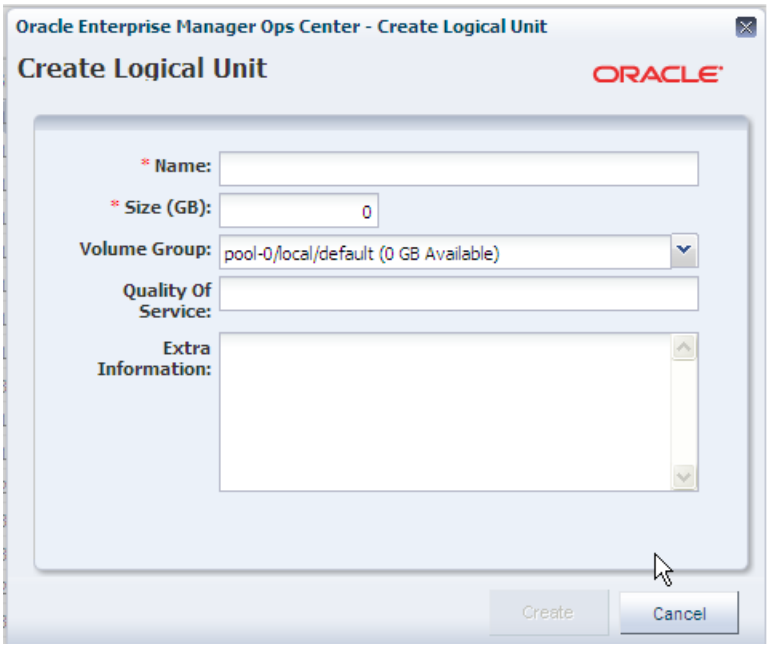

4. Enter the name and size for the new LUN.
5. Click **Create.**

Cloning a LUN

In a dynamic block storage library, you can make a copy of a LUN.

1. In the Asset section of the Navigation pane, select an asset in the Storage Array Group.
2. In the center pane, click the **Logical Units** tab.
3. Select the LUN to copy.
4. Click the **Clone LUN** icon.

Figure 6–11 Cloning a LUN

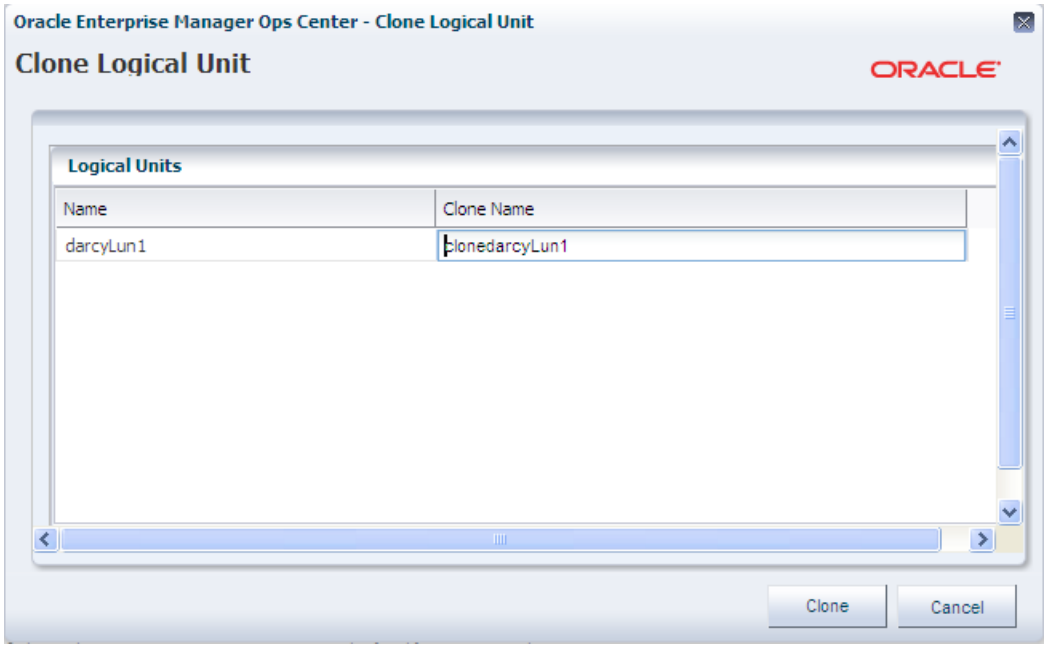

5. Enter the name and size for the new LUN.
6. Click **Clone.**

Oracle VM Storage Connect Plug-ins

Oracle Enterprise Manager Ops Center shares the capability of Oracle VM Manager to manage storage devices of various vendors. Oracle VM Storage Connect is an application programming interface (API) that exposes the storage device's features and attributes to Oracle Enterprise Manager Ops Center.

Displaying Storage Connect Plugins

When the plugin is installed on the Proxy Controller that supports the storage device, the storage devices can be managed through the Oracle Enterprise Manager Ops Center's user interface.

Use the following procedure to see the plug-ins that have been installed:

1. Expand **Administration** in the Assets pane.
2. Select **Enterprise Controller**.
3. Click the **Proxy Controllers** tab in the center pane.
4. Select one of the Proxy Controllers.
5. Click the down-arrow icon, located in the title bar of the display and shown in Figure 6–12.

Figure 6–12 Proxy Controllers

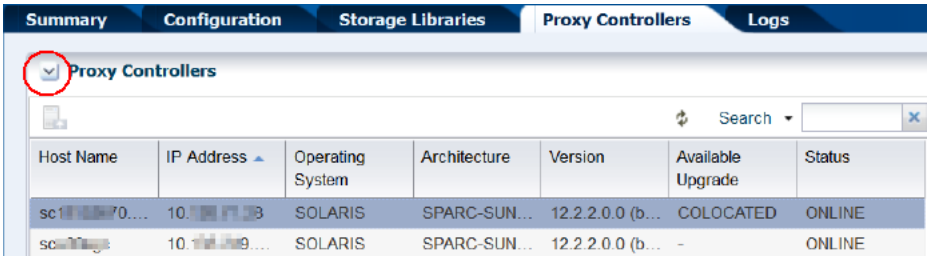

The center pane displays the Installed Storage Connect Plugins table, as shown in

Figure 6–13 Storage Connect Plugins

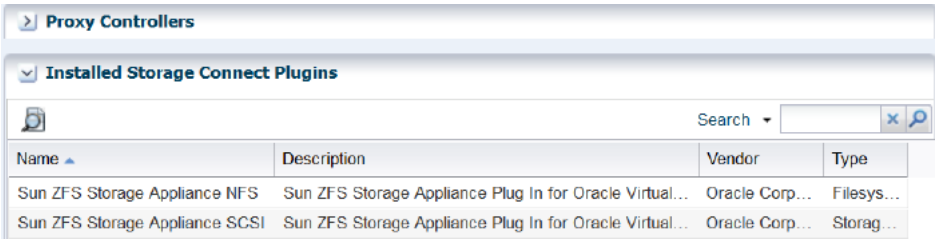

6. For information about one of the plugins, select the plugin and the click **View Plugin Details** icon

Adding a Storage Connect Plugin

Oracle Enterprise Manager Ops Center uses the same plugins as Oracle VM Manager. The plug-in software is available from the vendor and from Oracle's site, http://www.oracle.com/us/technologies/virtualization/storage-connect-partner-program/overview/index.html.

Install the plug-in in Oracle VM Manager software, according to the vendor documentation.

Storage Hardware

- RAID Controller
- NAS Storage Appliance
- File Server
- Storage Array
- Storage Server: Oracle ZFS Storage Appliance
- Storage Server: Exadata

The storage features are supported by various types of storage devices. The tabs in the center pane contain specific information about each type.

RAID Controller

RAID Controllers embedded in other types of storage report the following in addition to the hardware information:

- RAID Volumes

- RAID Levels
- Stripe size
- Number of disks

NAS Storage Appliance

Software and Storage Libraries can reside on the shares of an NFS server. Because the Enterprise Controller does not mount the NFS share, use an NFS server on a system that is close to the systems where the NFS share must be used, that is, the systems on which the virtualization hosts reside. The systems on which the Enterprise Controller and virtualization hosts reside must be able to write to the NAS shares as root and the files must be owned by root.

The procedure for setting up the share for a library depends on several site-specific factors such as the version of NFS protocol and name service management. The example in this section provides one method of configuring the share on an NFS server running on the Oracle Solaris 10 operating system. See *Managing Network File Systems* at http://docs.oracle.com/cd/E26502_01/html/E28997/index.html for the information about the Oracle Solaris 11.1 procedure.

Setting Up a Share on an NFS Server

1. On the NFS server, edit the `/etc/dfs/dfstab` file.
2. Add an entry to share the directory with share options that enable the NFS clients to have read and write root-level access to the share, such as:

   ```
   share -F nfs -o rw,root=<access_list> -d "<description>" /<directory>
   ```

   where <*access_list*> specifies the clients that can access the share as the root user, <*description*> is text to identify the purpose of the share, and <*/directory*> identifies the directory that you want to share on the NFS server. For example, to allow root access to the `/export/lib/libX` directory for all systems on the 192.168.1 subnet, add the following entry:

   ```
   share -F nfs -o rw,root=@192.168.1 -d "Share 0" /export/lib/libX
   ```

   See the `share_nfs(1M)` man page for information about NFS share options, and how to specify the access list.
3. Share the directory and verify that the directory is shared. For example:

   ```
   # share export/lib/libX
   # share
   -               /export/lib/libX   rw,root=@192.168.1 "Share 0"
   ```

After setting up a share on the NFS server, prepare the NFS client to mount the share.

Setting Up an NFS Client

1. On each NFS client, edit the `/etc/default/nfs` file.
2. Locate the `NFSMAPID_DOMAIN` variable and change the variable value to the domain name.
3. Verify the NFS share is visible on the client.

   ```
   # showmount -e <server-name>
   export list for <server-name>:
   /export/virtlib/lib0 (everyone)
   ```

To create a storage library, see *Configuring NAS Libraries* in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm.

File Server

For file-based storage devices such as NAS appliances and other NFS devices, Oracle Enterprise Manager Ops Center reports the following in addition to hardware information:

- File systems – The location in the file system hierarchy that is presented as a physical device. For example, an NAS share.
- Backing Devices – The physical device that supports a file system. For example, a RAID disk.

Storage Array

For block-based storage such as iSCSI SAN and Fibre Channel SAN devices, the storage array assets include the following tabs in the center pane:

- Volume Groups – Indicates a collection of physical disks or portions of physical disks.
- Logical Units (LUN) – Virtual disks created from the volume groups.

For a storage server and Oracle Enterprise Manager Ops Center to identify each other as eligible initiators and targets, each one's Fibre Channel World Wide Number (WWN) or iSCSI IQN must be registered with the other one. To create LUNs and make them available to Oracle Enterprise Manager Ops Center, see the storage server's documentation for instructions and *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing* at http://docs.oracle.com/cd/E26502_01/html/E29008/index.html for the procedures to perform the following:

- **Configure the initiator and the targets.** The initiator (Oracle Enterprise Manager Ops Center) must be able to recognize the targets (LUNs) and the targets must be able to recognize the initiator. Oracle Enterprise Manager Ops Center recognizes the targets because the WWNs of the storage server are recorded when the storage server is discovered. Any LUNs that have been assigned to that WWN are eligible to be used in a storage library. On the storage server, you must specify Oracle Enterprise Manager Ops Center's WWN as an initiator and assign LUNs to that initiator.
- **Enable multipathing on the Fibre Channel ports.**To use LUNs backed by Fibre Channel disks, you must enable multipathing on the Fibre Channel storage device or on its individual ports. Multipathed I/O (MPxIO) allows I/O devices to be accessed through multiple host controller interfaces. Multipathing is enabled by default on Oracle Solaris x86-based systems, but is disabled by default on Oracle Solaris SPARC-based systems. Use the `stmsboot -e` command to enable multipathing.
- **Create new LUNs.** It can take several hours for a new LUN to be displayed in Oracle Enterprise Manager Ops Center's user interface.

Storage Server: Oracle ZFS Storage Appliance

See the Oracle ZFS Storage Appliance documentation for specific information about the storage appliance. To view the *Sun ZFS Storage 7000 System Administration Guide*, log in to the Unified Storage System software interface and click **Help** in the top right corner of any screen.

This type of storage appliance contains a service processor, which must also be discovered. See Oracle ZFS Storage Appliance for this procedure.

Managing an Oracle ZFS Storage Appliance

After the storage appliance is discovered, you can manage it as you do other assets with the additional capabilities that the Oracle ZFS Storage Appliance provides. From the Oracle Enterprise Manager Ops Center UI, you can launch the storage appliance's UI. Use the following commands in the Action pane, to launch a specific page of the appliance's user interface. For each one, enter the credentials for the appliance and then perform the appliance tasks.

- **Launch Appliance UI** opens a new browser window or tab for the main page.
- **Launch Detailed Dashboard** opens a new browser window or tab for the status page.
- **Launch Analytics** opens a new browser window or tab for the dynamic analysis page.
- **Manage Shares** opens a new browser window or tab for the share configuration page.
- **Manage Services** opens a new browser window or tab for the data services configuration page.

Displaying Volume Capacity

The Oracle ZFS Storage Appliance product uses the term `project` to indicate a volume group. The Oracle ZFS Storage Appliance reports it current capacity to Oracle Enterprise Manager Ops Center, where it is displayed on the Dashboard.

- Allocated Size is the sum of all allocated space to LUNs and files systems in the select project or volume group.
- Total Size is the quota of the project or volume group. If no quota has been specified, the Total Size is the same as the Total Storage Space, which is all space reserved for other project/volume group quotas.
- Free Size is the Total Size - Allocated Size
- Used Size is the sum of the used space in all LUNs/files systems in the selected project/volume group

Auto Service Requests for the Oracle ZFS Storage Appliance

Oracle ZFS Storage Appliances have the capability to contact My Oracle Support, to provide analytic data, and to generate Oracle Auto Service Requests (ASR). The capability does not conflict with Oracle Enterprise Manager Ops Center's capability to create ASRs. If you prefer, exclude an Oracle ZFS Storage Appliance from Oracle Enterprise Manager Ops Center's communication with My Oracle Support by using the blacklisting option described in the *Oracle Enterprise Manager Ops Center Administration Guide*.

Provisioning and Updating an Oracle ZFS Storage Appliance

Note: You can update the storage appliance using provisioning profile and deployment plan. The Update Firmware action is not available.

---

1. Download the appliance update from My Oracle Support and uncompress the file to obtain the image for the appliance update.
2. Upload/import the appliance update image into a software library and accept the default action of creating a profile for this image.
3. Create a plan from the Update Storage Appliances plan template. Select the software library and then select the profile for the image.
4. To apply the deployment plan, select the plan and then select the storage appliances to update.

If it is necessary to downgrade a storage appliance, use the following procedure:

1. Expand Assets and then Storage in the Navigation pane.
2. Select the storage appliance.
3. Click the **System Details** tab in the center pane.
4. In the Appliance Updates table, select the update to remove.
5. Click the **Rollback Appliance Software** icon.
6. Confirm the action to submit the job.

Storage Server: Exadata

The Oracle Exadata Storage Server is a fast, high capacity storage server. Each server comes has 12 disks connected to a storage controller. Exadata has three disk groups that span all of the disks so that each physical disk is in each disk group. Having all physical disks participating in data retrieval and storage maximizes performance.

Oracle Enterprise Manager Ops Center discovers the Exadata Storage Server and its disks when it manages a SuperCluster Engineered System. See Chapter 23, "Oracle Engineered Systems" for more information.

Opaque Storage and Opaque Filesystems

When a virtual host, managed by Oracle Enterprise Manager Ops Center, uses storage resources that are not managed, Oracle Enterprise Manager Ops Center cannot retrieve information about the storage resources. The storage is opaque.

In versions before Release 12.2.0.0, Oracle Enterprise Manager Ops Center could not monitor opaque storage or filesystems and no operations could be performed. In this release, when Oracle Enterprise Manager Ops Center discovers a virtualization host or virtual host, the product software can also identify the virtual disks in use. For Oracle VM Server for SPARC and Oracle VM Server for x86, the storage remains opaque until you perform one of the following:

- You enable sharing on the virtual disk in a shared storage library. The opaque disk or filesystem can be shared among Oracle VM Servers.
- If the virtual host gets access to storage through a SAN, you can add the opaque LUN to the virtualization host's SAN storage library. Oracle Enterprise Manager Ops Center manages the LUN as a raw disk.

Note: Oracle Enterprise Manager Ops Center makes a distinction between "shared storage" and "sharable storage."

- Shared storage: You inform Oracle Enterprise Manager Ops Center that the opaque disk is accessible through the storage library. You declare the disk as shared, using the Enable Sharing action.
- Sharable storage: You can attach the same disk to more than one virtual host. The disk is sharable storage.

To Declare a Logical Domain's Opaque Storage as Shared

1. Select the Oracle VM Server for SPARC in the Assets tree.
2. Select the logical domain that uses the opaque disk.
3. Click the **Storage** tab. For each opaque disk, the entry in the Disk Type column is `opaque` and the entry in the Shared column is empty.
4. Select the disk.
5. Click the **Enable Sharing** icon and then confirm that you want the storage to be shared.

Figure 6–14 Enable Sharing on a Virtual Disk

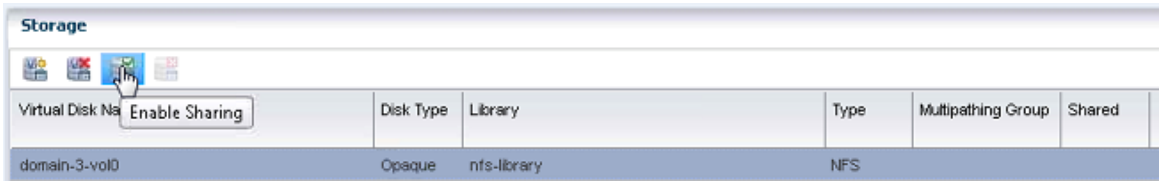

When the job is completed, the entry for the virtual disk's Shared column has a check symbol.

Figure 6–15 Shared Virtual Disk

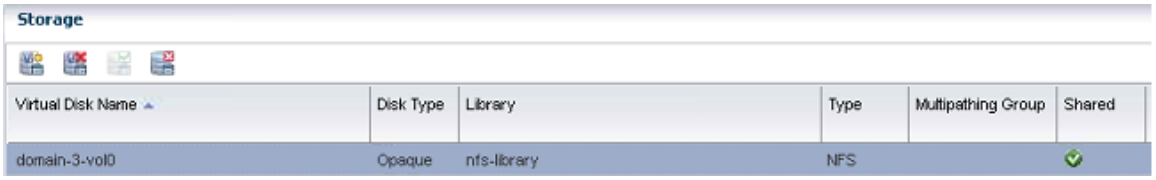

Storage Profiles

Oracle Enterprise Manager Ops Center provides default profiles for the following operations:

- Configuring RAID Controllers
- Updating the Oracle ZFS Storage Appliance

Multipath Storage for Logical Domains

Logical domains use storage libraries to store their metadata and their operational data. The path to the storage consists of the following:

- Name of virtual disk
- Name of the backing device such as the LUN or raw file.

Creating more than one path from a logical domain to its storage ensures that the logical domain can continue to operate if one path is not accessible. Using multiple I/O Domains provides this redundant access. The path to the storage then consists of the following:

- Name of I/O Domain
- Name of virtual disk
- Name of the backing device such as the LUN or raw file.

The control domain and the I/O Domain must have access to the same LUNs. The operation requires Oracle Solaris MPxIO

Use the **Add Storage** icon on the logical domain's Storage tab to create multipathed storage for a logical domain.

Use the **Edit Storage** icon on the logical domain's Storage tab to add a redundant path to existing storage.

To Set Up Multipath Storage for a Logical Domain

1. Select the Oracle VM Server for SPARC in the Assets tree.
2. Select the logical domain.
3. Click the **Storage** tab. For virtual disk, the entry in the Multipathing column is empty.
4. Select the disk.
5. Click the **Edit Storage** icon.

Figure 6–16 Enabling Multipath Storage on a Virtual Disk

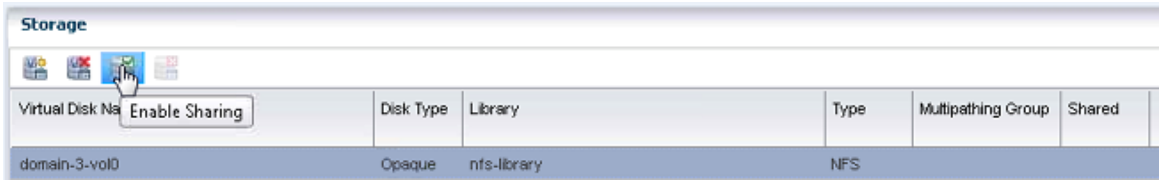

Enter the path to the storage in the Multipathing Group field.

High Availability for Storage Resources

Storage devices for an HA configuration must meet these requirements:

- Storage must be transferable between the primary and secondary Enterprise Controller systems. Do not attempt to use local file systems for high-availability storage.
- Storage must offer data redundancy capability, such as mirroring or RAID 5.
- Storage must offer performance that is sufficient to support operations.

- Storage must have the capacity to hold the data that the Oracle Enterprise Manager Ops Center software stores in the `/var/opt/sun/xvm` directory structure.

A variety of storage solutions meet these criteria, including NAS appliances, hardware RAID arrays and external JBODs. Storage can be attached directly to the Enterprise Controllers or through Storage Area Networks.

You must determine what storage solution offers the required capacity, performance, connectivity, and redundancy capabilities. Configuration procedures vary greatly among the available storage solutions, and among operating systems.

You must determine the specific failover procedures to use for the HA storage solution. Contact My Oracle Support to determine the procedures to use for your particular installation.

Related Resources for Storage

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- Chapter 16, "Storage Libraries for Virtualization"
- Chapter 11, "Hardware" for information about discovering and monitoring storage hardware
- *Deploy Storage Workflow* in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm`
- To view the *Sun ZFS Storage 7000 System Administration Guide*, log in to the Unified Storage System software interface and click Help in the top right corner of any screen. You can also access this guide at the host name or IP address of the storage system:
  - `https://`*`hostname`*`:215/wiki`
  - `https://`*`ipaddress`*`:215/wiki`

  See the Oracle ZFS Storage Appliance Software product information page at `http://www.oracle.com/us/products/servers-storage/storage/unified-storage/sun-storage-7000-uss-103104.html` for links to more information.

7

Networks

The following information is included:

Introduction to Networks

Oracle Enterprise Manager Ops Center manages network resources, from the physical to the virtual. Fabrics provide the physical infrastructure and network domains provide the logical infrastructure. Networks are created from the resources of a network domain.

Oracle Enterprise Manager Ops Center supports Ethernet and InfiniBand network protocols. While the Ethernet interconnect is the established and common interconnect, InfiniBand is popular in high-performance computing environments because it maximizes the speed of transactions using the short, multiple connections found in clusters and data centers.

- For an Ethernet network, both tagged and untagged VLANs are supported. An untagged VLAN has no VLAN IDs. Use tagged VLANs to create multiple networks on a fabric that use the same network address but different VLAN IDs. The network instances are independent of each other. However, in a server pool, use either all tagged VLANs or all untagged VLANs; do not mix the types of network in a server pool. For more information, see Mixed Network Tagging

Mode Configurations in Server Pool.

Note: Previous versions of the product software did not create independent network instances.

- For an InfiniBand network, partitions are supported.

Note: If you use an InfiniBand switch in an Ethernet network, the ports on the switch have Ethernet names.

Roles Required for Networks

Table 7–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 7–1 Network Tasks and Roles

Task	Role
Add Fabric	Network Admin
Remove Fabric	Network Admin
Discover and Manage the Switches	Network Admin
Configure Network for Server Deployment	Server Deployment Admin

Actions for a Networks

After a network is discovered or created, you can perform the following actions, depending on the requirements.

- Discover and manage the switches
- Add a fabric to network domain

Location of Network Information in the User Interface

Table 7–2 shows where to find information.

Table 7–2 Location of Network Information in the BUI

Object	Location
Fabric	Expand Networks in the Assets pane. Then select Fabrics.
Physical Fabric	Expand Networks in the Assets pane. Then select Fabrics and select Network Switches.
Network	To see all networks, regardless of type, expand Networks in the Assets pane. Then select Network Domains.
Services of a Network	Network Services tab: time server, WINS, DNS, and NIS. To modify these services, edit the network services. You cannot change the network's IP address or name.
Network Domain	Expand Networks in the Assets tree. The Default Network Domain is the first item.
Physical switch	Expand Assets and expand Network Switches. To see each port, click the Connectivity tab.

Fabrics

The fabric is the physical network infrastructure, such as switches, ports, host bus adapters, that provides network resources, through a network domain, to virtual assets.

When you use Oracle Enterprise Manager Ops Center to discover a physical switch or the host of switch, all the switching fabrics that the switch supports are also discovered. One physical fabric supports many fabrics, also called data links. The physical fabric is the collection of all switch ports, links, and physical interfaces or endpoints.

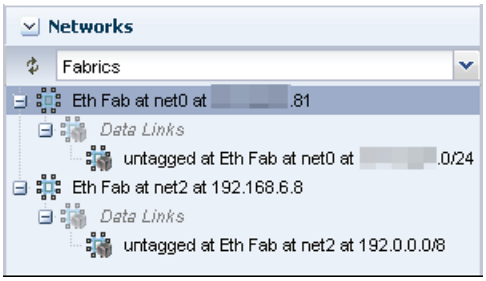

For each Ethernet fabric, the maximum VLAN ID range is 4096, which allows you to create 4096 networks.

For each InfiniBand physical fabric, the maximum number of partitions keys is 32000 so you can create 32000 partitions. Each partition is a logical fabric. For example, if a server has two partition keys, it participates in two different partitions.

Fabrics provide resources to the virtual networks they support in a manner that depends on their type: fully-managed, host-managed, or unmanaged. Table 7–3 shows the types of fabrics.

Table 7–3 Fabrics and the Network Domain

na	What Is Managed	Capability	Comments
Fully-managed switched fabric	The switch is discovered and managed.	For each VLAN ID or partition key, you can create a static or dynamic private network.	This type of fabric can be achieved on only the Sun Ethernet 10GbE Fabric switch or the Sun Datacenter InfiniBand switch and gateway.
Host-managed fabric	The host connected to the Ethernet switch is discovered and managed. A range of VLAN IDs has been assigned.	For each VLAN ID, you can create a static or dynamic private network.	To create a host-managed fabric, use the **Define Ethernet Fabric** action to specify the fabric and its VLAN ID range.
Unmanaged fabric	The Ethernet fabric is discovered or declared during the discovery of another asset, but the switch is not managed.	If the fabric has existing networks with VLAN IDs, you can create static private networks.	To convert an unmanaged fabric to a host-managed fabric, use the **Assign VLAN ID Range** action.

Network Domains

- Default Network Domain
- User-Defined Network Domains
- Editing Attributes of a User-Defined Network Domain

A network domain is a container for fabrics, managed networks, and private networks. The network domain handles the relationship between the physical fabrics and the virtual assets, such as virtualization hosts or server pools. The fabrics provide data links and IP subnets to the network domain, which then provides networks to the virtualization hosts and server pools.

Within the network domain, networks that have been discovered or specified are available for assignment. These are called public networks because their IP address space has been specified for their exclusive use. Another type of network is private, that is, the network is created using an IP address space that the network domain allocates to it.

A fabric can contribute to more than one network domain. When a network domain has more than one fabric, you designate one of the fabrics as the anchor fabric, which is the fabric from which new networks are created.

Public networks can be members of more than one network domain because their IP addresses are specific and dedicated. Private networks exist only within a specific network domain so two network domains could construct a private network with the same IP address without a conflict.

In Oracle Enterprise Manager Ops Center, networks become part of a network domain in the following ways:

- An asset that has a network is discovered.
- A user creates a network.
- A network is created when it is required. This is a dynamic network.

Oracle Enterprise Manager Ops Center operates on more than one layer of the Open Systems Interconnection model, using the network domain. Table 7–4 shows what the network domain manages in the physical to logical stack.

Table 7–4 Elements of a Network Domain

Layer	Asset	What Is Managed	Capability
Layer 3 Network: IP address	For Ethernet: fabric networks For InfiniBand: non-fabric networks	IP subnet and mask IP address range VLAN or Partition Services Routing	The network provides connectivity.
Layer 2 Data links	For a tagged Ethernet: VLAN For an untagged Ethernet: portID For InfiniBand: partition	VLAN IDs Partition keys (P-key)	A virtual host uses the virtual NIC and a virtual switch in a VLAN or partition.
Layer 1 Physical: switches, ports, host bus adapters	Fabrics	Varies, by type of fabric. See Table 7–3	Varies by type of fabric. See Table 7–3

Default Network Domain

The Oracle Enterprise Manager Ops Center software always has a Default Network Domain and all networks are members of it. If you have upgraded your product software from the previous release, the existing managed networks are now in the Default Network Domain. A new network becomes a member of the default network

domain. If you direct the new network to a user-defined network domain, the network is also a member of that network domain.

User-Defined Network Domains

Like the default network domain, a user-defined network domain provides network resources to a server pool or virtualization host. You create a network domain to support the use of virtualization hosts, server pools, or a virtual datacenter. For example, a virtual datacenter uses server, storage, and network resources in a dynamic way, allocating and releasing resources whenever necessary. The network domain provides the network resources to the virtual datacenter.

When you create a network domain, you set a limit on the number of networks that can be created in the network domain. Increase the number of networks when accounts in a virtual datacenter are not able to create vnets.

A new user-defined network domain includes the address space specified as private by the RFC 1918 specification. These addresses cannot be routed to the Internet and provide a way for organizations to create intranets. If you organization uses a portion of this private address space, reserve these IP addresses when you create a network domain so that the network domain does not use them.

Editing Attributes of a User-Defined Network Domain

You can change the name and description of the network domain and you can change the number of dynamic networks that are in use simultaneously.

To Edit Attributes of a Network Domain

1. Expand **Networks** in the Navigation pane.
2. Click the network domain.
3. Click **Edit Attributes** in the Actions pane.

   The Details tab is displayed in the center pane. You can now change the Name Description, and Number of Networks fields.
4. Edit the name or description or increase the number of networks.
5. Click **Save**.

Networks

- Requirements for a Network
- Limitations of Networks
- Public Networks and Private Networks
- Assigning Networks to a User-Defined Network Domain
- Bandwidth Management
- Creating IPMP Groups
- Creating Link Aggregation

In Oracle Enterprise Manager Ops Center, networks are the discovered and managed IP subnets. Oracle Enterprise Manager Ops Center manages network resources for its virtualization hosts.

Note: These networks are part of Oracle Enterprise Manager Ops Center's virtualization services. For a description of the networks that support the product, see Oracle Enterprise Manager Ops Center's Networks.

Networks are associated with a single virtualization host or a server pool, which contain multiple virtualization hosts. When you assign a network to a server pool, the network is accessible to each virtualization hosts in the pool and every guest of each virtualization host.

You can use networks to do the following:

- Manage individual virtualization hosts
- Connect virtualization hosts to the Proxy Controller
- Allow guests to communicate with each other or with the Internet
- Connect remote JMX with the public API

A network depends on the physical network interface card (PNIC) that is available to the host. You can create one network for each physical network interface card. If one host has two PNICs, it is a good practice to create two networks: a management network and a data network. Then place all virtual hosts on the data network, keeping them separate from the management network. The management network is dedicated to giving access to internal resources of the data center.

Figure 7–1 shows how two virtualization hosts participate in two networks. The actual network connection is made to the PNICs in the virtualization host. Network A is connected to PNIC 1 of both hosts and Network B is connected to PNIC 2 of the hosts.

Figure 7–1 Network with Virtual Hosts

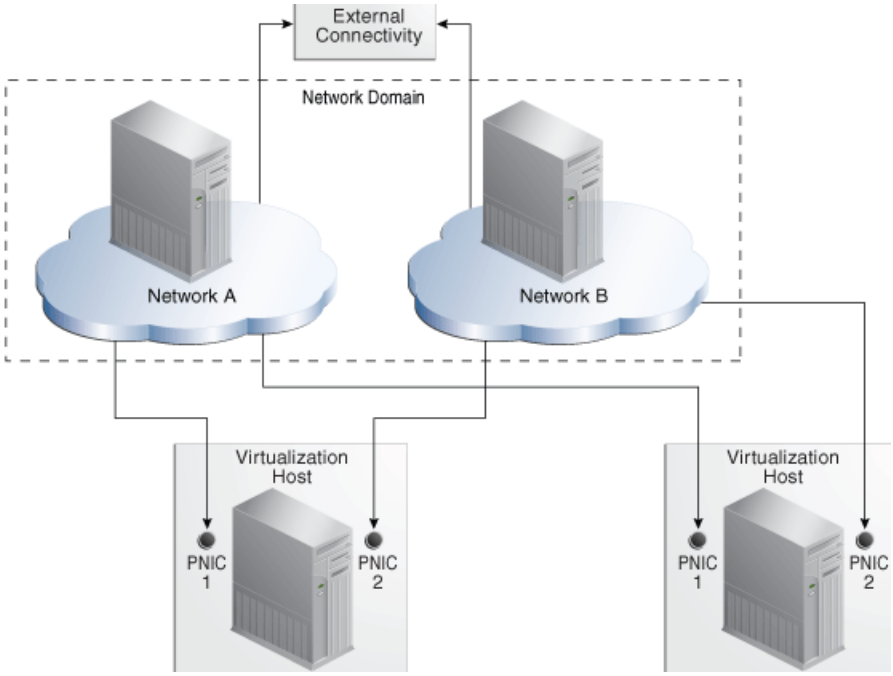

Requirements for a Network

A network requires a physical network interface or a link aggregation and the following specifications:

- IP address and netmask or CIDR format
- If you use static IP addressing, the IP address of the management interface

  If you use dynamic IP addressing, the range of allowed IP addresses and the gateway address

Before you attach a network to a server pool, verify that each virtualization host in the server pool has a physical network interface to the network so that all members of the pool can continue to share the network resources of the server pool.

Limitations of Networks

Ethernet networks and InfiniBand networks in the default configuration must not have the same CIDR (Classless Inter-Domain Routing) or with sub-blocks of the same CIDR. When discovering assets, verify that the Ethernet networks or InfiniBand networks in the default configuration comply with the following constraints. If so, reconfigure the asset before it is discovered.

- No assets with overlapping management networks. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping. However, you can use the same CIDR (not sub-block) for different assets. For example, you can use 192.0.2.1/22 as a CIDR for the Ethernet network for two assets.
- No overlapping private networks. For example, two private networks cannot have the same CIDR.
- No overlapping public networks. However, you can use the same CIDR (not sub-block) for different assets. For example, you can use 192.2.0.0/22 as a CIDR for the public EoIB network for multiple engineered systems.

For more examples of valid network specifications, see Example Oracle SuperCluster Network Configurations.

Starting in Release 12.2.2.0.0, you can configure InfiniBand networks to share a CIDR. However, when you allow networks to have the same CIDR, you must ensure that all the NFS shares used by storage libraries have a unique pathname. Oracle Enterprise Manager Ops Center identifies an NFS share by its CIDR, its NFS server, and the share name. When the CIDR addresses are the same, either each NFS server must have a unique name or each share must have a unique name. Use the following procedure to allow overlapping InfiniBand networks. You must have the Ops Center Admin role to change the value of a product property variable.

1. Log in to the Enterprise Controller.
2. Click **Administration** in the Navigation pane.
3. Click **Enterprise Controller.**
4. Click **Configuration** in the center pane.
5. In the Subsystem field, click **Network/Fabric Manager**. The `oem.oc.networkmgmt.ib.overlapping.enabled` property's default value is `false`.
6. Click in the **Value** field to edit it. Change the value to `true`.
7. Click the **Save Properties** icon.
8. Stop and restart the Enterprise Controller:

```
/opt/SUNWxvmoc/bin/ecadm stop
/opt/SUNWxvmoc/bin/ecadm start
```

Public Networks and Private Networks

Networks are introduced into Oracle Enterprise Manager Ops Center in the followings ways:

- By discovering the fabric that supports existing networks. All the attributes are discovered but, other than the name and description, they cannot be changed. All networks of a discovered fabric are in the Default network domain.
- By specifying the network completely, using the resources provided by a fabric. Use the **Define Network** action to specify the IP addresses and the VLAN IDs for an Ethernet network, based on what the fabric can provide. To create untagged networks, specify a VLAN ID of -1. For InfiniBand networks, the P-keys are assigned automatically.

  Note: In previous versions of the product software, this action was called **Manage Network**.

- The **Create Private Network** action creates a network from the resources of a user-defined network domain. Oracle Enterprise Manager Ops Center allocates the IP address from the addresses available within the network domain.

If you are creating networks for the use of virtual datacenters, create a private network to ensure that the virtual datacenter has exclusive use of the IP address space that it gets from the network domain. See *Creating Private Networks in a Virtual Datacenter* in the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.

Note: The Automated Installer for Oracle Solaris 11 uses the `installadm` service to identify all network interfaces and adds them to the `/var/ai/ai-webserver/listen-addresses.conf` file.

When you add a network interface, run the following command to update the `installadm` service and ensure the Automated Installer has access to all network interfaces:

```
svcadm refresh system/install/server
```

To see the list of network interfaces handled by the `installadm` service, view the `/var/ai/ai-webserver/listen-addresses.conf` file.

Assigning Networks to a User-Defined Network Domain

All networks are in the Default network domain, but a network can be a member of more than one user-defined network domains. You place a network into a user-defined network domain so that the network is available to the server pools associated with the network domain. Use the **Assign Network** action to place a network into a specific network domain.

1. Expand **Networks** in the Navigation pane.
2. Select the user-defined network domain.
3. Select **Assign Network** in the Actions pane.

Figure 7–2 Assign Network

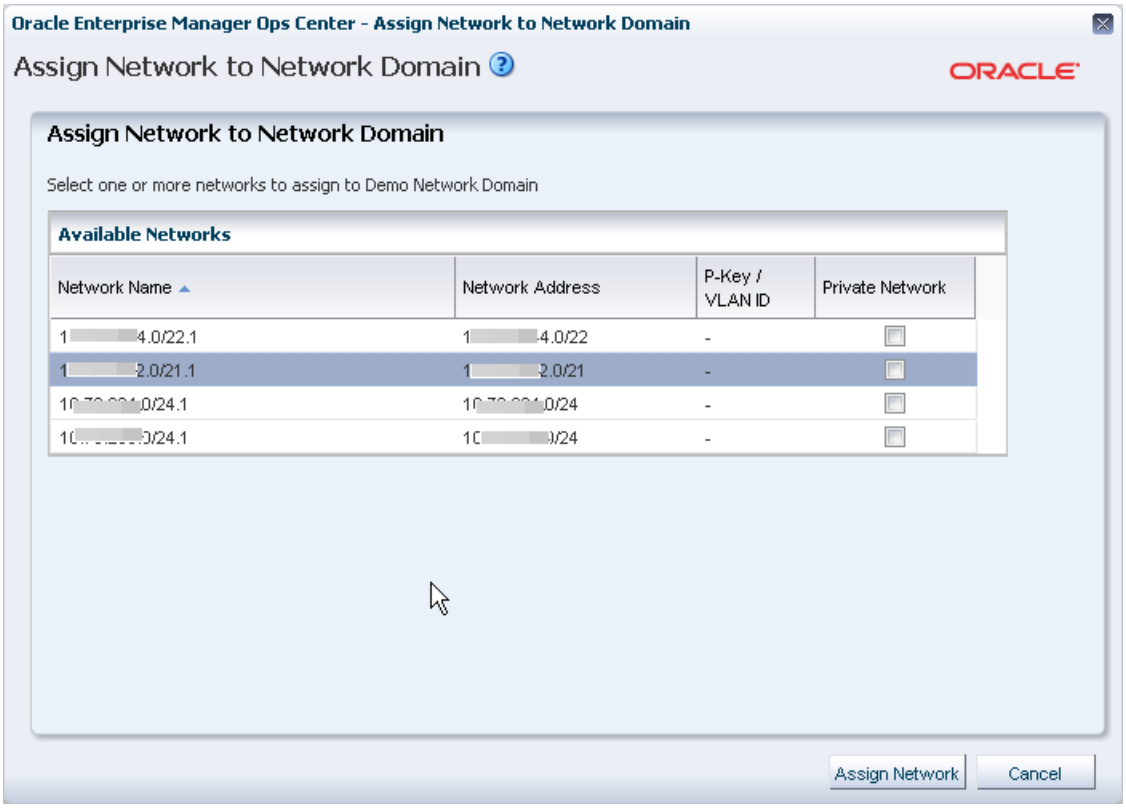

4. Select a network. If the network is does not share its IP addresses with any other network, select the Private Network option.
5. Click **Assign Network** to submit the job.

Bandwidth Management

A data link is a physical NIC, an aggregated link, or a virtual NIC. When a new data link is created, the operating system sets the default bandwidth flow. You cannot remove this flow. The flow is removed only when the physical link is removed.

In Oracle Solaris 11 operating system environments, you can manage the bandwidth flow of a data link, prioritizing the network traffic on the link and setting the maximum bandwidth limit.

Managing the Bandwidth Flows for a Data Link

1. Expand **Assets** in the Navigation pane.
2. Select an Oracle Solaris 11 operating system.
3. Click the **Networks** tab in the center pane.
4. Click the **Bandwidth Management** subtab in the center pane.
5. To modify a flow, click the **Modify** icon. To create a new link, click the **Add** icon, then specify a name for the flow and the physical network interface.

   The name of flow must meet the following requirements:

   - The first character must be alphabetic.

- All characters must be alphanumeric: a-z, A-Z, 0-9, underscore ('_'), period ('.'), or hyphen ('-').
- Maximum number of characters is 127.

6. Set the new bandwidth properties, as described in Properties of Bandwidth Flow.

Properties of Bandwidth Flow

- Priority: Set the priority of the network traffic on the link as high, medium or low.
- Bandwidth Limit: Enable the bandwidth limit to allocate guaranteed bandwidth to the specified link. Enter the maximum value for bandwidth limit in Kbps, Mbps, or Gbps.
- Set attributes for the data flow to identify its network traffic:
  - Local and Remote IP: The source and destination IP address.
  - Transport: The Internet Protocol used such as TCP, UDP. SCTP, ICMP.
  - Ports: The source and destination ports for TCP, UDP, and SCTP.
  - DS Field: The type of service field in the IP packets' header.

Creating IPMP Groups

For information about how IPMP groups work in Oracle Solaris 11.2, see http://docs.oracle.com/cd/E36784_01/html/E37476/index.html. For information about how IPMP groups work in Oracle Solaris 11.1, see http://docs.oracle.com/cd/E26502_01/html/E28993/index.html. For Oracle Solaris 10 documentation, see *IP Services* at http://docs.oracle.com/cd/E26505_01/html/E27061/index.html.

From the Network tab, view and manage IPMP groups as shown in Figure 7–3.

Figure 7–3 IPMP Groups

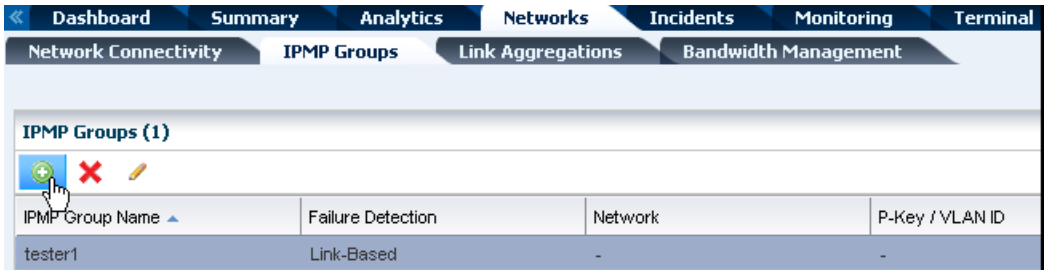

For all types of networks, you create an IPMP group by specifying the following:

- The link-based failure detection is enabled by default. To use, Probe-Based failure detection, select the **Probe-Based** option and provide the test address to track the interface status.
- You must assign the data addresses for the physical interfaces in the IPMP group. Data traffic flow uses the data addresses hosted on the IPMP interface and flows through the active interfaces of the group.
- The active and the standby interfaces of the group. By default, an interface added to an IPMP group is active. You can configure as many standby interfaces as you want for the group. The list of available network interfaces contains the interfaces

that qualify, depending on the operating system, the type of network you select, and the network's existing attributes.

- For an Ethernet network on Oracle Solaris 10:
  - If the network has a VLAN ID, you can select the Tagged mode and you can keep or change the VLAN ID.
  - If the network has a VLAN ID, you can select the Untagged mode and you can keep or change the VLAN ID.
  - If the network does not have a VLAN ID, the option to make it a tagged or untagged network is not available.
- For an Ethernet network on Oracle Solaris 11, you can specify its media type:
  - For the Ethernet media type, the resulting options are the same as for an Ethernet network in the Oracle Solaris 10 environment.
  - For the InfiniBand media type for a network with a VLAN ID, the P-Key field is displayed.
- For an InfiniBand network on Oracle Solaris 10, no network interfaces are included in the list of available network interfaces.
- For an InfiniBand network on Oracle Solaris 11:
  - For the InfiniBand media type for a network with a VLAN ID, the P-Key field is displayed and you can keep or change its value.
  - For the Ethernet media type, the P-Key field is not available.

To Create an IPMP Group

1. Select the Oracle Solaris OS in the Assets section.
2. Click the **Network** tab in the center pane.
3. Click the **IPMP Groups** subtab in the pane. The existing IPMP groups are listed in the subtab.
4. Click the **Create IPMP Group** icon to open the Create IPMP Group wizard.

Figure 7–4 Specify IPMP Group Details for an Ethernet Network on Oracle Solaris 11

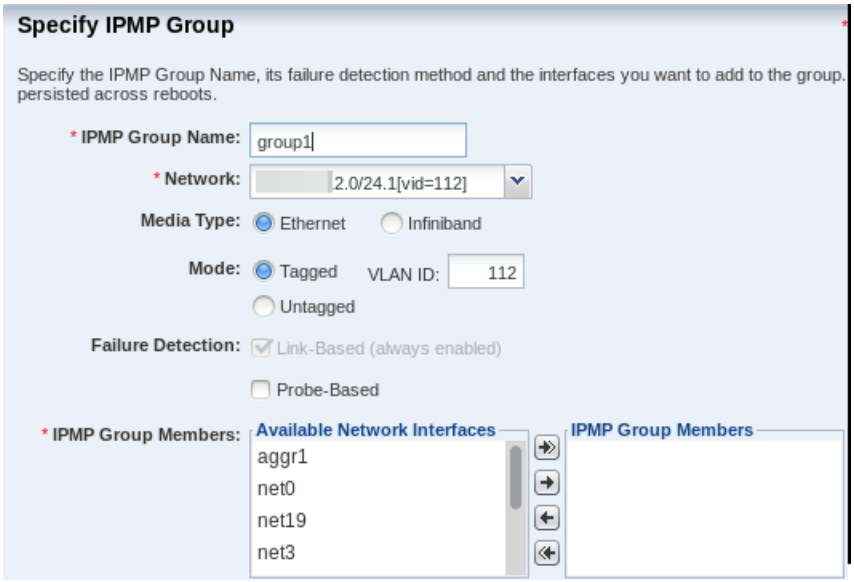

Figure 7–5 Specify IPMP Group Details for an InfiniBand Network on Oracle Solaris 11

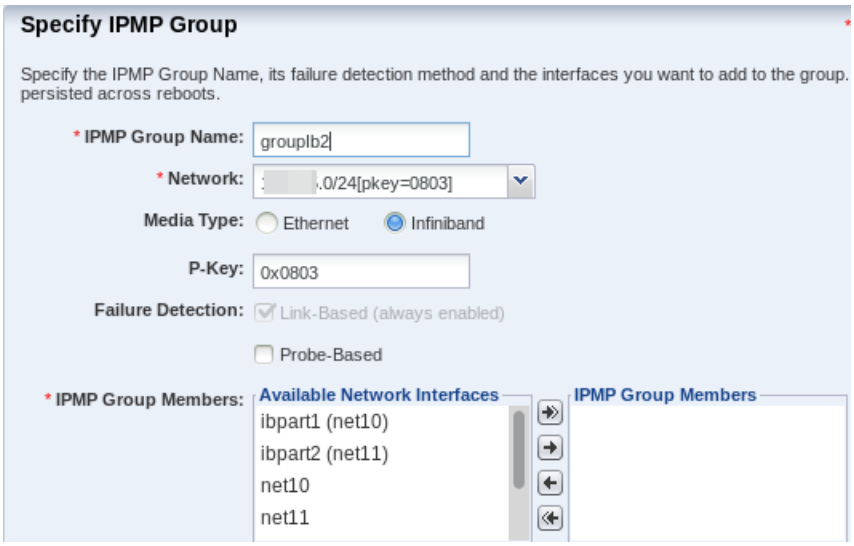

5. Enter the following details for the IPMP group, as shown in Figure 7–4:
   a. Provide a name for the IPMP group.
   b. Select a network from the list of available network interfaces.
   c. Depending on the type of network you select, specify the characteristics of the network interfaces.
   d. The Link-Based failure detection is always enabled by default. Select whether you want to also enable Probe-Based failure detection.
   e. Select each interface you want in the IPMP group from the Available Network Interfaces list and click the right arrow to include the interface in the group.

   Click **Next** to specify the NIC settings.
6. If you enabled probe-based failure detection, enter the test address for the NICs.

7. Select the interfaces that are in standby mode.

   You must have at least one active interface in the group. Click **Next**.

8. Enter the data address for the active interfaces of the group and select whether the interface has a failover and click **Next**.

9. Review the information and click **Finish** to create the IPMP group.

Creating Link Aggregation

The link aggregation conforms to the Link Aggregation Control Protocol (LACP) as described in the IEEE 802.3ad Link Aggregation Standard specification. The switch that communicates with the network interface must also support LACP.

To create a link aggregation, specify the following:

- Load balancing policy
- Link Aggregation Control Policy (LACP) mode and timer
- MAC address policy and if required, the MAC address

From the Network tab, you can create and manage the link aggregations as shown in Figure 7–6.

Figure 7–6 Link Aggregations

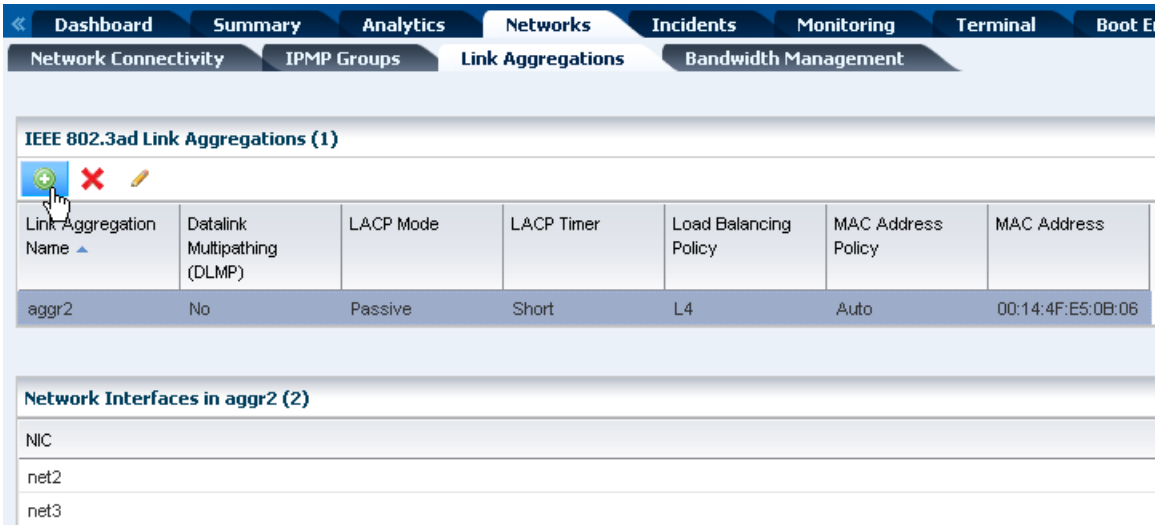

To Create a Link Aggregation

1. Click the **Network** tab.
2. Click the **Link Aggregation** subtab.
3. Click the **Create Link Aggregation** icon to open the wizard.
4. Enter the name of the link aggregation. By default, the name starts with `aggr` Append a number to make the name unique.
5. Select the network interfaces to be in the aggregation by selecting each one from the list of available interfaces and clicking the right arrow to include them in the list of network interfaces in the link aggregation. Click **Next.**
6. Specify the following information:

- Policy for load balancing by setting the type of packet identification in outgoing traffic. Packets with the same identification ar e routed to the same network interface in a link aggregation. The L4 policy is the fastest and the default.
- LACP mode, which is the type of LACPDU or packet required between the link aggregation and the switch. The value of `off` requires no packet, the value of `active` sends packets at intervals set by the LACP timer, and a value of `passive` sends a packet when the switch sends one.
- LACP timer, which sets the time for the LACP active mode. The default is 1 second.
- MAC address policy is either Auto or Fixed. The Auto policy generates the MAC address. The Fixed policy uses the MAC address you enter.

Click **Next** to view the summary.

7. Review the information and click **Finish** to create the link aggregation.

Properties of a Network

IPv4 and IPv6 Protocols

Some environments have a mix of IPv4 and IPv6. Oracle Enterprise Manager Ops Center is "IPv6-aware." If an asset has an IPv6 network interface, Oracle Enterprise Manager Ops Center can read it and displays its information, but it cannot provision an IPv6 network or use IPv6 networks to discover, monitor, or provision assets.

Routing Mode

A virtual host uses the network assigned to it according to the host's routing mode. You specify a virtual host's routing mode during its initial configuration if you do not accept the default mode, Automatic Routing. Oracle Enterprise Manager Ops Center supports the following routing modes:

- Automatic Routing – This is the default routing mode. Applying the static routes depends on the following conditions:
  - If your site defined a default gateway or static route or retrieved one from the DHCP server, this route is used and dynamic routing is disabled.
  - If no default gateway or static route is available, dynamic routing is enabled.
- Dynamic Routing Off – The virtual host uses the default gateway and any static routes configured for the network. The default gateway is retrieved from the DHCP server.
- Dynamic Routing On – The virtual host uses routes provided by the dynamic routing service. The default gateway and any static routes configured for the network are ignored.

Static Route for the Network

Static routes specify the route for external access. Although you define a default gateway for a network, it might not reach a particular subnet. In this case, you must also provide a static route for the subnet.

When you create a network, you can specify the static route. To add static routes after the network has been created, use the following procedure.

To Add a Static Route for the Network

1. Click **Managed Networks** in the Navigation pane.
2. Select a network from the list of networks.
3. Click **Edit Network Attributes** in the Actions pane.
4. Click the **Add** icon in the Static Routes table. A row is added to the table.
5. Enter the values for destination IP, netmask, and gateway.
6. Click **Finish**.

You can delete a static route and change the order of the routes using the icons in the Static Routes table.

Address Allocation Method

When you define a new network, you specify how its IP address is assigned:

- Static IP: You enter a specific IP address.
- Use System Allocated IP: Oracle Enterprise Manager Ops Center assigns an available IP address.
- Do not allocate IP: No IP address is assigned.
- DHCP: Use the DHCP service to acquire an IP address.

When you create an Ethernet network without an SR-IOV connection for a control domain, you have an additional option: Do Not Plumb Interface.

Maximum Transmission Unit (MTU)

The default size for the network's Maximum Transmission Unit (MTU) is 1500 bytes. If your network interface card is one of the following types, you can change the size of the MTU to a size between 576 and 9216 bytes. However, to assign the network to a logical domain, the minimum MTU size is 1500 bytes.

- `e1000g`
- `ce`
- `nxge`
- `nge`
- `bge`
- `xge`
- `hme`
- `ixgbe`
- `hxge`

- ipge
- igb

When you specify a size greater than 1500 bytes, Oracle Enterprise Manager Ops Center modifies the network interface card's MTU size. For other types of network interface cards, the MTU is changed when the card's driver firmware is updated to support the new MTU size. However, to change the MTU value for an IPMP group, you must edit the MTU value manually.

Note: When you provision an operating system, the MTU size resets to the default value. You must change the MTU again after you provision the system.

Network Utilization

Oracle Enterprise Manager Ops Center collects information every five minutes on every managed asset and displays the last hour of data on the asset. To see utilization data for a network over longer periods of time, up to six months, create a Network Utilization chart, which includes operating system, operating system for a virtual machine, virtual host, and server pool. You can also create a network utilization chart for an OS group or host group.

Network Connectivity

Connectivity is the network interface of the system. You can view information about a hardware asset's Network Interface Card (NIC) on the Connectivity tab of the asset's dashboard, including name, connection status, MAC address, and the corresponding IP address.

For switch hardware, the Connectivity tab shows information about each port.

For an Oracle Solaris OS, the Connectivity tab includes IPMP groups and aggregated links.

- The IPMP Groups subtab shows the group's name, its assigned network, and the type of failure detection, either link-based, probe-based, or both. For each IPMP group, the details include the state of the connection for each NIC, whether it is in standby mode or failover mode, and the IP address the NIC supports.
- The Link Aggregation subtab shows the aggregation's name, its MAC address, and its attributes. For each aggregated link, the subtab shows the state of the connection for each NIC, whether it is in standby mode or failover mode, and the IP address the NIC supports.

When you attach or assign networks or when you create virtual hosts, Figure 7–7 shows an example of a step in the wizard where you configure the network connection.

Figure 7–7 Configure Interfaces

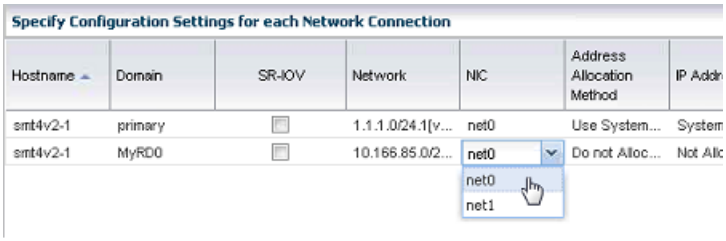

Network Hardware

- PCIe Endpoints
- Single Root I/O Virtualization
- Network Switches
- Virtual Network Switches

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches and Sun Datacenter InfiniBand switches. These switches reside in the system or blade system and provide the switch fabric.

The InfiniBand Gateway switch can expose the ports of a server that resides on an InfiniBand partition to an Ethernet network. To create an Ethernet on InfiniBand (EoIB) interface on the switch, you associate the switch's external port (eport) with the InfiniBand partition where the server resides, creating a virtual NIC (vNIC). The server's ports are displayed on the Switch Connectivity tab in the center pane.

For more information about these switches, see Switch Details or see Related Resources for Networks for links to the switch documentation.

PCIe Endpoints

A PCIe bus consists of the PCIe bus itself and all of its PCI switches and devices. Oracle VM Server for SPARC software can assign a PCIe bus (also known as a root complex) to a domain. An I/O domain that is configured with an entire PCIe bus is also known as a root domain.

Oracle Enterprise Manager Ops Center also supports the NIU-compatible cards in T5x20, T3 or T4 systems if an XAUI card is present. You can assign the NIU device to a Logical Domain in the same way you assign a PCIe bus or End Point.

Figure 7–8 PCIe and Buses

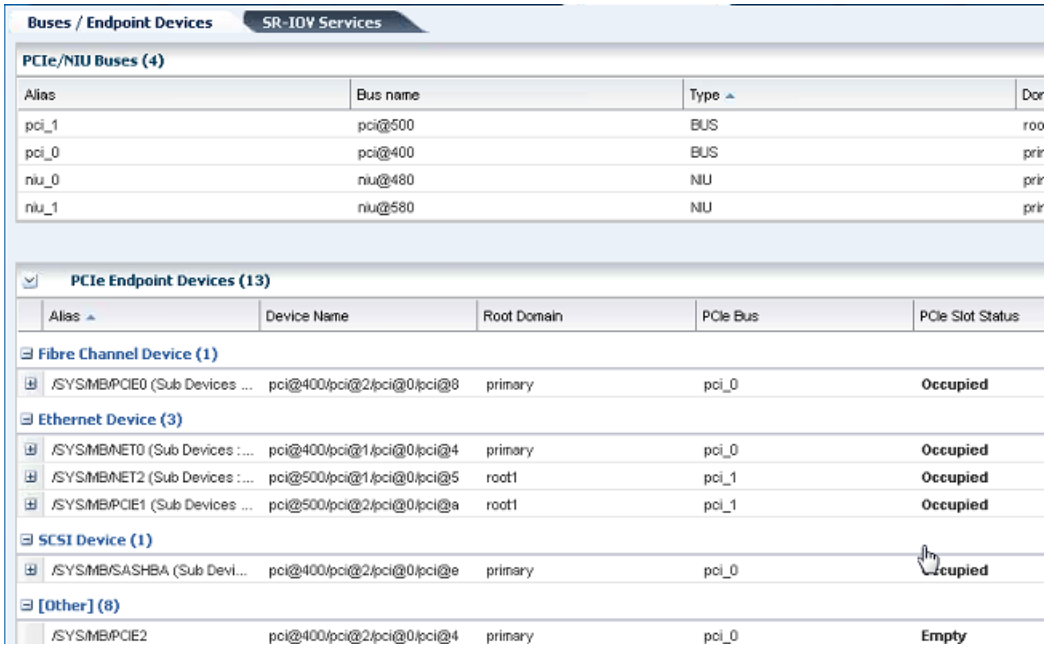

Single Root I/O Virtualization

InfiniBand switches support Single Root I/O Virtualization (SR-IOV), so that a single PCIe device (a physical network card) is presented as multiple PCIe devices. SR-IOV relies on both the hardware BIOS and the hypervisor layer to create these virtual PCIe devices. Each virtual PCIe device specializes in an operation called its virtual function (VF), but does not have the complete functionality of the physical PCIe device.

By defining a network on a virtual PCIe device, one physical PCIe device supports several networks as if each network had exclusive access to the device. Figure 7–9 shows the physical PCIe devices available to a virtualization host. When one of the physical functions is selected, its virtual functions are also displayed.

Figure 7–9 SR-IOV Tab

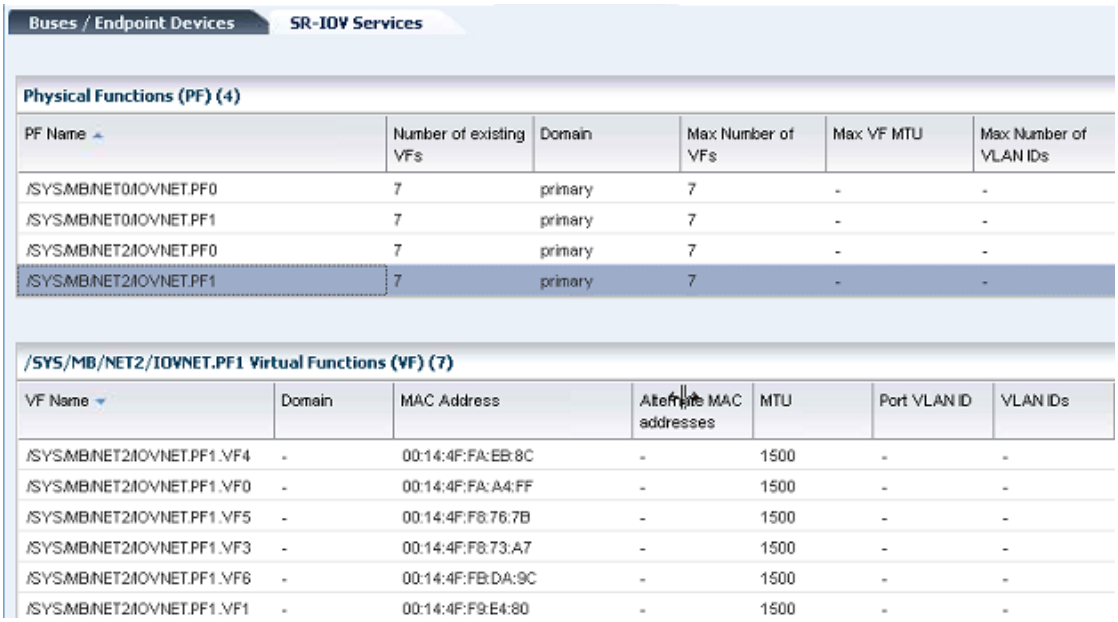

When you create a virtual host, you specify its network connection. If the networks are backed by an InfiniBand fabric, you can specify that the virtual host uses a virtual function by checking the SR-IOV option.

Network Interface Card (NIC)

The Network Interface Card (NIC) is the physical connection between a network switch and a network. When you create a network or attach an asset to a network, you select the NIC. You can create one network for each physical network interface card. To see the NICs for a server, select the server and then click the Connectivity tab. The Server Processor Connectivity table lists all of the NICs.

Network Switches

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches and Sun Datacenter InfiniBand switches. These switches reside in the system or blade system and provide the switch fabric. The Cisco Catalyst® 4948 Switch is also supported.

For the Ethernet switches, both tagged and untagged VLANs are supported.

If you use an InfiniBand switch in an Ethernet network, the ports on the switch have Ethernet names.

For more information about these switches, see the product documentation:

- For the Sun Ethernet 10GbE Fabric switch, see http://docs.oracle.com/cd/E19934-01/index.html
- For the Sun Network QDR InfiniBand Gateway Switch, see http://docs.oracle.com/cd/E36256_01
- For the Sun Datacenter InfiniBand Switch 36, see http://docs.oracle.com/cd/E26698_01

Virtual Network Switches

Oracle Solaris network virtualization provides an internal virtual network solution in which a virtualization host communicates with its virtual hosts as if using a network hardware. A virtual network consists of virtual network interface cards (VNICs) and virtual switches. A VNIC depends on a physical NIC and handles incoming and outgoing data in the same ways as a physical NIC. A virtual switch is created from the hypervisor layer of Oracle Solaris to provide the data path between the virtual hosts that reside on the same server and so must communicate with each other using the same ports. Figure 7–10 shows the relationship among the elements of a virtual network built on a single system with one NIC. Three VNICs support three zones. The virtual switch handles the communication, both between the VNICS and between the VNICs and the physical NIC.

Figure 7–10 Virtual Network

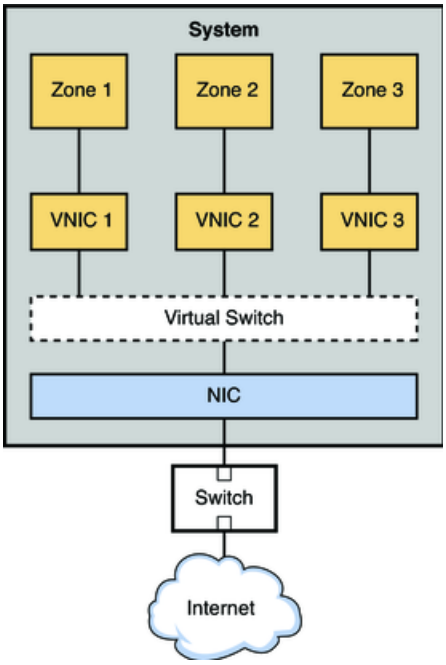

In Oracle Enterprise Manager Ops Center, the virtual network switches are listed on the Virtual Services tab, as shown in Figure 7–11. For a specific switch, you can also see VNICs that use the virtual switch.

Figure 7–11 Virtual Switches

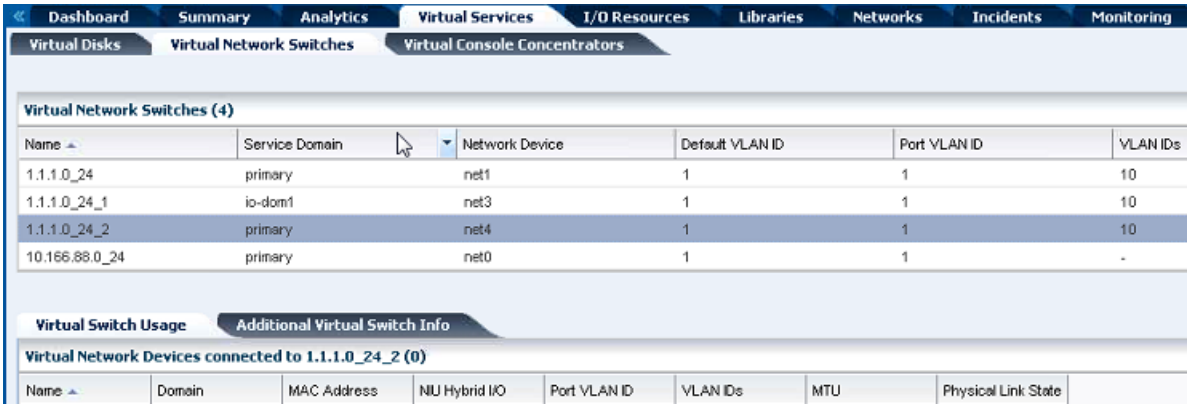

Network Profiles

Oracle Enterprise Manager Ops Center provides default profiles for the following operations:

- Monitor Network hardware – Reports Cisco switch's connection to assets on the Switch Connectivity tab.
- Discover a switch – Use a discovery profile with Cisco iOS credentials.

Oracle Enterprise Manager Ops Center's Networks

This section describes the requirements for the networks that Oracle Enterprise Manager Ops Center uses. This section does not discuss the networks that support virtual hosts and server pools.

In the product environment, a unique network is a combination of an IP subnet address and a subnet mask. You can implement Oracle Enterprise Manager Ops Center's network connections using any combination of VLAN-tagged and untagged networks.

- A network is considered VLAN-tagged when the interface is tagged at the Operating System. This is host-based tagging and the interface on the switch is the trunked interface.
- A network is considered untagged when the interface is untagged at the Operating System. This is port-based tagging because VLAN tagging is configured at the switch.

Network Switch Configuration

Use these guidelines to configure a network switch for a system running the Oracle Enterprise Manager Ops Center software.

- Use an Virtual LAN (VLAN)-capable switch.
- Discover and manage the switch. Provide the credentials to log in to the switch management controller's `ilom-admin` account. Do not use the `root` account.

Separate Networks for LOM Management, Enterprise Manager Ops Center, and Applications

Create a separate VLAN for asset management and provisioning networks, as shown in Figure 7–14. For Ethernet connectivity:

- The network used for managing a server's LOM must be a 10/100 MB connection to the server's net MGMT port. This is a requirement of the physical server.
- The network used for managing a server's LOM must be a 10/100 MB connection to the server's net MGMT port. This is a requirement of the physical server.

DHCP Servers on Proxy Controllers

Each Proxy Controller has a DHCP server process that is used for OS provisioning. This DHCP server process does not provide general DHCP services. Instead, this DHCP server responds to requests only from the specific MAC address of the asset that is being provisioned and only for the duration of the provisioning job. When OS provisioning an asset on a network that is not connected directly to the Proxy Controller, you must enable DHCP Forwarding on each intermediate network switch.

Types of Network Configurations for Oracle Enterprise Manager Ops Center

Simplest Configuration: Test System

Figure 7–12 shows a minimal configuration:

- Enterprise Controller (EC)
- Co-located Proxy Controller (PC), that is, the PC runs on the same system as the Enterprise Controller.
- Network connection (net0) can be a physical NIC, a link aggregate, or an IPMP group.

Figure 7–12 Configuration of a Test System

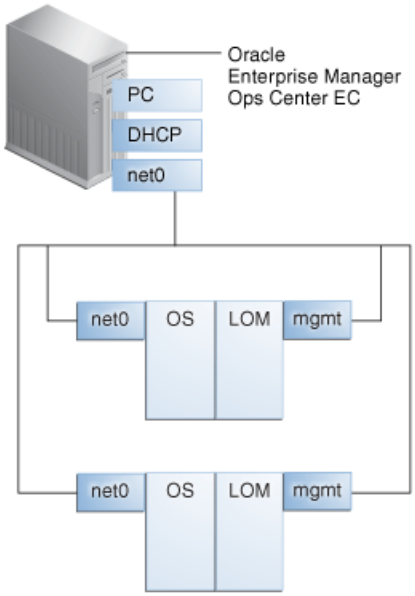

Simple Configuration: Datacenter on Same Network

Figure 7–13 shows a common configuration:

- Enterprise Controller (EC)
- Remote Proxy Controller (PC), that is, one or more Proxy Controllers running on different systems.
- Network connection (net0) can be a physical NIC, a link aggregate, or an IPMP group.

Figure 7–13 Configuration Using One Network

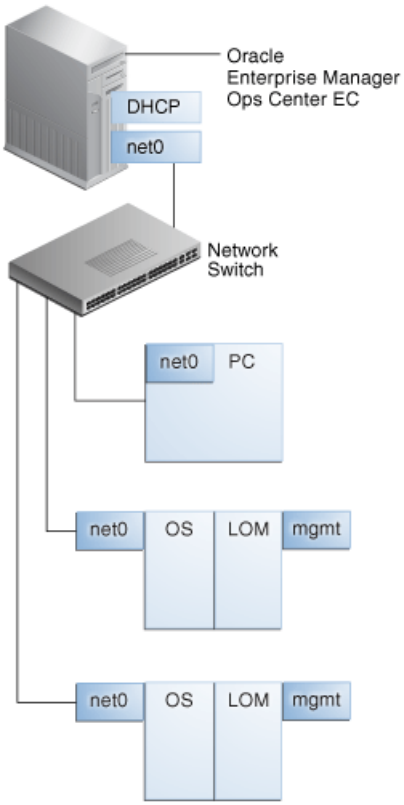

Good Practice: Separated Networks and Security

Figure 7–14 shows the preferred configuration:

- Enterprise Controller (EC)
- Remote Proxy Controllers (PC) in security zones.
- Firewalls protect security zones.
- Network connection (net0) can be a physical NIC, a link aggregate, or an IPMP group.

Figure 7–14 Configuration Using Separate Networks

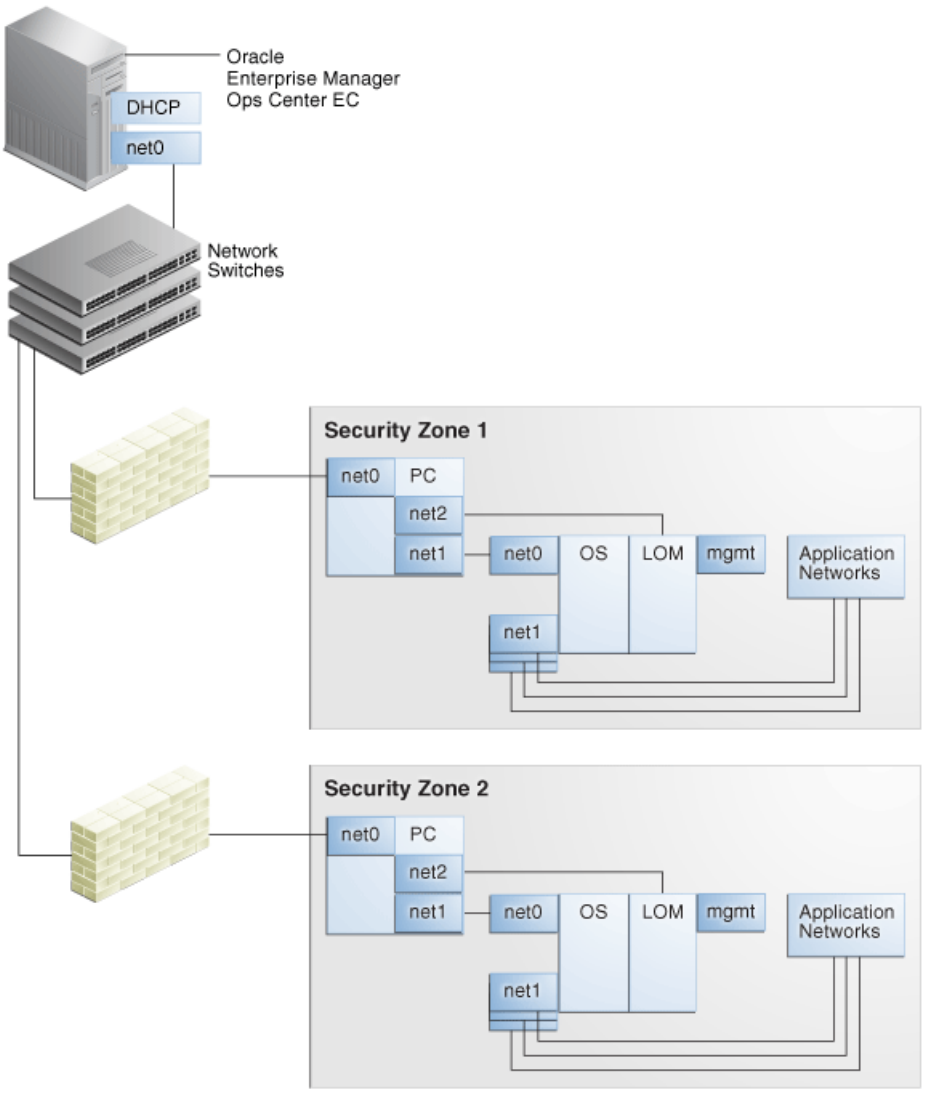

Related Resources for Networks

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- Chapter 17, "Networks for Virtualization"
- *Deploy Networks Workflow* in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm
- For the Sun Ethernet 10GbE Fabric switch, see http://docs.oracle.com/cd/E19934-01/index.html
- For the Sun Network QDR InfiniBand Gateway Switch, see http://docs.oracle.com/cd/E36256_01
- For the Sun Datacenter InfiniBand Switch 36, see http://docs.oracle.com/cd/E26698_01

8

Plans and Profiles

This chapter provides an overview of the concepts of operational plans, deployment plans, and profiles. Detailed information is covered in the corresponding feature chapters.

The following information is included:

Introduction to Plans and Profiles

Oracle Enterprise Manager Ops Center uses a combination of plans and profiles to reduce complexity and increase consistency when you perform standard management and operational activities, such as configuring hardware, installing servers, updating operating systems, and creating virtual systems.

A plan defines the actions and the targets. A profile defines a task, how the task is performed and enables you to define what is allowed, and not allowed, to be installed on a system. Together, plans and profiles enable you to create a reusable set of procedures to perform tasks, such as configuring hardware, upgrading firmware, installing and patching operating systems, and creating virtual systems and guests.

You create, manage, and access all plans and profiles in the Plan Management section of the user interface. Plan Management contains three basic components:

- Profiles and Policies: Profiles define the configuration of components for a specific type of system and task, such as the naming schema and configuration options to use when creating a zone or logical domain. Update policies define the level of interaction you want when applying patches and packages.

- Deployment Plans: Perform standard management activities. A Deployment Plan provides a framework of steps that you need to complete one or more tasks. You customize the plan to include specific profiles for the steps. Complex deployment plans enable you to add operational plans as a step. When you apply a Deployment Plan, you select one or more targets, or group of targets, on which to complete the tasks.
- Operational Profiles and Plans: Perform one or more operational activities, such as scripts and utilities to fix common problems, monitoring rule configurations and thresholds, and a knowledge base that you can create based on the incidents that occur in your environment.

Roles for Plans and Profiles

Table 8–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 8–1 Plan Management Roles and Permissions

Task	Role
View a profile or plan	Read
Create a profile or plan	Profile Plan Administrator
Edit a profile or plan	Profile Plan Administrator
Copy a profile or plan	Profile Plan Administrator
Delete a profile or plan	Profile Plan Administrator
Create a profile or plan	Profile Plan Administrator
Apply a plan	The plan determines the role required. See Deployment Plans for links to the chapters that go into more detail about specific plans.

Actions for Plans and Profiles

Oracle Enterprise Manager Ops Center uses plans and profiles to perform many tasks in your data center.

The following types of plans and profiles are available:

- Deployment plans and Profiles: Several types, or categories, of profiles are available. Deployment plans and profiles perform a variety of tasks, including discovering and adding assets to the UI, deploying or updating operating systems and firmware, creating zones, and creating logical domains.
- Operational Plans and Profiles: Operational profiles are a specific type of profile that you can create to store a user-defined script. The associated operational plan deploys the script to selected targets.

You can create, copy, edit, and delete profiles and create, copy, edit, deploy, and delete plans. In addition to deploying operational plans on their own, you can add an operational plan as a step in a complex deployment plan. See Operational Plans and Profiles, Profiles and Policies, and Deployment Plans for more information about the actions you can perform.

Location of Plans and Profiles in the User Interface

Profiles, plans, and templates are available in the Plan Management section of the UI.

Table 8–2 Location of Profile and Plan Information in the UI

Object	Location
Deployment Plans	Expand the **Plan Management** section of the Navigation pane. Click **Deployment Plans** for a list of existing plans. To locate a specific plan, expand the folder for the plan you want, such as Configure Server Hardware.
Profiles	Expand the **Plan Management** section of the Navigation pane, then scroll down to Profiles and Policies. Click **Profiles and Policies** for a list of existing profiles. Expand the folder for the type of profile you want, such as RAID Controller.
Update Policies	Expand the **Plan Management** section of the Navigation pane, then scroll down to Profiles and Policies. Click **Profiles and Policies**, then scroll to the bottom of the section to view the OS Update Policies. Click **Update Policies** to list the existing OS Update Policies in the center pane.
Update Profiles	Expand the **Plan Management** section of the Navigation pane, then scroll down to Profiles and Policies. Click **Profiles and Policies**, then scroll to the bottom of the section to view the OS Update Profiles. Click **Update Profiles** to list all default and user-defined OS Update Profiles in the center pane.
Operational Plans	Expand the **Plan Management** section of the Navigation pane, then scroll down to Operational Plans. Expand **Operational Plans** for a list of user-defined operational plans.
Operational Profiles	Expand the **Plan Management** section of the Navigation pane, then scroll down to Operational Plans. Expand **Operational Plans**, then click **Operational Profiles** for a list of user-defined operational profiles.
Monitoring Policies	Expand the **Plan Management** section of the Navigation pane, then scroll down to Operational Plans. Expand **Operational Plans**, then expand **Monitoring Policies** to display a list of system-defined and user defined monitoring policies. System defined policies all begin with OC.

Overview of Version Control

Profile and plan versions are numbered sequentially. When you create a profile or plan, the version number is one. When you edit a profile or plan, you create a version that is referenced by a new number.

When you create a version, you have the option to automatically change the related plan during the update. When you do not choose that option, the operational and deployment plans that are using the profile are not updated with the new version. For example, you have a firmware profile that is using firmware image A, the profile version is one. You want to update the profile to use the latest firmware image, B. When you update the profile to change the image to B, version two of the profile is created. When you have five plans that are using version one of the firmware profile, all of those plans continue to use version one and image A. When you want the plans to use image B, you must manually update the plans to reference version two of the profile.

When you edit a deployment plan that is referred to by another plan in a complex plan, the referring plan is not automatically updated to refer to the edited plan's version unless you choose the option when you edit the plan.

You can view version details by highlighting the plan, then clicking the View Version Details icon. Use the arrows to view the different versions. When you want a plan to use the new version of the profile, edit the plan and associate it with the correct version of the profile. You can edit the following fields in an Operational Plan: Description, Failure Policy, and Associated Profile. You cannot edit the name of a

profile or plan. When you want a different name, you must copy or create a profile and plan.

You can delete a version of a profile or plan. When you delete a version and more than one version exists, the previous version becomes the default. When only one version of the plan exists, both the profile and plan are deleted. When the plan is used in other plans, then the Delete Deployment Plan option is not enabled.

Operational Plans and Profiles

Operational profiles are designed to assist you in the day-to-day operation of your data center.

An Operational Profile performs one or more tasks needed to operate your environment. The profile contains a single shell script and can include asset attributes as environmental variables. For each profile, you can choose one of the following types of shell scripts:

- **EC Shell**: The script runs only on the Enterprise Controller and is executed with the logged-in user's credentials.
- **Remote Shell**: The script can run on any managed system that contains a remote agent and is executed with root permissions.

You have the option of adding two types of variables to operational profiles:

- System defined variables: Use these variables to define specific information in the script. The following are software-specific variables:
  - $OC_TARGET_NAME
  - $OC_JOB_ID
  - $OC_SASL_FILE
  - $OC_UFN
  - $OC_TARGET_TYPE
- Additional Variables: Use these variables to add information to the script and enable users to change that information when they execute the operational plan. Additional variables are user-defined.

A list of supported asset attributes is available in *Oracle Enterprise Manager Ops Center Feature Reference Appendix Guide*.

You must have an Operational Plan to execute an Operational Profile. When you create an Operational Profile, the default action also creates an Operational Plan. You use the plan to execute the profile on a managed resource or group of resources, such as performing state changing actions.

You can create a simple plan, such as disabling print capabilities, and deploy the plan across a group of resources in your data center. A more sophisticated example is to create several Operational Plans and add them as steps in a complex type Deployment Plan. For example, you can create an Operational Plan to shut down all logical domains and another Operational Plan to shut down an Oracle VM Server for SPARC. You can then add these plans as steps in a complex Deployment Plan.

You can perform the following actions, depending on the role:

- Creating an Operational Profile and Plan
- Editing an Operational Profile or Plan

Creating an Operational Profile and Plan

An Operational Plan defines how an Operational Profile is deployed, and against which targets. By default, creating a profile also creates an Operational Plan.

You can save a shell script on the Enterprise Controller and download it into the plan, or you can enter the script in a field when you create the plan. Both types of shell scripts are executed by the user. They differ in the location, either on the Enterprise Controller or on the remote Agent, and the user credentials needed to execute the profile.

Note: The uploaded shell script file cannot exceed 2 GB.

An operational plan defines the targets and failure policy for an operational profile. The profile defines one or more operations that are to be performed on a managed resource or group of resources. For example, deploying thresholds onto a managed resource, or performing state changing actions such as shutting down all logical domains and then shutting down an Oracle VM Server for SPARC.

The profile uses a shell script to define the operations. When you do not have an operational profile, you can create a profile when you create the plan. When you create the profile, either download a shell script that is on the Enterprise Controller (EC Shell script), or enter the shell script (Remote Shell script) in the profile. The EC Shell and Remote Shell are both shell scripts that are executed by the user. They differ in the location (on the Enterprise Controller or on the remote Agent) and manner (user credentials) of the execution. The EC Shell is executed with the credentials of the user that is logged in. The Remote Shell can run on any managed system that contains a remote agent and is executed with root permissions.

To Create an Operational Profile and Plan

Perform the following steps to create an operational profile and plan:

1. Expand **Plan Management** in the Navigation pane, then click **Operational Profiles**.
2. Click **Create Profile**.
3. Name the new profile and add a description. Select a subtype from the list to identify the type of target for this profile. Click **Next**.
4. Define the script, then click **Next**.
   a. Define the type, either **EC Shell** or **Remote Shell**.
   b. Define the time out parameters.
   c. Enter the script in the Script field, or click **Load** if you are loading an EC Shell script that is saved on the Enterprise Controller. Click **View System Variables** to see the available software-specific variables.
5. (Optional) Add user-defined variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Editing an Operational Profile or Plan

When you edit a profile or plan, you create a version. When you create a version, you have the option to automatically change the related plans during the update. See Overview of Version Control for how versions are maintained.

You can view version details by highlighting the plan, then clicking the View Version Details icon. Use the arrows to view the different versions. When you want a plan to use the new version of the profile, edit the plan and associate it with the correct version of the profile. You can edit the following fields in an Operational Plan: Description, Failure Policy, and Associated Profile. You cannot edit the name of an Operational profile or plan. When you want a different name, you must copy or create a Operational Profile and Plan.

Copying an Operational Profile

You can copy an existing operational profile, rename it, and create a new profile and plan.

To Copy an Operational Profile

1. Expand **Plan Management** in the Navigation pane, then click **Operational Profiles**.
2. Select the profile in the center pane that you want to copy, then click the **Copy Profile** icon. Or you can double-click the profile in the center pane to display actions in the Actions pane.
3. Rename the new profile you are creating and revise the description, then click **Next**.
4. (Optional) Edit the script. Click **Next**.
5. (Optional) Edit the variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Copying an Operational Plan

You can copy an existing operational plan, rename it, and create a new profile and plan.

To Copy an Operational Plan

1. Expand **Plan Management** in the Navigation pane, then click **Operational Plans**.
2. Select the plan in the center pane that you want to copy, then click the **Copy Plan** icon. Alternatively, double-click the plan in the center pane to display actions in the Actions pane.
3. Rename the new plan you are creating and revise the description. You can choose to change the Failure Policy for the new plan and change the Operational Plan steps. You can change the associated profile or plan for each step. To change the profile or plan, click the associated profile or plan to select from a list of available options. To add additional steps, click the **Replicate Step** icon, then select the profile or plan from the list to associate with that step. Click **Save**.
4. (Optional) Edit the script. Click **Next**.
5. (Optional) Edit the variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Deleting an Operational Profile or Plan

You can delete a version of a profile, or you can delete all versions of the profile and the associated plan.

You cannot delete a version of an operational plan version unless you created the version. Deleting a version of a plan might impact the Incident Knowledge Base or deployment plans that reference the version. Before deleting a version, verify that the version is not being used.

Note: When you delete a version and more than one version exists, the previous version becomes the default. When only one version of the plan exists, the operational profile and plan are deleted.

Profiles and Policies

Profiles and policies define how a job is performed and the level of interaction.

The following information is discussed in this section:

About Profiles

Profiles define how standard management tasks are performed. With the help of wizards, you create a customized set of profiles. Once created, authorized users can use the profiles to perform tasks, such as discovery and provisioning. The software has pre-defined profiles for some common OS tasks, such as reboot and check for installed security patches. You can choose to use the pre-defined profiles for those actions or create your own.

Profiles appear in the Plan Management section of the UI and are organized by category. The profiles are discussed in more detail in various sections of the documentation. The following are the profile categories and where you can find more specific information about that task and profile:

- Discovery: See Chapter 2, "Asset Management"
- Service Processor: See Chapter 11, "Hardware"
- RAID Controller: See Chapter 11, "Hardware"
- Firmware: See Chapter 11, "Hardware"
- Dynamic System Domain: See Chapter 11, "Hardware"
- OS Provisioning: See Chapter 13, "Operating System Provisioning"
- OS Configuration: See Chapter 13, "Operating System Provisioning"
- Logical Domain: See Chapter 19, "Oracle VM Server for SPARC"
- Oracle Solaris Zone: See Chapter 18, "Oracle Solaris Zones"

- Virtual Machine: See Chapter 20, "Oracle VM Server for x86"
- BIOS Configuration: See Chapter 11, "Hardware"
- Boot Environments: See Chapter 12, "Operating System Management"
- Storage Appliance Update: See Chapter 11, "Hardware"
- Update Policies: See Chapter 14, "Operating System Updates"
- Update Profiles: See Chapter 14, "Operating System Updates"

With the exception of the Discovery profile, these types of management profiles require a deployment plan to execute the tasks on specific targets. When you create a profile, the default setting is to create a corresponding deployment plan. See Deployment Plans for more information on these types of plans.

Viewing Profile Details and Associated Plans

Details for each profile, including the associated plans and version number are readily available. To display the information in the center pane, expand Plan Management, then select the profile from the list of Profiles and Policies. The following tabs appear in the center pane:

- Details
- Referrers
- Version History

Details

The **Details** tab displays the profile configuration details for the current version. As shown in Figure 8–1, the details include Name, Description, Target Type, Subtype, Version, and date and time last modified. The Profile Details and File Systems sections include the wizard selections and settings for this version of the profile.

Figure 8–1 Profile Details

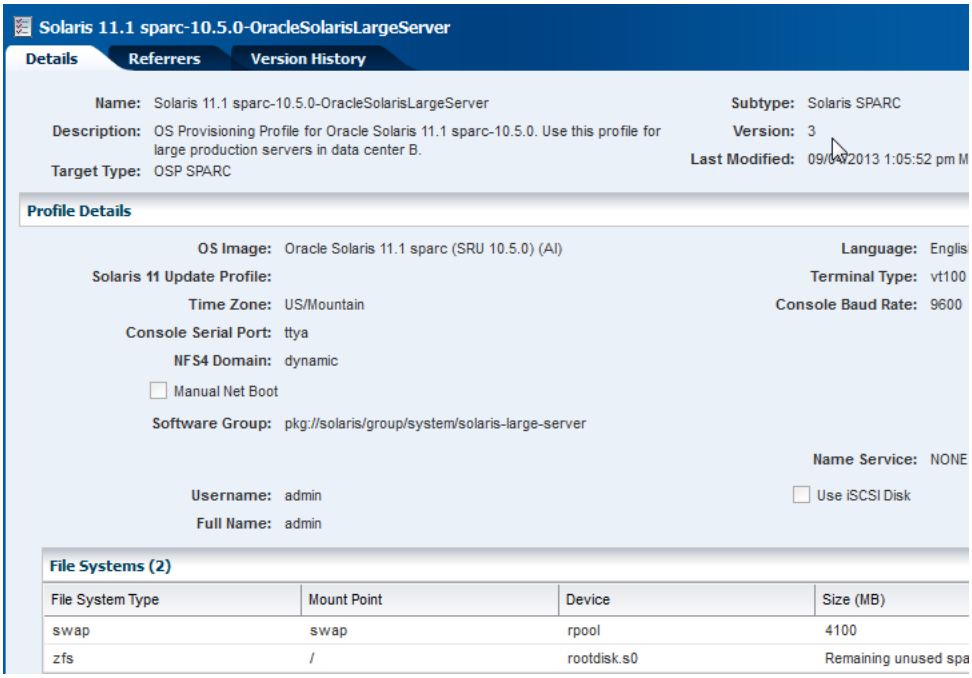

Referrers

The **Referrers** tab displays all deployment plans that use the profile and the profile version number used. For example, in Figure 8–2 two deployment plans use the profile. The first deployment plan listed uses version 3 of the profile and the second deployment plan listed uses version 1 of the plan. This occurred because the profile was edited after the profile was part of the *S11.1 SPARC Large* deployment plan. Because the profile was part of the plan, the plan automatically updated to use the latest version of the profile. When the *S11.1 SPARC Large - Data center A* deployment plan was created, three versions of the profile existed and the user had the option of which version to associate with the plan. In this case, version 1 of the profile was selected to be part of the plan.

Figure 8–2 Profile Referrers

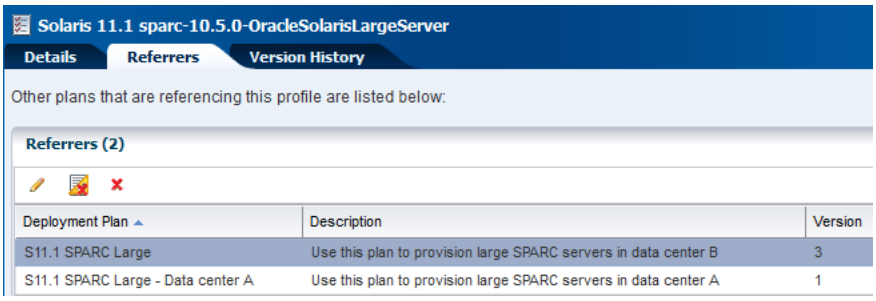

Version History

The **Version History** tab shows the number of versions, the profile description, the target type, and when the version changed. Click the view icon to display the profile details for that version. Figure 8–3 shows a Solaris 11 OS Provisioning profile with three versions. By clicking the view icon in the upper left corner, you can display the details for the selected profile.

Figure 8–3 Profile Version History

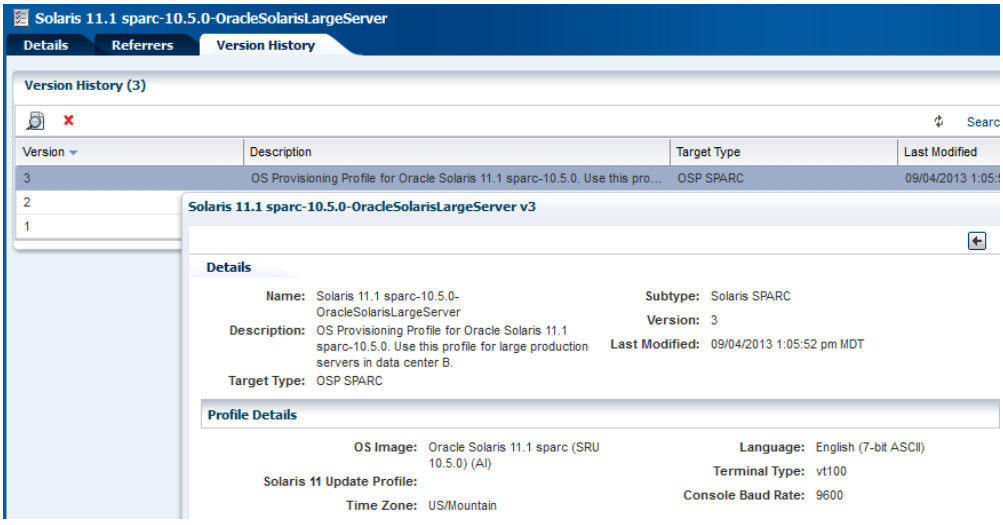

Creating a Profile

Profiles define the core information for a task. The profile categories appear in the Plan Management section. You create a profile by completing the fields in a wizard for the specific type of profile.

When you import images into the software library, a default OS profile is created for the image. You can use the default profile, copy the default profile to create your own profile, or create a new profile. To preserve the default profile configuration, you might want to copy a profile and provide a name and description that describes the profile and any unique properties for the profile.

The prerequisites and steps to create a profile differ, depending on the type of profile you want to create.

To Create a Profile

The following steps are an overview of how to launch the create profile wizard:

1. Expand **Plan Management** in the Navigation pane.
2. Expand the **Profiles and Policies** navigation tree, then select a profile.
3. Click **Create Profile**.
4. Complete the wizard, then click **Finish**.

Copying a Profile

Copying a profile enables you to create a new profile with the same configuration, and then edit the profile to make a unique profile. For example, you might want to copy a default operating system profile to create a user-defined profile and retain the default profile as a template.

To Copy a Profile

The following steps are an overview of how to launch the copy profile wizard:

1. Expand **Plan Management** in the Navigation pane.
2. Expand the **Profiles and Policies** navigation tree, then select a profile.
3. Click **Copy Profile**.

   The Create Profile - OS Provisioning Wizard appears.
4. Edit the name and description of the profile.
5. Edit any parameters that you want to change.
6. Click **Finish**.

Editing a Profile

Editing a profile enables you to change the profile configuration. Changing values other than the name or description creates a new version. When a profile is referenced by one or more deployment plans, the default action is to update the referring plans to the new version. Editing the name or description of the profile creates a new profile. Before editing a profile or plan, see Overview of Version Control for how versions are maintained.

To Edit a Profile

Perform the following steps to edit a profile:

1. Expand **Plan Management** in the Navigation pane.
2. Expand the **Profiles and Policies** navigation tree, then select a profile.
3. Click **Edit Profile**.
4. Complete the changes, then click **Finish**.

About Policies

Policies are lists of actions that are explicitly approved or denied. Policies are specific to OS update and define the amount of user interaction you want when applying OS patches and packages. For example, you might want to automate applying fixes without user interaction, but you might want to pause a job and require user approval before performing a downgrade or uninstall. You can also deny certain actions, such as installing patches or packages that are not certified.

Policy settings are hierarchical. When you have not defined a policy for a component, the policy for that component's parent applies. For example, it is possible to create a policy that allows the system to install a given component, but prohibits installing specific versions of that component.

For more information about update policies, see Creating Update Policies.

Deployment Plans

Deployment Plans use profiles to perform standard management activities in a consistent and repeatable manner. A deployment plan defines the sequence of operations or steps that must be carried out on an asset to deploy it together with the specification or profile that each step applies and the resources that are required to apply it such as network addresses and system names.

Deployment plans are all based on a set of pre-defined templates. Some templates are designed for simple tasks and other, more complex plans, are designed to perform a series of tasks. A comprehensive set of deployment templates is available for you to use to create a variety of deployment plans. Each template is an unbound deployment plan which defines the steps of execution, but not the profiles and assets.

A deployment template enable you to define a task and the associated resources (such as images and network addresses), that are needed to complete the task. You can use the templates and customized profiles to create plans for your data center. For each plan, you can define the course of action to take when a step cannot be completed on a target. You can choose to stop the job at the first failure, or attempt to complete as much as possible. In some cases, you can build complex plans by combining existing plans. For example, you might add an operational profile and plan to the end of a deployment plan.

Several plans contain more than one step. In some cases, you can associate steps in a plan with another plan. The associated plan is referred to as a nested plan. You can use nested plans as shared building blocks, much in the same way as profiles are used. Configuring a single nested plan once and reusing it in many other plans reduces the number of individual operations that you must complete in the UI. The templates appear in the user interface in alphabetical order, not by type. However, it is useful to categorize the plans as simple, multi-step, and complex. See the different categories for a description of each plan.

The following lists the plans that are located in each category:

- About Simple Deployment Plans

  - Configure RAID
  - Configure Service Processor
  - Update Firmware
  - Update BIOS Configuration
  - Update Storage Appliances
  - Create Oracle Solaris Zones
  - Create Virtual Machines
- About Multi-Step Deployment Plans
  - Install Server
  - Create Dynamic System Domain
  - Create Logical Domains
  - Provision OS
  - Software Deployment/Update
  - Update Solaris 11 OS
  - Create Boot Environment
  - Update Firmware and Install Oracle VM Server for SPARC
- About Complex Deployment Plans
  - Configure M-Series Hardware, Create, and Install Domain
  - Configure Server Hardware and Install OS
  - Configure and Install Dynamic System Domain
  - Configure and Install Logical Domains
  - Configure and Install Oracle Solaris Zones
  - Configure and Install Virtual Machines

Use the templates to create your own deployment plans and configure the plans using profiles. The settings and values in the profiles bound to each step are defaults. You can modify the plan before it is actually applied. You can further constrain the profile settings and values by the target systems to which the plan is applied.

With the exception of complex plans, the templates do not allow you to add steps to a plan. You can use only those steps that are defined in the template from which the plan is derived. Complex plans enable you to add one or more deployment and operational plans inside a complex type plan.

About Simple Deployment Plans

The software provides you with the ability to create, configure, manage and execute deployment plans which drive the hardware, firmware and software provisioning activities in a repeatable fashion.

You create plans from defined templates. Each plan defines the sequence of steps that must be carried out for configuration or provisioning of a system. Plans might contain a single step or a sequence of multi-steps. Each step in the plan is configured by associating a profile or another plan.

A simple plan contains a single step with a single image. You can define the image used by the plan, but you cannot add more than one image or add steps. In most cases, when you create a profile, the default action also creates a simple deployment plan.

The following plans are simple deployment plans:

- Configure RAID: Use this plan to configure the RAID controller on a server. See Configuring a RAID Controller for how to use the profile and plan.
- Configure Service Processor: Use this plan to configure the service processor on a chassis. See Configuring the Service Processor for how to use the profile and plan.
- Update Firmware: Use this plan to update firmware. See Firmware Provisioning for how to use the profile and plan.
- Update BIOS Configuration: Use this plan to update the BIOS configuration of servers. See Creating a BIOS Configuration Profile and Plan for how to use the profile and plan.
- Update Storage Appliances: Use this plan to update storage appliance software. See Chapter 16, "Storage Libraries for Virtualization" for how to use the profile and plan.
- Create Oracle Solaris Zones: Use this plan to create zones. See Creating a Zone Profile and Creating and Deploying Zone Plans for how to use the profile and plan.
- Create Virtual Machines: Use this plan to create virtual machines for Oracle VM Server for x86. See Create Virtual Machines and Deploying Virtual Machine Plan for how to use the profile and plan.
- Create Dynamic System Domain: Use this plan to create dynamic system domains. See Configuring a Dynamic System Domain for how to use the profile and plan.
- Create Logical Domains: Use this plan to create logical domains. See Creating a Guest Domain Profile for how to use the profile and plan.

About Multi-Step Deployment Plans

Several deployment templates with multi-step sequences are available. These plans are designed to provide you with a customized and repeatable way to perform common operations with a single click. The steps in a multi-step plan are defined in the template. You cannot skip steps or add steps.

The following plans are multi-step deployment plans:

- Install Server: Use this plan to provision the server and update the OS.
- Software Deployment and Update: Use this plan to apply script based update profiles. See Updating an Operating System From a Deployment Plan for how to use the profile and plan.
- Provision OS: Use this plan to provision and configure an operating system. See Operating System Provisioning for how to use the profiles and plan.
- Update Solaris 11 OS: Use this plan to update Oracle Solaris 11 operating systems. See Updating an Operating System From a Deployment Plan for how to use the profile and plan.
- Create Boot Environment: Use this plan to create Oracle Solaris boot environments. See Creating an Oracle Solaris 11 Boot Environment and Creating an Oracle Solaris 10 Boot Environment for how to use the profile and plan.

- Update Firmware and Install Oracle VM Server for SPARC: Use this plan to update firmware and then install Oracle VM Server for SPARC on the system. See Overview of Oracle VM Server for SPARC Installation for how to use the profile and plan.

About Complex Deployment Plans

You can use a combination of profiles, deployment plans, and operational plans to create a complex deployment plan that enables you to automate a variety of detailed workflows into a single plan. Complex plans provide flexibility to structure plans that meet your local requirements, increasing consistency and allowing for a greater level of automation.

When you create complex deployment plans, you can choose to skip a step in the plan. Skipped steps are not processed when the plan is applied. You can replicate certain steps to perform the same operation but using a different profile or plan. You can also add one or more deployment plans and operational plans.

The following plans are complex deployment plans:

- Configure M-Series Hardware, Create and Install Domain: Use this plan to configure an M-Series server, create dynamic system domains, provision OS on the domains, and update the domains.
- Configure Server Hardware and Install OS: Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- Configure and Install Dynamic System Domain: Use this plan to create dynamic system domains, provision and update OS on the domains.
- Configure and Install Logical Domains: Use this plan to create logical domains and provision OS on the logical domains.

Managing Deployment Plans

The following management actions are available:

- Copying a Deployment Plan
- Editing a Deployment Plan
- Deleting a Deployment Plan

Copying a Deployment Plan

You can copy an existing deployment plan, rename it, and create a new plan.

To Copy a Deployment Plan

Perform the following steps to copy a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select **Copy Deployment Plan**:
   - Method 1: Select the deployment type from the tree and select a plan from the list to enable the Copy Deployment Plan icon. Click the **Copy Deployment Plan** icon.
   - Method 2: Expand the selected deployment type and select a plan from the list. Select **Copy Deployment Plan** from the Actions pane.
3. Edit the following details of the plan:

- Description: Provide a description of the plan.
- Plan Name: By default, the plan name is Copy of <*plan name*>. For example, Copy of Firmware Update. You can modify the name.
- Failure Policy: Select whether the plan execution is to stop at failure or complete as much as possible.

4. Configure a step of the plan by setting or changing the associated profile or by creating a new profile.
5. Edit the plan by replicating the steps and associate targets, depending on the type of plan selected.
6. Click **Save the new plan**. A new plan is saved with the version v1.

Editing a Deployment Plan

You can edit the deployment plan details, alter the plan configuration by skipping steps in the plan, change the profile or plan bound to each step, or save the plan under a different name to create a new plan.

Note: When you edit a deployment plan that is referred to by another plan, for example, in a nested plan, the referring plan is not automatically updated to refer to the edited plan's version. You must manually modify the referring plan if you want it to use the modified version.

Before editing a plan, see Overview of Version Control for how versions are maintained.

To Edit a Deployment Plan

Perform the following steps to edit a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select the **Edit Deployment Plan** option:
   - Method 1: Select the deployment type from the tree and select a plan from the list to enable the icon. Click the **Edit Deployment Plan** icon.
   - Method 2: Expand the selected deployment type and select a plan from the list. The plan details appear. Select **Edit Deployment Plan** from the Actions pane.
3. Edit the following details of the plan:
   - Plan Name: Edit the name to create a new plan. You create a new plan when you modify the plan name.
   - Description: Provide a description of the plan.
   - Failure Policy: Select whether you want the plan execution to stop at failure or complete as much as possible.
4. Configure a step of the plan by setting or changing the associated profile, or by creating a new profile.
5. (Optional) Edit the plan by replicating the steps and associate targets depending on the type of plan selected.

6. Click **Save** to save any changes made to the plan. When you have changed the name, a new plan is saved with the version v1.

Deleting a Deployment Plan

You can delete a deployment plan or only a version of the plan. When the selected deployment plan is not referenced by any other plans, you can confirm deleting the plan or its version. When the plan is used in other plans, the Delete Deployment Plan option is not enabled.

To Delete a Deployment Plan

Perform the following steps to delete a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select the Delete Deployment Plan option:
   - Method 1: Select the deployment type from the tree. The plans of that type are listed in the center pane. Select a plan from the list. The Delete Deployment Plan and Delete Version icon is enabled. Click **Delete Deployment Plan** or **Delete Version**.
   - Method 2: Expand the selected deployment type and select a plan from the list. Select **Delete Deployment Plan** or **Delete Version** from the Actions pane.
3. Click **Delete** to confirm the delete action.

Applying a Deployment Plan

When you deploy a plan, you select the target assets against which the plan is executed. After you select the targets, you have the option to temporarily override the profile configuration for specific assets before you deploy the plan.

Many actions use deployment plans. See the documentation in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm and the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm for workflows and end-to-end examples.

Related Resources for Plans and Profiles

For profile and plan details, and how to use individual profiles and plans, go to one of the following resources.

- Chapter 2, "Asset Management"
- Chapter 11, "Hardware"
- Chapter 12, "Operating System Management"
- Chapter 13, "Operating System Provisioning"
- Chapter 14, "Operating System Updates"
- Chapter 18, "Oracle Solaris Zones"
- Chapter 19, "Oracle VM Server for SPARC"

For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm and the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.

Part III

Operate and Maintain

Part III contains the following chapters:

9

Incidents

This chapter describes how you can use the software to identify, assign, and resolve incidents.

The following information is included:

Introduction to Incidents

When an asset is not operating within the parameters defined in the monitoring rules and policies, Oracle Enterprise Manager Ops Center generates an alert and an incident.

An alert indicates that a monitored asset is not performing as expected. The monitoring rule parameters determine when an alert is triggered and the severity: Informational (info), Warning, or Critical.

An incident is raised when the monitoring rule is asserted by the raising of one or more alerts that the monitoring rule requires. New alerts will update an open incident. One or more subsequent alerts trigger the monitoring rule, which notifies the incident management system. The incident management system detects that there is already an open incident for the monitoring rule for that asset and correlates the alerts under the open incident and the worst severity level is associated with the incident. For example, when the incident is at a Critical severity level and a new Warning alert is added to the incident, the incident severity remains at the Critical level.

Oracle Enterprise Manager Ops Center uses a help desk approach to manage the incidents in your data center. All open incidents appear in the Message Center. You can assign incidents to others for resolution, add comments, provide a list of possible causes and impacts, provide recommendations, add utilities or scripts to resolve an issue, view progress, and open a service desk ticket.

The following are the main components that help you to track and manage known issues on your monitored assets:

- Using the Message Center: Central location for details on incidents, notifications, and service requests.
- Using Annotations: Comments, suggested actions, and operational plans that enable you to effectively manage an incident.
- Building an Incidents Knowledge Base: Database of your annotations and actions for specific types of incidents and severity levels.
- Using Oracle Services and Service Requests: Service request details on all requests submitted to My Oracle Support through ASR or the Oracle Enterprise Manager Ops Center UI. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about Auto Service Requests and how to enable them.

When an incident appears, you can assign it to a user for resolution and use annotations to add comments and suggested actions. You can build an Incident Knowledge Base that contains your annotations from specific incidents, or add suggested fixes or automated fixes for a specific type of incident.

If you cannot resolve an incident, you can open a service request with My Oracle Support inside the Oracle Enterprise Manager Ops Center UI. The information gathered by the software is automatically populated into the service request. You can track the status of any service request opened from within Oracle Enterprise Manager Ops Center, whether they belong to you or one of your co-workers.

Roles for Incidents

The following table lists the tasks that are discussed in this section and the role required to complete the task. You can restrict privileges to specific targets or groups of targets. Contact your administrator when you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 9–1 Incident Management Tasks and Roles

Task	Role
View Incidents	All
Assign Incidents	Fault Administrator
Add Annotation to incidents	Fault Administrator
Acknowledge incidents	Fault Administrator
Take Actions on Incidents	Fault Administrator
Mark Incidents as Repaired	Fault Administrator
Close Incidents	Fault Administrator
Take Actions on Notification	Fault Administrator
Delete Notifications	Fault Administrator

Actions for Incidents

Incident management is automatically enabled after an asset is discovered and managed. Alerts and incidents are generated based on the monitoring profiles and policies.

Actions for Incidents

You can perform the following actions for a specific incident:

- View Alerts: View all alerts that are generated when the state is outside the defined monitoring parameters.
- View Annotations: View the scripts or suggested actions that you or others in your organization have associated with an incident.
- View Possible Impacts and Causes: View the possible impacts and causes for an incident.
- View Comments: View comments that you or others in your organization have associated with an incident.
- View Suggested Actions: View suggested actions that you or others in your organization have associated with an incident.
- Add Annotation to the Incident: Add a script, suggested action, or comment to a specific incident.
- Assign the Incident: Assign an incident to a user.
- Acknowledge the Incident: Acknowledge, or accept, and incident that has been assigned to you.
- Take Actions on the Incident: Take action to resolve an incident. Options include executing a suggested action, executing a script that is part of an operational profile, executing a command, or executing a script.
- Mark Incident as Repaired: Identifies the incident as being fixed.
- Close the Incident: Closes the incident and removes it from the Message Center.
- Open a Service Request: Open a service request for the incident with My Oracle Support.

Actions for Service Requests

When the Enterprise Controller is connected to My Oracle Support and service requests are enabled, you can perform the following service request actions:

- View all open service requests: View all open service requests that were filed with the Oracle Enterprise Manager Ops Center software.
- View your service requests: View open service requests that you filed with the Oracle Enterprise Manager Ops Center software.

Actions for Incidents Knowledge Base

The Incidents Knowledge Base is located in the Operational Plans section of the Plan Management section. You can perform the following actions for the Incidents Knowledge Base:

- **View and sort the annotations by type**: The Incident Knowledge Base appears as a table with the name, description, subtype, target type, version and date the annotation was last updated.
- **Add an annotation for an Incident type**: Add a comment, suggested action, or automated action for a type of incident and monitored attribute. You define the severity levels that the annotation is applicable.
- **Edit an annotation for an incident type**: Edit an annotation for a type of incident.
- **Delete an annotation for an incident type**: Delete an annotation for a type of incident.

Location of Incident and Service Request Information in the User Interface

Incident and service request information is available in the Message Center in the Navigation pane. Use the Message Center to view and manage view notifications, service requests and warranty information for an asset. You can also view incident details by expanding the asset and clicking the Incident tab.

Table 9–2 Location of Incidents and Service Request Information in the UI

Object	Location
Unassigned Incidents	Expand the Message Center in the Navigation pane, then click Unassigned Incidents. Or, click the asset in the Assets section, then click the Incidents tab.
My Incidents	Expand the Message Center in the Navigation pane, then click My Incidents. Or, click the asset in the Assets section, then click the Incidents tab.
Incidents Assigned to Others	Expand the Message Center in the Navigation pane, then click Incidents Assigned to Others. Or, click the asset in the Assets section, then click the Incidents tab.
Open Service Requests	Expand the Message Center in the Navigation pane, then click Open Service Requests.
My Service Requests	Expand the Message Center in the Navigation pane., then click My Service Requests.
Notifications	Expand the Message Center in the Navigation pane, then click Notifications.
Relayed Incidents	Expand the Message Center in the Navigation pane, then click Relayed Incidents.
Relayed Service Requests	Expand the Message Center in the Navigation pane, then click Relayed Incident Requests.
Incident Knowledge Base	Expand Plan Management in the Navigation pane, then click Incident Knowledge Base.

Using the Message Center

Each time a monitored attribute does not meet its monitoring rule, a new alert is generated. The first alert raises an incident, which appears in the Message Center. Subsequent alerts for the same rule and asset are correlated with the same open incident.

When a value for an attribute exceeds its monitoring rule and then later meets the rule, the alert is cleared automatically. The change in the attribute value does not automatically clear the incident. When an attribute's value is moving in and out of its monitoring rule's parameters, alerts are generated and cleared continuously. A new incident is only generated when the original incident is closed. When an incident is not yet closed, the new alerts are aggregated into the existing incident. You must close an incident to clear it from the user interface.

Incidents, notifications, and service requests that are generated by managed assets appear in the Message Center. Your assigned user role determines the actions that you can perform in the Message Center.

Incidents, notifications and service requests appear in the following categories in the Message Center:

- Unassigned Incidents: Newly created incidents and those that have not been assigned an owner.
- My Incidents: All incidents that are assigned to you. You can perform additional actions to these incidents to manage their status, such as: Take Action, mark as being repaired, acknowledge, and open a service request for the incident.
- Incidents Assigned to Others: Incidents that are currently assigned to other users. You can view these incidents, but you cannot perform specific actions on them.
- Open Service Requests: All Service Requests that have been filed on assets that are submitted through Oracle Enterprise Manager Ops Center.
- My Service Requests: All Service Requests that you submitted on assets that are managed by Oracle Enterprise Manager Ops Center.
- Service Requests Opened by Others: All Service Requests that have been submitted on assets that are managed by Oracle Enterprise Manager Ops Center.
- Notifications: Information that is automatically generated by services running in the backend while monitoring the assets. Typically, notifications are less important than incidents.
- Relayed Incidents: All incidents reported from any discovered Oracle Engineered System. You must log in to the Oracle Enterprise Manager Ops Center instance that manages each Oracle Engineered system to fix any incidents related to its assets.
- Relayed Service Requests: All open Service Requests for any discovered Oracle Engineered System.

Closed incidents and service requests do not appear in the Message Center.

Figure 9–1 Message Center

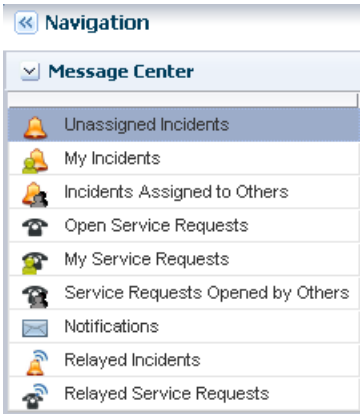

A bar chart on the page visually displays the number of new incidents by severity and the total number of unassigned incidents. The page also contains a table that categorizes the incidents by severity and asset type. Select a row in the table to display all unassigned incidents for the selected severity category. A third table, at the bottom of the dashboard provides details and action icons.

The All Unassigned Incidents table shows all open unassigned incidents. Each incident receives an ID to help you to track the issue. The table includes a description, source, Creation Date, and URL field for each incident. Hover over the URL icon for an incident to display a pop-up window. The default sort is by the creation date; however, you can sort by any column.

Figure 9–2 Unassigned Incidents Dashboard

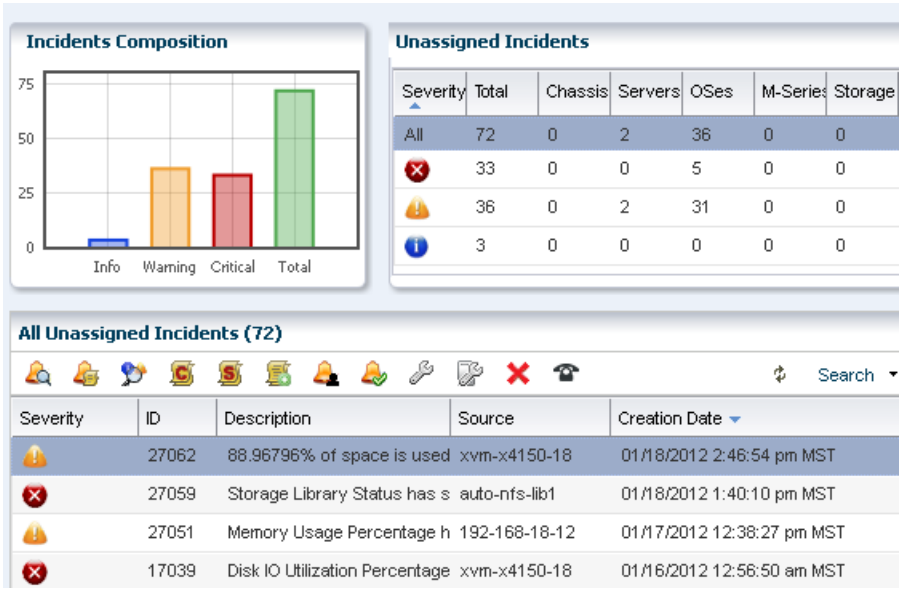

The following information is available in the Unassigned Incidents table:

- Severity: The severity icon shows the severity level, either informational (info), warning, or critical.
- ID: A generated ID assigned to the incident to help track the issue.
- Description: A brief description of the incident.

- Source: The asset that is generating the incident.
- Creation Date: The date and time that the incident generated.
- URL: Contains incident details. Hover over the URL icon to display the duration of the incident, when the incident was assigned, to whom the incident was assigned, any suggested actions, the source of the incident, and a larger description field.

A row of icons at the top of the Unassigned Incidents table provides actions that are available to you, based on your user role. The following action icons are available for incidents, the actions are in the order that the icons appear in Figure 9–3:

- View Alerts
- View Annotations
- View Possible Impacts and Causes
- View Comments
- View Suggested Actions
- Add Annotation to Incidents
- Assign Incidents
- Acknowledge Incidents
- Take Actions on Incident
- Mark Incidents as Repaired
- Close Incidents

Figure 9–3 Incident Actions

Incidents Dashboards in the Message Center

Incidents assigned to you appear in the My Incidents dashboard in the Message Center. Anybody with the appropriate permission level can see the incidents in your queue, can reassign incidents to another user, and can assign you new incidents. Unassigned Incidents and Incidents assigned to others have their own categories in the Message Center. Each Incident dashboard is laid out the same.

A bar chart on the page visually displays the number of new incidents by severity. The page also contains a table that categorizes the incidents by severity and asset type.

Figure 9–4 My Incidents

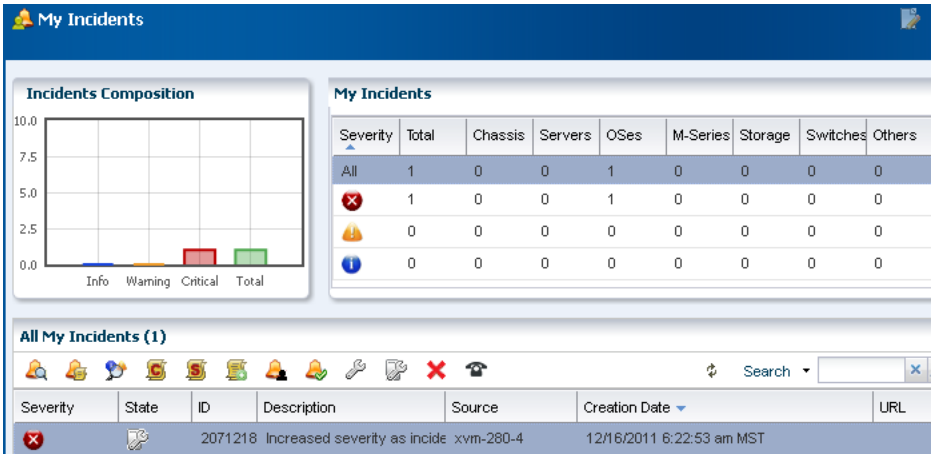

Select a row in the table to display all incidents for the selected severity category and to drill down for more details. The software assigns each incident an ID for tracking purposes. Hovering your mouse over any field in the incident row displays details.

Figure 9–5 Incident Details

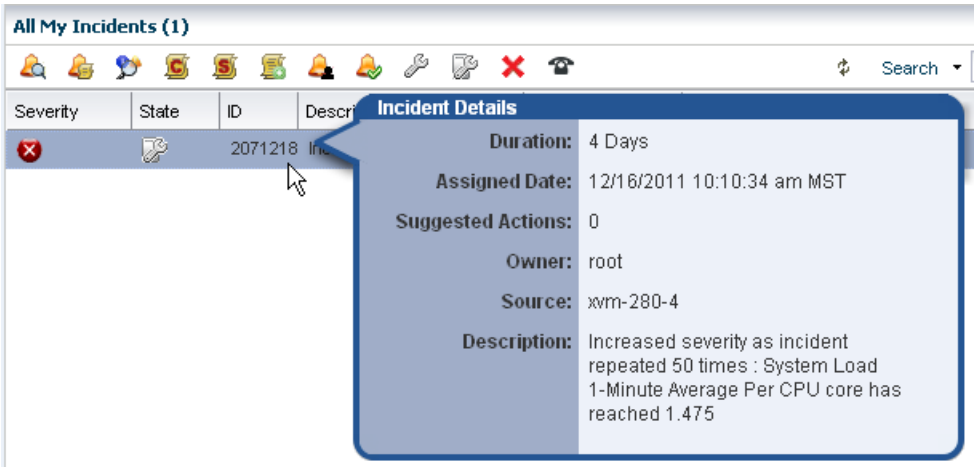

The icons in the center pane enable you to perform actions on a specific incident or view the alerts that make up an incident in the center pane. Each icon has text to define the action.

Figure 9–6 Incident Icons

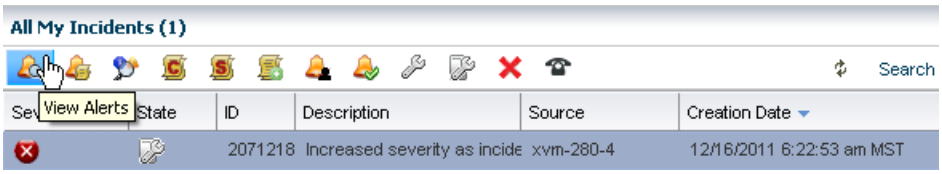

The following actions, as described in Actions for Incidents, are available:

- View Alerts
- View Annotations
- View Possible Impacts and Causes

- View Comments
- View Suggested Actions
- Add Annotation to Incidents
- Assign Incidents
- Acknowledge Incidents
- Take Actions on Incident
- Mark Incidents as Repaired
- Close Incidents
- Open a Service Request

The Notifications appear in the Message Center. You can also configure the software to send you e-mail or pager notification of incidents for critical assets or severity levels.

About Incident Severity Badges

Incident severity badges are an option that you can enable to display an incident severity badge next to an asset show the incident status. Badges also appear in the dashboard membership graph to show the status of an asset. See `Badges` in the *Oracle Enterprise Manager Ops Center Concepts Guide* for a description of each badge.

To Display Incident Severity Badges

Perform the following to display the badges in the Navigation pane:

1. Click **Setup** in the upper right corner of the UI.
2. Click **My Preferences**, then click **User Interface Preferences**.
3. Click the check box to make the badges visible in the user interface.

When an asset has an incident, the severity badge appears next to the asset icon in the Asset hierarchy. When it is the highest severity incident in the membership of a group, the severity badge also appears next to the parent assets. In Figure 9–7, the operating system for xvmbrm-t5140-2 has a critical incident. The critical incident badge also appears on the system and service processor. Any group that this OS is a member of, such as All Assets and Operating Systems, also display the badge.

Figure 9–7 Critical Incident Badge

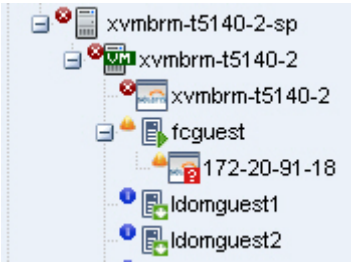

When the incident is the only critical incident, the badge is removed when the incident is acknowledged, marked repaired, or closed. When open incidents are still present, the next highest severity badge displays. For example, when both a Critical and a Warning incident is detected and the Critical incident is acknowledged, the Critical badge is replaced with the Warning badge because that is now the highest level unacknowledged incident.

Figure 9–8 Warning Incident Badge

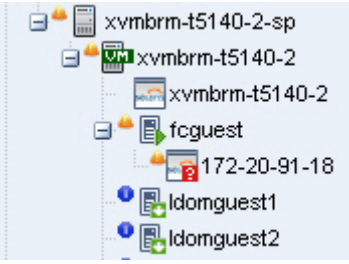

After the incident is closed, its severity badge is not displayed in the asset hierarchy.

See Chapter 4, "Monitoring Rules and Policies" for more information about creating and maintaining monitoring rules and profiles.

An incident is one or more alerts on the same monitored attribute and asset. A ticker at the top of the UI contains the number of unassigned critical incidents, unassigned Warning incidents, number of relayed incidents, the number of critical incidents and the number of warning that are assigned to you. You can click an icon to display the incidents dashboard for that category.

Figure 9–9 Incidents Ticker

Using Annotations

Annotations are comments, suggestions, or automated operations with associated scripts that you can use to document an incident. They are defined by an asset type, asset resource type, an attribute, or an incident type.

Annotations are as powerful as you want to make them. You can View Annotations, View Possible Impacts and Causes, View Comments, View Suggested Actions, and Add Annotations that are associated with an incident in the Message Center or from the Incidents tab in the Asset view, as shown in Figure 9–10.

Figure 9–10 Annotation Icons

You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, and when you close an incident.

Figure 9–11 Add Comment Type Annotation

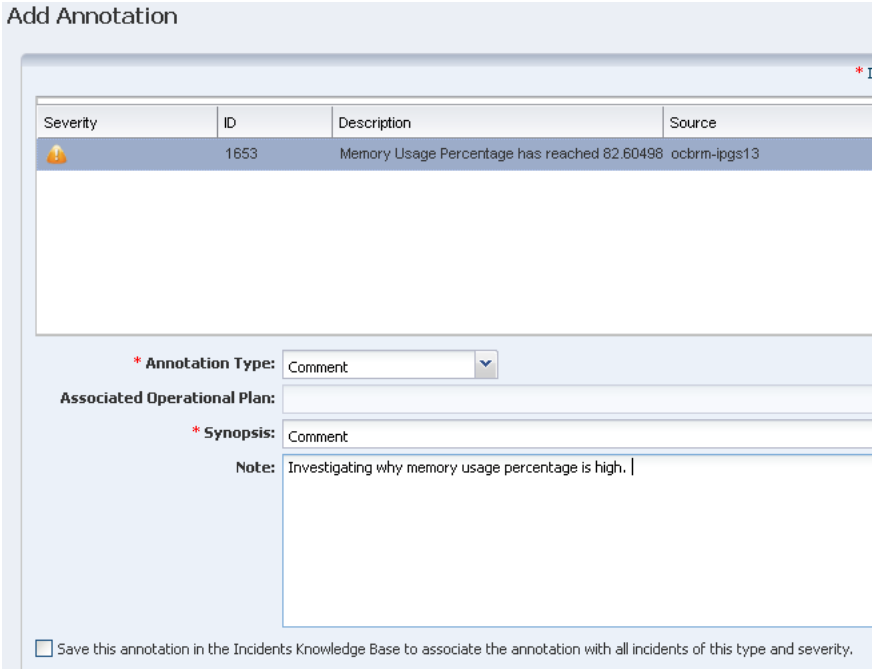

When you create an annotation, other users who have set up notification profiles to use e-mail or a pager are informed about the new annotation and see its content.

You can associate annotations with an incident instance or an asset type. The Incidents Knowledge Base contains your annotations, by asset type, and stores the information on the Enterprise Controller. You can view this annotation by browsing the Incidents Knowledge Base.

For more robust incident management, associate an annotation with an incident and store the annotation in the Incidents Knowledge Base to provide an automated solution or a suggested action when a similar incident is detected. To provide an automated solution, create an operational plan that contains a shell script. When a specific incident occurs, the script is executed automatically. To provide a suggested fix or course of action for a specific incident, you can create a text only annotation that provides a suggested course of action, or you can include a shell script.

When a rule is triggered, and an alert or incident is identified, the software checks the Incidents Knowledge Base and Incident Profiles for the Incident Type and any associated annotations. The software generates an incident of the defined type and severity level and attaches any annotations. When an automated operation annotation is associated with the incident, the script is executed. When a suggested action annotation is associated with the incident, the text and script (when available) appear in the incident details.

Annotations are available from several different views within the UI:

- In the Message Center, when an annotation is associated with an incident
- In the Asset view, when an annotation is associated with an incident
- In the Incident Knowledge Base (in the Plan Management section), when an annotation is associated with an asset type

When an annotation is associated with an incident, you can click the View Annotations icon to see the associated annotation.

You can create annotations by an asset type, asset resource type, an attribute, or an incident type. The annotations are located in the Incidents Knowledge Base (KB). Annotations are automated operations with associated operational profile, suggested fixes or actions, or text only comments. You can associate Automated Action and Suggested Action annotations with an operational profile, which can contain a script.

You can also add an annotation to the Incidents KB when you add an annotation to a specific incident.

A comment is a type of annotation. You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, and when you close an incident. You can also use annotations to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions. To add a comment, see Adding an Annotation.

Annotations in the Incidents Knowledge Base (KB) are associated with types of incidents. You can delete an annotation in the Plan Management section of the UI.

Building an Incidents Knowledge Base

Annotations enable you to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions.

You can use your annotations to build an Incidents Knowledge Base that is based on your company policies, procedures, and incidents.

You can add annotations to the Incidents Knowledge Base in Plan Management or from a specific incident. Either way, annotations enable you to provide solutions or recommended actions for specific incidents.

The following types of Annotations are available:

- **Automated Operation**: You can create an annotation that references an Operational Profile and automatically executes the profile when a specific incident occurs.
- **Suggested Action**: You can provide a suggested fix or course of action for a specific incident. The suggested course of action can be text only, or you can reference an Operational Profile.
- **Comment**: You can write notes about the situation, such as the current status of an incident. This is a text only field.

You can build an Incident Knowledge Base that contains your annotations from specific incidents, or add suggested fixes or automated fixes for a specific type of incident. You can add possible causes and impacts or you can add an annotation. An annotation is either a comment or a suggested action. Annotations enable you to associate an operational plan that contains a utility or script to correct a specific issue. You can choose to associate the annotation with all incidents of the same type and severity and save it in the Incidents Knowledge Base. The next time the incident occurs, you can access the possible causes and impacts and annotations identified in the previous incident and resolve the issue more quickly.

When an incident is detected, Oracle Enterprise Manager Ops Center checks the Incidents Knowledge Base and Incident Profiles for the asset type and corresponding incident type. Any associated annotations are added to the Incident and Operational Profiles referenced in Automated Operation annotations are executed against the asset on which this incident was open. You can create, update, and delete annotations in the Incidents Knowledge Base.

Using Annotations in the Incidents Knowledge Base

For more robust incident management, you can associate an annotation with a rule and store the annotation in the Incidents Knowledge Base to provide an automated solution or a suggested action when a specific incident is detected. To provide an automated solution, create an operational plan that contains a shell script. When a specific incident occurs, the script is executed automatically. To provide a suggested fix or course of action for a specific incident, you can create a text only annotation that provides a suggested course of action, or you can include a shell script.

When a rule is triggered, and an alert or incident is identified, the software checks the Incidents Knowledge Base and Incident Profiles for the Incident Type and any associated annotations. The software generates an incident of the defined type and severity level and attaches any annotations. When an automated operation annotation is associated with the incident, the script is executed. When a suggested action annotation is associated with the incident, the text and script (when available) appear in the incident details.

Annotations are automated operations with associated scripts, suggested fixes or actions, or text-only comments. You can associate annotations with an incident instance or an asset type. The Incidents Knowledge Base contains your annotations, by asset type, and stores the information on the Enterprise Controller. This annotation can be viewed by browsing the Incidents Knowledge Base.

1. Open the incident from the Message Center or Assets view.
   - Message Center View: Click **Message Center**, then click an Incident category: **Unassigned Incidents**, **My Incidents**, or **Incidents Assigned to Others**.
   - Assets View: Click the asset in the Assets section of the Navigation pane, then click the **Incidents** tab in the center pane.
2. Click the incident in the center pane.
3. Click the **View Annotations** icon.

Example 9–1 Example of Using Annotations

```
The CPU usage on a Sun Fire x4150 host is exceeded and an incident is generated.
You assign the incident to Lee. Lee is concerned because these systems are often
used to host Oracle Solaris Zones. Lee adds the following comment to the incident:
"This asset is not powerful enough and cannot cope with the load". Lee also wants
to associate an annotation with the Global Zone asset type to recommend checking
for processes that are consuming excessive CPU usage on the Global Zone. Lee adds
the following annotation to the asset type: "Run the 'prstat 1 1' command to check
which processes are taking CPU." The annotation is saved in the Incidents
Knowledge Base and appears the next time CPU usage is exceeded on a global zone
asset type.
```

Managing Incidents

Incident management in Oracle Enterprise Manager Ops Center consists of several components that are designed to work together to simplify managing incidents for a large number of assets. The components include monitoring rules, suggested actions, and methods for automating incident identification and resolution.

Monitoring includes a standard set of monitoring rules, consisting of an asset's attribute and the threshold value for that attribute. When Oracle Enterprise Manager Ops Center performs monitoring, it generates alerts, which connect to both the incident management and notification features.

When an asset is not operating within the parameters defined in the monitoring rules and policies, Oracle Enterprise Manager Ops Center generates an incident and displays the information in the Unassigned Incidents section. Incidents appear as Informational (Info), Warning, or Critical severity. All incidents appear in the Message Center. When an incident is first detected, it appears in the Unassigned category. When you assign an incident to yourself, it moves from the Unassigned Incidents to My Incidents. When an incident is assigned to someone else, it appears in Assigned Incidents.

For example, the CPU usage on a Sun Fire x4150 host is exceeded and an incident is generated. You assign the incident to Bob. Bob is concerned because these systems are often used to host Oracle Solaris Zones.

Bob reviews the incident and adds the following comment to the incident: This asset is not powerful enough and cannot cope with the load. Bob also wants to associate an annotation with the Global Zone asset type. He wants to add a recommended action annotation to the asset type to check for processes that are consuming excessive CPU usage on the Global Zone. He adds the following annotation to the asset type: `Run the prstat 1 1 command to check which processes are taking CPU.` The annotation is saved in the Incidents Knowledge Base and displays the next time CPU usage is exceeded on a global zone asset type.

Methods of Incident Management

Oracle Enterprise Manager Ops Center uses a help desk approach to managing incidents. The following are the key tools available for taking action on an incident:

- Message Center: View the status of incidents and assign incidents.
- Annotations: Add notes and change status. Use annotation options to provide recommended actions or fixes, or add custom scripts to provide an automated response to an incident.
- Operational Plans: Deploy a shell script against a specific asset, or asset sub-type to automate incident resolution.
- Incidents Knowledge Base: Collect comments and suggested actions for known issues for future use.

When you want to receive e-mail or pager notification each time an incident is reported in the Message Center, you can create notification rules to send a message advising you of a new critical or warning incident.

The Message Center contains a detailed list of unassigned incidents, incidents assigned to you, and incidents assigned to other users.

You can manage incidents from either the Message Center or from the Asset view. You can view and add comments and annotations, take action on an incident, and close incidents.

- The Message Center provides a list of all incidents. Select an incident to see its details and activity.
- From the Asset tree in the Navigation pane, select the asset and then click the Incidents tab to see a list of incidents for that asset.

When you have the Manage or Admin role for the asset, you can take action on the incident. The person assigned to the incident must also have the Manage or Admin role. When the icon is not active, you do not have the appropriate role.

Viewing Unresolved Incidents

You can view unresolved incidents for a specific asset or by incident:

- To view unresolved incidents for a specific asset, click the asset in the Navigation pane, then click the Incidents tab in the center pane.
- To view unresolved incidents from the Message Center, click one of the following:
  - Unassigned Incidents
  - My Incidents
  - Incidents Assigned to Others

    The number of unresolved incidents for an asset appears in a bar chart and in a summary by severity. All Unresolved incidents appear in a table.

View high-level details by hovering your mouse over the incident or clicking the incident in the Unresolved Incidents table. You can drill down to view the alerts that make up the incident by clicking the incident, then clicking the Alerts icon in the center pane.

Viewing Incident Details

A incident consists of one or more alerts. You can view incident details, including the individual alerts that are part of the incident. The following incident details are available:

- How long the incident has been open, or the duration
- When the incident was assigned
- The number of suggested actions for the incident
- Who is assigned to the incident
- Which resource is affected, or the cause
- A description of the incident

Perform the following steps to view incident details:

1. Select **Assets** in the Navigation pane.
2. Select an asset that has an incident badge next to the icon. The Dashboard page displays with the status of the asset.
3. Click the **Incidents** tab.
4. Hover over the incident to display the incident details.
5. To display the alerts that are associated with the incident, click the **Alerts** sub-tab or click the **Alerts** icon in the center pane. The alerts that make up the incident are displayed, including the current and highest alert status, and the alert history.

Assigning an Incident

You can assign an incident to a user who has Manage or Admin role for the asset.

Assigning an incident might affect the asset's Incident severity badge. When an incident was previously acknowledged or marked as being repaired, its severity was not propagated up to antecedent assets in the navigation pane. After assigning an incident (to a user or to no one), the severity is propagated up again to antecedent assets in the navigation pane.

1. To display an incident from the Message Center, click Message Center, then click **Unassigned Incidents** in the navigation pane.

   To display an incident from the asset view, click the asset in the Navigation pane, then click the **Incidents** tab.

2. Select one or more incidents in the center pane, then click the **Assign Incidents** icon.

3. Select a user name from the **Assign To** list, which is the list of users who have either the Manage or Admin role for the asset. To relocate an assigned incident back to the Unassigned Incidents queue, select **No One** from the list.

4. (Optional) Add a note in the text field.

5. Click **Assign Incidents**.

Acknowledging Incidents

Acknowledging an incident indicates that you are investigating the issue. You can acknowledge an incident when you have the Admin or Manage role for the asset on which the incident is identified.

When severity badges are enabled and an incident occurs on an asset, a severity badge appears next to the asset in the Asset hierarchy. When it is the highest severity incident in the membership, it also appears next to the parent assets. In Figure 9–12, the operating system for xvmbrm-t5140-2 has a critical incident. The critical incident badge also appears on the system and service processor. The badge appears on any group that this operating system is a member of, such as All Assets and Operating Systems.

Figure 9–12 Critical Incident Badge on Asset and Parent Assets

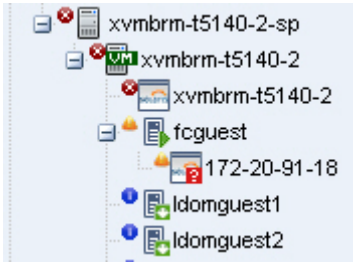

When the incident was previously in an Unassigned state or was assigned to someone else, the severity was taken into account in the computation of the highest severity to propagate up to antecedent assets in the navigation pane. When you acknowledge an incident, it is moved into your queue in the Message Center and the severity is no longer propagated up to antecedent assets in the navigation pane.

When you acknowledge the Critical incident, the badge is replaced with the Warning badge because that is now the highest level unacknowledged incident.

Figure 9–13 Effect of Acknowledging a Critical Incident

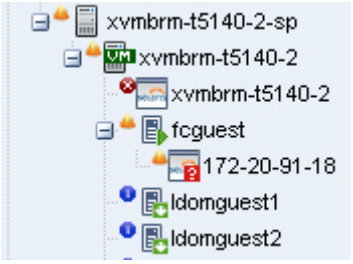

1. Open the incident from the Message Center or Assets view.
   - Message Center View: Click Message Center, then click an Incident category: Unassigned Incidents, My Incidents, or Incidents Assigned to Others
   - Assets View: Click the asset in the Assets section of the Navigation pane, then click the Incidents tab in the center pane.
2. Select one or more incidents, then click the **Acknowledge Incidents** icon in the center pane.

Adding an Annotation

Annotations are defined by the asset type. Annotations are comments, a suggested action, or a reference to an operational profile. Any user can add an annotation to a Incident. To add an entry to the Incidents Knowledge Base requires Oracle Enterprise Manager Ops Center Admin permissions.

1. Open the incident from the Message Center or Assets view.
   - Message Center View: Click **Message Center**, then click an Incident category: **Unassigned Incidents**, **My Incidents**, or **Incidents Assigned to Others**
   - Assets View: Click the asset in the **Assets** section of the Navigation pane, then click the **Incidents** tab in the center pane.
2. Select the incident, then click the **Add Annotations** icon in the center pane.
3. Select one of the following types from the Annotation Type from the drop-down list:
   - **Comment**: Text only option that is designed to be used to add a note or editorial comment.
   - **Suggested Action**: Text required and a script is optional.
4. Select an operational plan from the drop-down list of operational profiles defined for the type of asset on which this incident is open.
5. The Synopsis field is completed based on the annotation type. Edit the synopsis, as needed. The UI does not have a character limit, but the API allows for 80 characters.

   Note: When you enter more than 80 characters, the synopsis is truncated to the first 80 characters when viewed in the annotation.

6. Type a description or instructions in the **Note** field. There is no character limit.

7. To add the annotation to the Incidents Knowledge Base and include the annotation for every incident of this type and severity, click the check box.

   Note: You must have the Oracle Enterprise Manager Ops Center Admin role to complete this operation.

8. Click **Save and Execute** or click **Save**.

Displaying Annotations

You can display annotations for an asset type in the Incidents Knowledge Base.

1. Click **Plan Management**.
2. Expand Incidents Knowledge Base in the Navigation pane, then select the asset type. The annotations associated with the asset type appear in the center pane.

Viewing Comments

A comment is a type of annotation. You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, or when you close an incident. You can also use annotations to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions. To add a comment, see Adding an Annotation.

1. Expand the **Message Center**, then click one of the following:
   - **Unassigned Incidents**
   - **My Incidents**
   - **Incidents Assigned to Others**
2. Click the incident in the center pane, then click the **View Comments** icon.

Taking Action on a Incident

When you have the Manage or Administration role for an asset that has an open incident, you can correct some incidents by using a script or command. In these situations, you might want to associate an automated action with some known issues. For other incidents, you might want to review the issue before deciding on the appropriate action.

The incident management functionality provides varying levels of control. You can build the Incidents Knowledge Base with annotations that contain a combination of automated actions, suggested actions, and comments. You can execute one or more of the suggested actions associated with the Incident or use an operational plan to correct the incident. When an action is not available in a suggested action or Operational Profile, you can execute a command or a custom script that is stored on the managed asset or on the Enterprise Controller.

See Adding an Annotation for information on adding a suggested action to a Incident.

See Building an Incidents Knowledge Base to add a suggested action or automated action in the Incidents Knowledge Base. See Creating an Operational Profile and Plan for information about creating an operational profile.

1. Open the incident from the **Message Center** or **Assets** view.

- **Message Center View**: Click Message Center, then click an Incident category: Unassigned Incidents, My Incidents, or Incidents Assigned to Others
- **Assets View**: Click the asset in the Assets section of the Navigation pane, then click the Incidents tab in the center pane.

2. Select the incident.
3. Click the **Take Actions** on a Incident icon in the center pane.
4. Select the action to perform:
   - When the Incidents Knowledge Base has provided a suggested action for the incident, select **Execute the Selected Suggested Action** option and then select the action from the table.
   - When an operational plan has a suggested action, select the **Execute an Operational Plan** option, then select the plan from the drop-down list.
   - To run a script or command that is not part of a suggested action or operational plan, select the **Execute a Command or Script File** option.
     - To execute a command, enter the command in the field.
     - To browse for a script, click **Browse** and then select the script from the **File Chooser** popup.
5. Select where to run the script, on the managed asset where the incident is open, or on the Enterprise Controller.
6. Define the time out period for the action, in minutes, hours, or days.
7. (Optional) Add a note describing the action taken.
8. Click **Execute Selected Action**.

Marking an Incident Repaired

The software cannot determine when an incident is repaired. However, you can open a known incident and manually add a note with the repair details and mark the incident as repaired. You must have the Manage or Admin role for the asset to perform this task.

When the incident was previously in an Unassigned state or Assigned to someone, the severity was taken into account in the computation of the highest severity to propagate up to antecedent assets in the navigation pane. After marking this incident as repaired, its severity badge does not appear in the assets list in the navigation pane.

1. From the Message Center, click **Message Center**, then click one of the following:
   - Unassigned Incidents
   - My Incidents
   - Incidents Assigned to Others

   or

   From the asset view, click the asset in the Assets section of the Navigation pane, then click the **Incidents** tab.
2. Select one or more incidents, then click the **Mark Incidents as Repaired** icon in the center pane.
3. (Optional) Select the incident, then add a Note.
4. Click **Tag Incidents as Being Repaired**.

Closing an Incident

The incident stays open until you close it, even if the alerting condition is cleared. To remove an incident from the Message Center and the asset view, you must close the incident or take no action on it for seven (7) days. When any action is taken on an incident, such as adding an annotation, the counter is reset.

> **Note:** Incidents with no activity for seven (7) days are closed automatically by Oracle Enterprise Manager Ops Center, and do not appear in the UI. You can edit this value in the public API.

When an incident is closed, its status changes to Closed, the incident is deleted from the list of active incidents, and the incident is no longer displayed in the UI. You can retrieve information about a closed incident for 60 days by using the public API. After 60 days, closed incidents are permanently deleted. To edit the time limit, you must edit the value in the public API. You can disable the time limit by setting the value for the number of days to 0.

> **Note:** When the monitoring condition is still true after the incident is closed, a new alert is raised and a new incident is created.

To Close an Incident

You can close an incident from the asset view or from the following categories in the Message Center:

- Unassigned Incidents
- My Incidents
- Incidents Assigned to Others

Perform the following steps to close an incident from the asset view:

1. Click the asset in the Assets section of the Navigation pane, then click the Incidents tab.
2. Select one or more incidents, then click the **Close Incidents** icon in the center pane.
3. (Optional) Select the incident, then add a Note.
4. (Optional) To temporarily disable the monitoring rule that identified the incident, click the **Action** check box, then define when to enable the monitors.

   This action does not disable the monitoring rule for all assets. The action disables the monitoring rule for only the assets that were related to the incident to avoid raising a similar incident on the same assets.
5. Click **Close Incidents**.

Disabling and Enabling Incidents and Alerts

Incidents and alerts are enabled by default. You can use one of the following methods to disable and enable incidents and alerts:

- Using Maintenance Mode: Disables incidents from generating on individual assets or a group of assets. This feature is designed to temporarily disable incidents while you perform maintenance tasks.

- Disabling and Enabling Alert Monitoring: Disables monitoring policies and prevents incidents and alerts from generating on all assets in your data center.

> **Note:** Neither option disables the collection of data on managed assets.

Using Maintenance Mode

Maintenance mode is designed to temporarily disable assets from generating incidents. This mode is useful when you plan to power off a hardware asset, reconfigure a system manually, or perform maintenance on a system and you do not want incidents to appear in the UI.

> **Note:** Monitoring still occurs and alerts are still generated on the asset while in maintenance mode. View alerts by selecting the **Alerts** tab, which is a subtab of the Incidents tab.

When you place an asset in maintenance mode, the severity badge of unassigned and assigned incidents affecting the asset and its children is not propagated up the asset membership hierarchy in the navigation pane.

When you place a Proxy Controller in maintenance mode, you disable incidents from generating on the Proxy Controller and you disable all jobs that go through the Proxy Controller, including discovering, managing, and migrating assets.

When you place an Oracle VM Server that is a member of a server pool in maintenance mode, all of the virtual machines running on the Oracle VM Server are automatically migrated to other Oracle VM Servers in the server pool, if they are available. When the Oracle VM Server is the master Oracle VM Server in the server pool, this role is moved to another Oracle VM Server in the server pool, if available. While in maintenance mode, you cannot create or place guests and guests cannot be recovered to the server from another control domain. When the Oracle VM Server is not a member of a server pool, or other servers are not available in the server pool, the virtual machines are stopped. To manually bring down the control domain and all of its guest domains, use the action **Disable Automatic Recovery** on the virtual machines to disable the auto-recovery, then put the control domain in maintenance mode. See Automatic Recovery in Chapter 21, "Server Pools" for more information.

To Place Assets in Maintenance Mode

1. Select an asset in the Navigation pane.
2. Click **Place in Maintenance** in the Actions pane.
3. Click **Place** to confirm the action.

To Remove Assets From Maintenance Mode

When the maintenance operations are completed, use the **Remove From Maintenance** action to begin sending incidents for the asset. When the asset is removed from maintenance mode, the severity badge appears in the asset membership hierarchy in the navigation pane.

1. Select the asset in the Navigation pane.
2. Click **Remove From Maintenance** in the Actions pane.
3. Click **Remove** to confirm the action.

Disabling and Enabling Alert Monitoring

You can disable alerts and incidents for all assets in your data center by disabling all of the monitoring policies.

When you disable the policies, the monitors are no longer deployed on the assets. When you enable monitoring that you previously disabled, the monitoring rules defined by the default monitoring policies are reapplied to all of the assets.

Note: Disabling monitoring disables the evaluation of monitoring rule conditions against collected data and prevents the deployment of monitors across your data center, it does not disable the collection of data on managed assets.

Disabling all monitoring for your data center is only available from the command line interface. See *Oracle Enterprise Manager Ops Center Command Line Interface Guide* for more information. See Using Maintenance Mode for information about temporarily disable the software from generating incidents.

Using Oracle Services and Service Requests

Oracle Services provides integrated methods for maintaining and displaying current contracts, warranty information, contract dates, and service requests for managed assets.

Use the Oracle Services feature to view the contract or warranty information and any service requests for a specific asset. You can also view service requests that were the result of an alert or incident in Oracle Enterprise Manager Ops Center, view service request details, and file a service request.

- Contracts and Warranties

  Maintaining a valid inventory of the assets in your data center, including contracts and warranties, can be a time-consuming and labor-intensive process. Use Oracle Enterprise Manager Ops Center to display current contract and warranty information for a specific asset, or view the entitlements associated with your Oracle online account. When a contract or warranty is about to expire, Oracle Enterprise Manager Ops Center generates an alert.

- Service Request

  You can create new service requests, review your requests, and review the requests of other users.

  You can file requests manually, or you can provide credentials and contact information for assets and configure Oracle Enterprise Manager Ops Center to generate service requests. When an incident occurs with an asset, a service request is automatically created using the credentials and contact information that you provided.

Note: You cannot display service requests created outside of Oracle Enterprise Manager Ops Center. To see the status of a service request filed outside of Oracle Enterprise Manager Ops Center, go to the Service Requests Home page on My Oracle Support.

Requirements for Oracle Services

To use these Oracle Services, you must take the following actions:

- Register your assets with My Oracle Support.
- Register your user account as a My Oracle Support user so that you can get access to the My Oracle Support database.
- Run Enterprise Manager Ops Center in Connected Mode.

To access the My Oracle Support database, your user must be registered as a My Oracle Support user. This is the same account that is used to access My Oracle Support at

To determine if you are running in Connected Mode and have access to My Oracle Support, check the icons in the upper right corner of the UI, as shown in Figure 9–14. If an icon does not contain color, you are not connected.

- The World icon indicates the status of the Internet connection.
- The Shield icon indicates the status of the connection to the Oracle Knowledge Base.
- The Phone icon indicates the status of the connection to My Oracle Support Services.

Figure 9–14 Connection Icons

Viewing Contract and Warranty Information

You can display contract information by asset, or you can obtain entitlements associated with all contracts that are associated with a user. Contract and warranty information is available for managed servers that have a serial number associated with a contract in the My Oracle Support database.

The contract and warranty information in Enterprise Manager Ops Center is updated each week so contract changes or new contracts might take up to seven days to appear in the user interface. When a contract or warranty is about to expire, an alert is displayed as an Incident in the Message Center and the contract details appear in orange text in the asset's Summary tab.

Note: Updating contract and warranty information requires running the product software in Connected mode. If you change to Disconnected Mode, the contract information becomes outdated.

To View Contract and Warranty Information For an Asset

1. Select a hardware asset in the Navigation pane, from either the **All Asset** list or from a group.
2. Click the **Summary** tab.

   When Enterprise Manager Ops Center is in Connected Mode and the serial number of the selected asset is associated with a contract, a Support row is added to the summary. The Support field contains the contract ID and an expiration date.

   - If the contract is within 90 days of expiration, the information is displayed in an orange font.

- If the contract has expired, the information is displayed in a red font.

To View All Contracts Associated with a My Oracle Support Account

1. Click the **Enterprise Controller** in the Administration section of the Navigation pane.
2. Click **Edit Authentications** in the Actions pane. The Edit Authentications window is displayed with online user names and associated contracts.

Viewing Service Requests

You can see all current and completed service requests that are filed through Oracle Enterprise Manager Ops Center. The service request contains information about the request, including the request number, severity, a summary of the problem, the date and time last updated, contact information, and status.

To View Service Requests

1. Click **Message Center** in the Navigation pane.
2. Click **Open Service Requests**, **My Service Requests**, or **Service Requests Opened by Others** to display a list of requests.
3. To view details of a particular service request, highlight a row, then click the **View Service Request** icon.

Figure 9–15 View Service Request

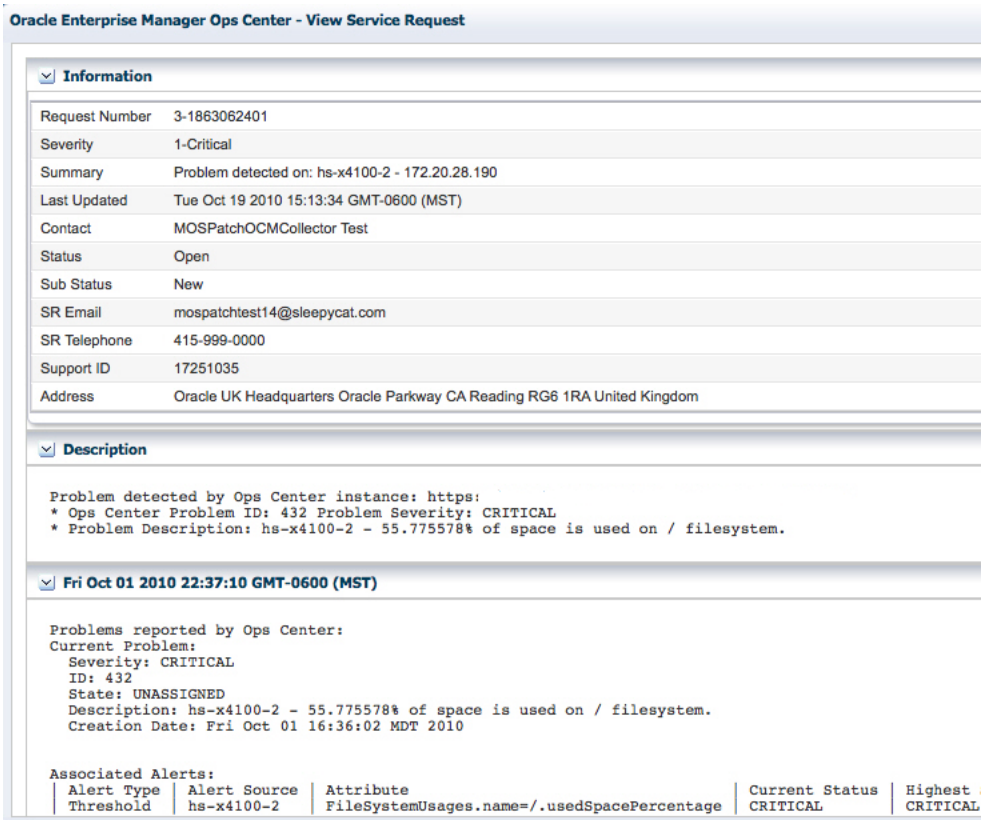

4. Click **Close**.

Filing a Service Request

When your assets are associated with a contract and registered in the Oracle database, you can create a service request from an incident or from an asset. See Requirements for Oracle Services for requirements that must be met before successfully filing a service request ticket. For example, if the asset is not registered in My Oracle Support, the service request job fails. If the Open Service request action is disabled, there is no connection to My Oracle Support.

To File a Service Request From a Incident

1. Click **Message Center** in the Navigation pane.
2. Click **My Incidents** or **Unassigned Incidents**.
3. Select the incident, then click the **Open Service Request** icon in the center pane.

To File a Service Request From an Asset

1. Select the hardware in the Assets section of the Navigation pane.
2. Click **Open Service Request** in the Actions pane.

Auto Service Requests

Instead of manually filing a service request, you can configure Oracle Enterprise Manager Ops Center to automatically create service requests for known issues. When Auto Service Requests (ASRs) are enabled, Oracle Enterprise Manager Ops Center automatically generates service requests based on critical incidents. Contact information for the ASR is taken either from Oracle Enterprise Manager Ops Center or from the Customer Service Identifier (CSI) associated with the asset. Annotations are added to the incident to indicate the status of the ASR creation. Once they are created, ASRs are identical to other service requests and can be viewed and managed using the same processes and tools.

A user with the Ops Center Admin role must enable the ASR feature. See the *Oracle Enterprise Manager Ops Center Administration Guide* for how to configure and use auto service requests.

Related Resources for Incidents

The Oracle Enterprise Manager Ops Center 12*c* Release 2 documentation is available at `http://docs.oracle.com/cd/E40871_01/index.htm`.

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

The *Oracle Enterprise Manager Ops Center Feature Reference Guide* has information about asset management, Oracle Solaris 11 Software Update library, storage, networks, zones, Oracle VM Server, and server pools. See Chapter 4, "Monitoring Rules and Policies" for information about the monitoring rules that trigger incidents.

For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` and the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm`.

See the following How To documents:

- *Oracle Enterprise Manager Ops Center Managing Incidents*
- *Oracle Enterprise Manager Ops Center Tuning Monitoring Rules and Policies*

- *Oracle Enterprise Manager Ops Center Understanding OS Performance and Capacity*
- *Oracle Enterprise Manager Ops Center Using Service Requests*

See the *Oracle Enterprise Manager Ops Center Concepts Guide* for a description of the icons.

The *Oracle Enterprise Manager Ops Center Administration Guide* has information about user roles and permissions.

10 Reports

This chapter describes how to create reports in Oracle Enterprise Manager Ops Center. This chapter includes the following information:

Introduction to Reports

Reports provide information about assets, such as job history, firmware, operating system updates, and incidents. You can use the Reports feature to consolidate changes to hardware, software, and job conditions. You can use reports to export the information or to start jobs on targeted assets.

Types of Reports

Reports are grouped in Oracle Enterprise Manager Ops Center in the following way:

- **OS Reports**: OS reports includes Oracle Solaris Update reports, Oracle Linux Update Reports, and Windows Update reports. They enable you to check for new

patches and security advisories. You can get a general report, or test a system or installed package for available fixes. See Creating an Operating System Report for more information.

- **System Information Reports**: System Information reports are used to obtain the information on assets such as OS, server, chassis, logical domains, global zone, non-global zone, and M-Series server. See Creating System Information Reports for more information.
- **Oracle Engineered System Reports**: Oracle Engineered System reports enables you to view the rack setup for each of the rack within the system including the asset details related to the rack. These reports provide information about your assets, such as job history, firmware, OS updates, and incidents. See Creating Oracle Engineered Systems Reports for more information.
- **Incident Reports**: Incident report summarizes information about all alerts and incidents for a specified category, such as alarm state, alarm owner, asset type, date range, severity levels, and affected asset groups. It also includes an audit trail consisting of state-change annotations, alert annotations, suggested-fix annotations, comment annotations, operation annotations. See Creating Incident Reports for more information.
- **Firmware Reports**: Firmware Reports enables you to maintain consistent firmware versions across your data center. The Firmware Report feature compares the firmware images specified in a firmware profile to the firmware images installed on one or more hardware assets. The report indicates whether the firmware on the asset complies with the profile's specifications. See Creating a Firmware Report for more information.
- **Additional Reports**: Additional Reports enables you to obtain information from Service Pack Compliance Report, Distribution Update Report, and Package Compliance Report. See Creating Additional Operating System Reports for more information.

Figure 10–1 displays the Reports of Oracle Enterprise Manager Ops Center.

Figure 10–1 Reports in Oracle Enterprise Manager Ops Center

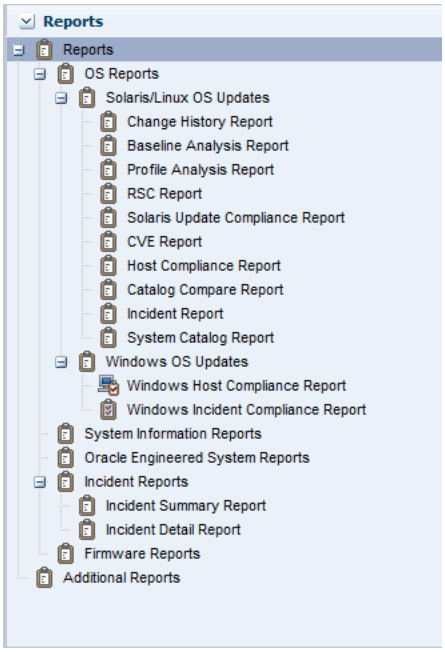

Scheduling Reports

You can schedule to run the report using any of the following parameters:

- Now: Select the current date and time to generate the report.
- At a later date/time: Select a date and time to generate the report.
- On a Recurring Schedule: Select the month and day when you want to generate the report. Select the Start Time, End Time, and Number of Hours between runs. This is to set the number of times the report is generated between the specified start and end time. For example, when you set the start time at 6.00 a.m, end time at 12.00 a.m, and the number of hours between runs as 2, then the report is run at 6.00 a.m, 8.00 a.m, 10.00 a.m, and 12.00 a.m.

Output of Reports

After the report is generated, the Report Results pane lets you to export the report result in CSV and PDF formats, view the report interactively, and delete the report. View Interactive option helps you to view the generated report in detail. Report result displays the Report name, Report type, Run Date, Targets of the Report, OS updates applicable to selected targets. Report parameters show the target name, product name, description. You can save the report as a template and also rerun the report.

Figure 10–2 displays an Interactive view of Report Result.

Figure 10–2 Interactive View of Report Result

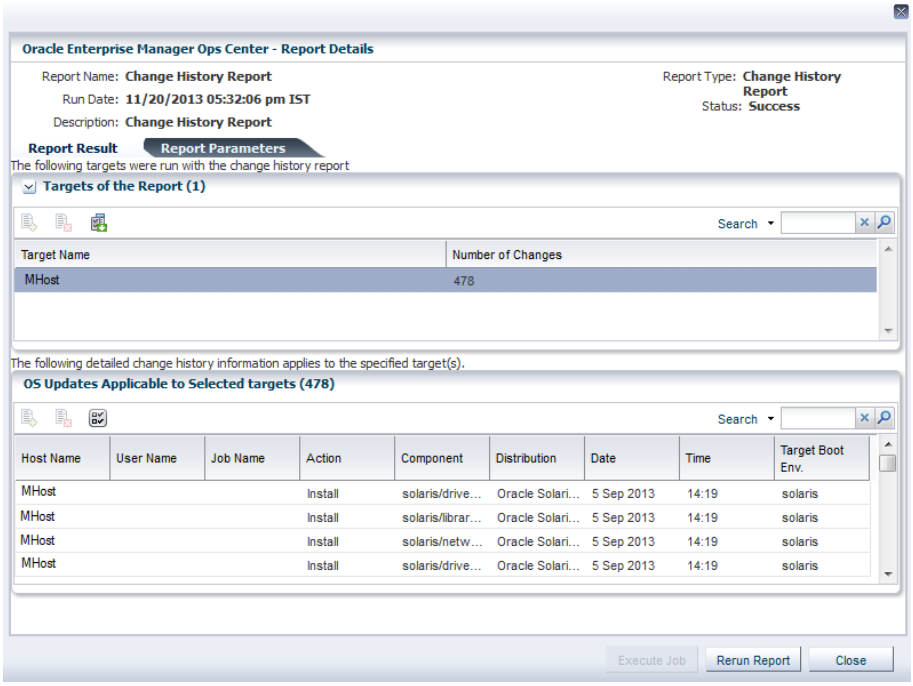

Roles for Reports

The following table lists the tasks and the role required to complete the task. Contact your administrator when you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.Table 10–1 illustrates the roles and permissions for reports.

Table 10–1 Reports Tasks and Roles

Tasks	Roles
View Reports	All
Create Reports	Asset Administrator

Actions for Reports

You can perform the following actions in the reports section.

- Generate a Report from a Report Template
- Update a Report
- Delete a Report
- View a Report Result
- Save a Report Result

Location of Report Information in the User Interface

To see report information, expand Reports in the Navigation pane. This displays OS reports, System Information reports, Incident reports, and Firmware reports. Click OS reports to view the various types of OS reports that run on Oracle Solaris/Linux OS updates and Windows OS updates. Click Incident reports to view the different types of incident reports. Table 10–2 illustrates the location of reports in the user interface.

Table 10–2 Location of Report Information in the UI

Object	Location
OS Reports	Expand the Reports in the Navigation pane. Click OS Reports to view the various reports that run on Oracle Solaris/Linux OS updates and Windows OS updates.
System Information Reports	Expand the Reports in the Navigation pane. Click System Information Reports. This displays the Create System Information Report wizard in the Actions pane.
Oracle Engineered System Reports	Expand the Reports in the Navigation pane. Click Oracle Engineered System Reports. This displays the Create Oracle Engineered System Report wizard in the Actions pane.
Incident Reports	Expand the Reports in the Navigation pane. Click Incident Reports. This displays the types of incident reports.
Firmware Reports	Expand the Reports in the Navigation pane. Click Firmware Reports. This displays the Create Firmware Report wizard in the Actions pane.
Additional Reports	Expand the Reports in the Navigation pane. Click Additional Reports. This displays the Distribution Update Report, Service Pack Compliance Report, and package Compliance Report wizards in the Actions pane.

Creating Templates

A report template is a pre-formatted file that serves as a starting point to create a new report. When you save a file created from a template, you are prompted to save a copy of the file so that you do not overwrite the template. Templates are provided within a software or a program or it is created by the user. Most major software support templates. If you want to create a similar document or report over and over again, it is a good idea to save one of them as a template. You can open the report template and

start creating reports from there. Parameters in the report template are specified when the report is created or run. You can save the criteria as a template, after creating the report criteria.

You can create a report template for any type of a report. In this example, you will create a report template for Change History Report. To create a report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select **Change History Report** and click **Create Change History Report** in the Actions pane. The **Create Change History Report** wizard is displayed.
3. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Date Range: Specify the start date and end date between which the report will cover.
   - Actions: Select the actions to be reported. You can select Install, Uninstall or both.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
4. Click **Next** to schedule the report.
5. Select a desired schedule to run and generate the report.
6. Click **Next** to display the Summary.
7. Review the report parameters and click **Save Template and Close** to save the report template.

   The created report template is displayed in **Report Templates** in the center pane.

Generating a Report from a Report Template

To generate a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report template from the center pane.
3. Click the **Generate Report** icon to run the report. The corresponding report is generated and the results are displayed under Report Results.

Deleting a Report

To delete a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report to delete from the center pane.
3. Click the **Delete Report** icon to delete the report.
4. Click **Ok** to confirm the delete action. The selected report is deleted.

Updating a Report Template

To update a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report template to edit from the center pane.
3. Click the **Edit View** icon to edit the selected report template. The corresponding report wizard is displayed.
4. Edit the report parameters as required in the wizard.
5. Click **Run and Close** to run the report or click **Save Template and Close** to save the report template. When you click **Run and Close**, the report is generated but the edits for the report are not saved.

Viewing a Report Result

Oracle Enterprise Manager Ops Center provides an interactive result viewer to view the results. The generated report results are displayed under the Report Results in the All Reports page. You can select a report result from the Report Results pane to view the report interactively. You can rerun the report, delete the report, view, export, and save the output of report in the format of CSV and PDF.

To view the report result, do the following:

1. Select **Reports** in the Navigation pane. The All Reports page in the center pane displays all the report templates and report results.
2. Select a report result under the Report Result section in the center pane.
3. Click the **View Interactive** icon to view the report result. The **Interactive Result viewer** opens the report result and displays the following information.
   - Report detail: This displays the name, type, run time, and the status of the report.
   - Report Result: This displays the targets on which the report is run and the corresponding operating system updates that are applicable.
   - Report Parameters: This displays the parameters that are used to generate the report.

Saving a Report Result

After you create and generate a report in Oracle Enterprise Manager Ops Center, you can save a report result in CSV or PDF format.

To save the report result, do the following:

1. Select **Reports** in the Navigation pane. The All Reports page in the center pane displays all the report templates and report results.
2. Select a report result under the Report Result section in the center pane.
3. Click the **View CSV or View PDF** icon. You can save or open the report in CSV or PDF formats.

Figure 10–3 illustrates the report output in PDF format.

Figure 10–3 Exporting a Report Result in PDF

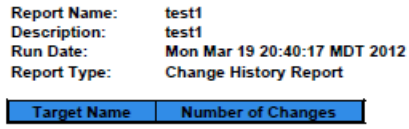

Creating an Operating System Report

Use Operating System update reports to check for new software updates and security advisories. For auditing purposes, create a change history report. Various types of update reports are available for Linux, Oracle Solaris, and Windows operating systems. You can export report results to CSV or PDF format.

Use operating system reports to obtain information about managed Oracle Solaris, Linux, and Windows operating systems.

Table 10–3 illustrates the various reports that are run in Oracle Enterprise Manager Ops Center and displays the type of report that are supported and run on an operating system.

Table 10–3 Compatibility of Reports on Operating System

Report Name	Linux OS	Oracle Solaris 8, 9, and 10	Oracle Solaris 11	Microsoft Windows
Change History Report	Yes	Yes	Yes	No
CVE Compliance Report	Yes	Yes	No	No
RSC Report	Yes	Yes	No	No
System Catalog Report	Yes	Yes	Yes	No
Oracle Solaris Update Compliance Report	No	Yes	No	No
Baseline Analysis Report	No	Yes	No	No
Update Compliance Report	Yes	Yes	No	No
Incident Compliance Report	Yes	Yes	No	Yes
Host Compliance Report	Yes	Yes	No	Yes
Distribution Update Report	Yes	Yes	No	No
Service Pack Compliance Report	Yes	No	No	No
Package Compliance Report	No	Yes	No	No

Update Compliance Reports

The following reports are available for Linux, Oracle Solaris, and Windows operating systems:

- Host Compliance: Provides information on whether your system is compliant with security and bug fixes incidents.

- Incidence Compliance: Provides information about the number of systems to which the selected operating system updates apply.

Oracle Linux and Oracle Solaris OS Update Reports

In addition to the reports created for all types of operating systems, the following reports are available for Oracle Linux and Oracle Solaris operating systems:

- Change History: Provides a history of Operating System update, install, and uninstall jobs completed on managed systems.
- CVE Compliance: Provides information on incidents that are related to specific Common Vulnerability and Exposure Identifiers (CVE IDs) and the systems that have these incidents installed. CVE IDs are unique, common identifiers for publicly known security vulnerabilities.
- Distribution Update: Provides a mapping between selected updates, CVEs, and selected distributions to find out whether the updates are installed.
- Package Compliance: Provides the details of the selected packages on managed system that are compliant or not compliant with the latest recommended version available.
- Recommended Software Configuration (RSC): Provides information about the system compliance for installing a specific application, such as the Oracle 11*g* Database, on an Oracle Solaris, or Linux Operating System.
- Service Pack Compliance (Linux only): Provides information on incidents created by the publication and release of a service pack by a vendor. This helps in determining whether the system has the latest service packs released by the vendor.
- Oracle Solaris Update Compliance (Oracle Solaris Operating System only): Provides information on whether an Oracle Solaris system is compliant with a specific update.
- Baseline Analysis (Oracle Solaris Operating System only): This helps to check the compliance of systems against newly released Oracle Solaris baselines.

Creating a Change History Report

The Change History Report provides a history of operating system update install and uninstall jobs completed on managed Oracle Solaris or Linux systems. The report also displays the deployments made by the specific user, enabling you to track a team of operators.

To Create a Change History Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Change History Report** from the Actions pane.

   The Create Change History Report Wizard is displayed.

Figure 10–4 Defining Report Parameters for Change History Report

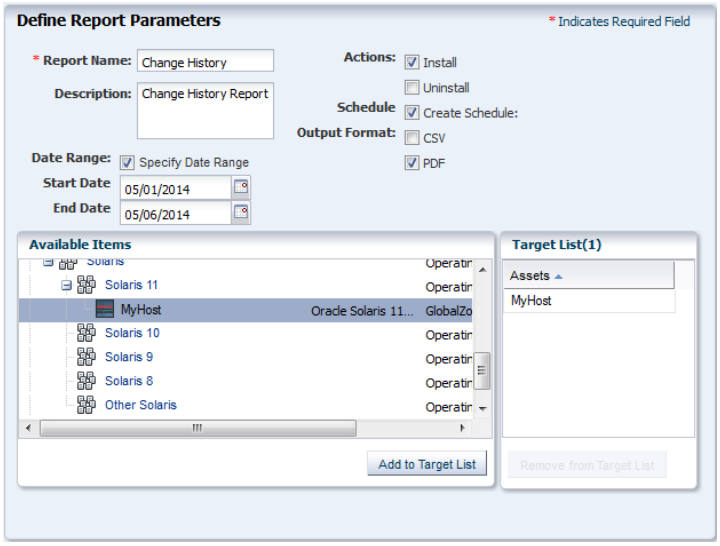

4. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Date Range: Specify the start date and end date between which the report will cover.
   - Actions: Select the actions to be reported. You can select Install, Uninstall or both.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to schedule the report.
6. Select a desired schedule to run and generate the report.
7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:
   - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
   - Run and Close: Runs the report and closes the wizard window.

   The report results are displayed under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result.

Creating a Baseline Analysis Report

An Oracle Solaris baseline is a dated collection of Oracle Solaris updates, update metadata, and tools. Oracle releases Oracle Solaris baselines on a monthly basis. A Baseline Analysis Report checks the compliance of Oracle Solaris systems against

newly released Oracle Solaris baseline. When you install the updates of a baseline on a host, that system is considered to be compliant with that baseline.

Each dated baseline contains these update sets:

- Full: Includes all Oracle Solaris updates
- Recommended: Includes Oracle Solaris recommended updates and security updates
- Security: Includes only Oracle Solaris security updates

All baselines include updates for a specific time frame. However, the Full baseline often contains Oracle Solaris operating system updates that are not included in the Recommended baseline. The Full baseline includes additional updates based on feedback from various customer support groups within Oracle. Recommended baseline does not includes these updates.

To install the Recommended and Security baselines, you must either deploy two jobs or have a job that includes multiple tasks. This might result in multiple reboots, for example, if both tasks (baselines) include updates that have Single User mode requirements.

Oracle Enterprise Manager Ops Center's Knowledge Base (KB) is updated with the information about the baselines. This is done a few days after the official release of baselines by Oracle.

Note: The Oracle Solaris 8 Operating System was placed into End of Service Live (EOSL) on March 31, 2009. Oracle Solaris 8 Operating System baselines are available through March 2009. The KB might contain artificial baselines after that date. Do not use baselines dated after March 2009.

Oracle Solaris baselines enables you to easily identify the update level of your hosts. For example, install some test hosts with a particular baseline. Test these hosts for a period to see whether the updates in this baseline are stable enough to be used on production hosts. When the testing reveals that this baseline is stable, install the same baseline on production hosts.

Oracle Solaris baselines are available as a component in the recommended component list. This contains a list of dated baselines.

The Baseline Analysis report helps to verify the compliance of your system against the newly-released baselines (as and when they are available in Knowledge Base).

The Baseline Analysis Report (BAR) enables you to determine whether the managed system is compliant with recently released Oracle Solaris baselines. Baselines pertain only to Oracle Solaris systems. This section describes Oracle Solaris baselines, white list, black list, and how to run a Baseline Analysis report in connected and disconnected mode of the Enterprise Controller.

The Baseline Analysis Report (BAR) describes how to generate a BAR. The report gives the compliance status of the managed system with the selected Oracle Solaris baseline that was released.

You can generate two types of BARs:

- Agent-based BAR
- Database-based BAR

In an agent-based BAR, a simulated job is run against the managed hosts. This type of report takes time to complete because it checks for dependent components and missing dependencies, and then downloads the updates that must be installed. When you run a compliance job from this report result, the job is completed quickly because the updates are downloaded. However, to improve the report performance of a BAR, skip the downloads in a simulated job by deselecting this option.

In a database-based BAR, the report is run against the database of the management server, the selected baselines are broken down into individual update IDs, and then formed into an incidents list. The report is generated based on the information that are available on the database. Based on the report result, run a compliance job.

White List

A white list is the list of updates that is required to install in addition to the updates in the baseline. To establish a white list, create a profile using the Required setting. You can also specify a white list when generating a Baseline Analysis Report. Select the white list either from the created profile or enter the update IDs separated by new lines.

For example, baseline B includes updates X, Y, and Z, and the white list has updates U, V, and W. When the Baseline Analysis report is created, the host is marked compliant only when all six updates X, Y, Z, U, V, and W are present.

Black List

A black list is a list of updates that you do not want to install. Create a black list by creating a policy with the specified action for the updates. Select a black list option while creating a Baseline Analysis Report. Select the black list either from the created policy or enter the update IDs separated by new lines.

If a particular update in the profile is set with the policy component setting as Never for an install action, then the update is not installed. If the update is installed, it will not be uninstalled or removed.

For example, if baseline A has updates X, Y, and Z, and the black list specifies only Y and Z, the system is compliant if X is installed. If the updates Y and Z are installed, they will not get uninstalled if you run a compliance job from the report results. If Y and Z are not installed, they are not listed in the non compliant result and are not added in the compliance job.

To Create a Baseline Analysis Report

This report provides information about the hosts that are compliant with a baseline operating system.

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Baseline Report** from the Actions pane. The Create Baseline Analysis Report Wizard is displayed.
4. Define the report parameters:
   - Report Name: Name of the report.
   - Description: The description of the report.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.

- Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.

5. Click **Next** to select the Oracle Solaris baselines.
6. In Select Baseline(s), select the following options:

Figure 10–5 Selecting Baselines for Baseline Analysis Report

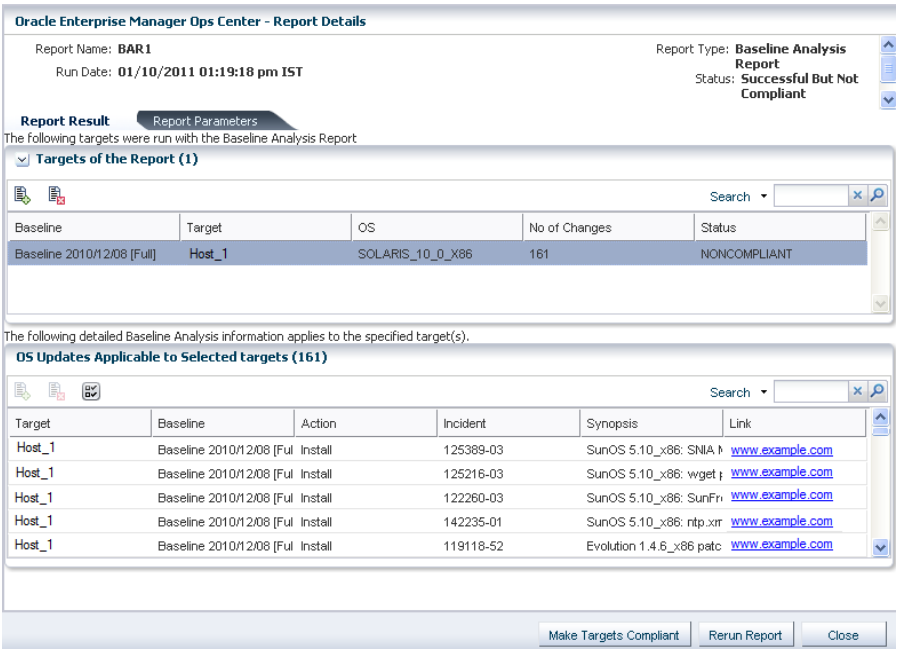

- Run Against Database or Run Report Against Agent.

  When you select **Run Report Against Agent,** check the **Download check box** to download the updates that are installed on the target.

- Select the distribution type and select the baselines from the list. You can select targets of multiple distribution. For each distribution, select the corresponding baselines. A warning message is displayed when the baselines are not selected for a distribution.

  Note: If you have multiple distributions, then you must select baselines for at least one distribution to continue further in the wizard. If you have not selected baselines for a distribution, then the targets of that distribution are not in the report result.

- Click **Add** or **Add All** to select all the baselines.

7. Click **Next** to modify the update lists that are applied to the report.
8. Select any of the following White List options:
   - None: No white list.
   - Manual Input: Enter a list of updates.
   - Specify with Profile: Select a profile to import as a white list.
9. Select any of the following Black List options:

- None: No black list.
- Manual Input: Enter a list of updates.
- Specify with Policy: Select a policy to import as a black list.

10. Click **Next** to schedule the report.
11. Select a desired schedule to run and generate the report.
12. Click **Next** to display the Summary.
13. Review the report parameters and select one of the options as required:
    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate reports later.
    - Run and Close: Runs the report and closes the wizard window.

    The report result displays under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating a Profile Analysis Report

A Profile Analysis Report provides information about Oracle Solaris or Linux systems' compliance with the Operating System Update Profiles that you define in Oracle Enterprise Manager Ops Center. The update profiles include both the system-defined and user-defined profiles in Oracle Enterprise Manager Ops Center.

> **Note:** Avoid running reports for system-defined profiles like Perform Reboot+Reconfigure and Perform Reboot as these profiles do not contain any updates.

Before You Begin

You can modify the update list that is applied to generate the report by selecting a white list and a black list.

A white list is the list of updates to install. To establish a white list, create a profile using the required setting. Select the white list either from the created profile or enter the update IDs separated by new lines.

For example, baseline B includes updates X, Y, and Z, and the white list has updates U, V, and W. When the Baseline Analysis Report is created, the host is marked compliant only when all six updates (X, Y, Z, U, V, and W) are present.

A black list is a list of updates that you do not want them to be installed. You create a black list by creating a policy with the specified action for the updates. Select the black list either from the created policy or enter the update IDs separated by new lines.

When a particular update in the profile is set with the policy component setting as Never for the install action, then the update is not installed. When the update is installed, it is not uninstalled or removed.

For example, when baseline A has updates X, Y, and Z, and the black list specifies only Y and Z, the system is compliant when X is installed. Even if the updates Y and Z are installed, they are not uninstalled when you run a compliance job from the report results.

To Create a Profile Analysis Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Profile Report** from the Actions pane. The Create Profile Report Wizard is displayed.
4. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the profiles.
6. Select the profiles from the list and click **Add** or **Add All** to select all the available profiles.

Figure 10–6 Selecting Profiles for Profile Analysis Report

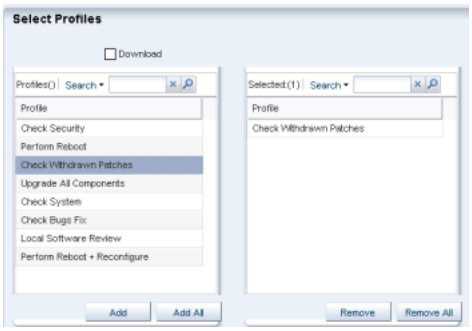

7. Check the Download check box to download the updates that must be installed for the system compliance.
8. Click **Next** to modify the update lists that are applied to the report.
9. Select any of the following White List options:
   - None: No white list.
   - Manual Input: Enter a list of updates.
   - Specify with Profile: Select a profile to import as a white list.
10. Select any of the following black list options:
    - None: No black list.
    - Manual Input: Enter a list of updates.
    - Specify with Policy: Select a policy to import as a black list.
11. Click **Next** to schedule the report.
12. Select a desired schedule to run and generate the report.
13. Click **Next** to display the Summary.

14. Review the report parameters and select one of the options as required:
    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
    - Run and Close: Runs the report and closes the wizard window.

    The report result displays under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating a Recommended Software Configuration Report

A Recommended Software Configuration provides information about the system compliance for installing a specific application, such as the Oracle 11g Database, on an Oracle Solaris or Linux operating system.

The Knowledge Base provides a list of application configuration requirements with which you can check your system compliance status.

For example, you can check the system compliance status of Oracle Solaris operating system for installing Oracle 11g Database. The report provides information about the updates that must be installed, uninstalled, or upgraded for installing the Oracle database.

For an Oracle Solaris operating system, you cannot upgrade a update component from the existing lower version to the recommended higher version. Such instances will be marked as Error in the RSC report result. In such scenarios, you cannot make the target system fully compliant with the recommended software components by the report.

Before You Begin

You can generate different types of RSCs:

- Agent-based RSC
- Database-based RSC

In an agent-based RSC, the report is generated based on the information from the target system. The dependencies for the updates are checked and downloaded when required. This report takes time to generate because it checks dependencies and downloads updates that must be installed.

In a database-based RSC, the report is generated based on the target system information that is available on the database of the Enterprise Controller. The dependencies are not checked and required updates are not downloaded. This type of report is generated quickly.

To Create a Recommended Software Configuration Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Recommended Software Configuration Report** from the Actions pane. The Create Recommended Software Configuration Report Wizard is displayed.
4. Define the report parameters:
    - Report Name: The name of the report.
    - Description: A description of the report.

- Schedule: Select **Create Schedule** to schedule the report.
- Output Format: Select the output format of the report result. CSV and PDF formats are available.
- Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.

5. Click **Next** to select the recommended software configurations.
6. In Select Recommended Software Configurations, select any of the following options:
   - Run Against Database or Run Report Against Agent.

     When you select Run Report Against Agent, then click the **Download check box** to download the updates that must be installed on the target.
   - Select the **Distribution type**.
   - Select the recommended software component from the list and select the required configuration. The recommended configuration describes the prerequisite list of updates for the selected application. You can select targets of multiple distribution. For each distribution, select the corresponding RSCs. A warning message is displayed when you have not selected RSCs for a distribution.

     Note: When you have multiple distributions, then you must select RSCs for at least one distribution to continue further in the wizard. When you have not selected RSCs for a distribution, then the targets of that distribution are not in the report result.
7. Click **Next** to schedule the report.
8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.
10. Review the report parameters and select one of the option as required:
    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
    - Run and Close: Runs the report and closes the wizard window.

    The report result displays under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating an Oracle Solaris Update Compliance Report

The Oracle Solaris Update Compliance report determines whether a specific Oracle Solaris system is compliant with a particular released Update.

To Create an Oracle Solaris Update Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select Additional Reports from the Reports section.
3. Select **Solaris Update Compliance** from the Actions pane. The Solaris Update Compliance Report Wizard is displayed.

4. Specify the report parameters:
   - Name: The name of the report.
   - Description: A description of the report.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the target asset.

   The Select Targets page is displayed.
6. Add the targets by selecting them from the list on the left by clicking **Add to Target List**. Click **Next** to display the Summary page.
7. Click **Save Report** to save the report for future use. This returns you to the Reports tab, where you can run the report by selecting it from the Saved Reports section and clicking **Re-run Report**.
8. Click **Run Report** to run and display the report.
9. Click **Export to CSV** to export the report result.
10. Click **Done** to close the report.

Creating an Incident Compliance Report

Incidents are the updates that are available for an application or feature. Incidents apply to one or more packages or RPMs. You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

Incident Compliance Report for Oracle Solaris or Linux

You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

To Create an Incident Compliance Report for Oracle Solaris or Linux

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Incident Report** from the Actions pane. The Create Incident Compliance Report Wizard is displayed.
4. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Criteria: You can select either Select Updates or Filter Updates, for selecting the updates that are used as a comparison. Depending on the selection of criteria the wizard steps vary.
   - Compliant: Select either **Compliant** or **Non-compliant** for compliance status.
   - Schedule: Select **Create Schedule** to schedule the report.

- Output Format: Select the output format of the report result. CSV and PDF formats are available.
- Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.

5. Click **Next** to select the updates.
6. When you have selected Select Updates in the previous step, the list of available incidents is displayed.

Figure 10–7 Selecting Incidents for Incident Compliance Report

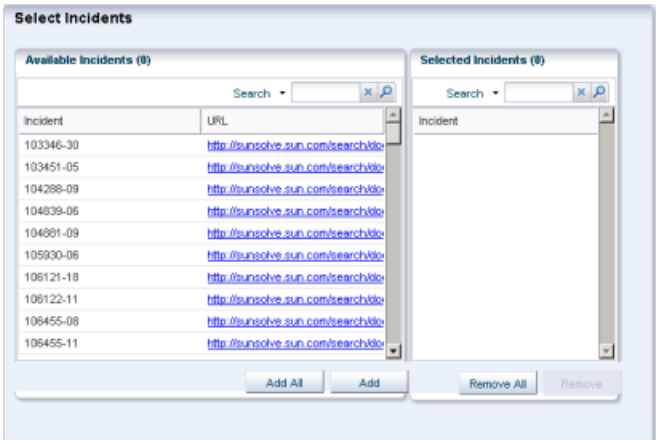

7. Select the incidents and click **Add** or **Add All** to select all the listed incidents.
8. If you have selected Filter Updates in the first step, then select the following:
   - Select Packages: You can select the updates based on the category, update type and releases date. Select the packages and click **Add** or **Add All** to select all the packages in the Available Packages list. Click **Next** to select the CAN IDs.
   - Select CAN IDs: Select from the list of Available CAN IDs. Click **Add** or **Add All** as required.
9. Click **Next** to schedule the report.
10. Select a desired schedule to run and generate the report.
11. Click **Next** to display the Summary.
12. Review the report parameters and select one of the options as required:
    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
    - Run and Close: Runs the report and closes the wizard window.

    The report result displays under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating an Incident Compliance Report for Microsoft Windows

You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

The Incident Compliance Report provides information about whether your systems are compliant with the Windows updates incidents. This report displays the number of

systems to which the selected Windows updates apply, how many systems have the updates installed, and how many systems require the updates to be installed to make the systems compliant. You can create a Windows update job based on the results of an Incident Compliance Report.

To Create an Incident Compliance Report for Windows

1. Select **Reports** from the Navigation pane.
2. Select Windows Incident Compliance Report from the Actions pane. The Windows Incident Compliance Report Wizard is displayed.
3. Specify the report parameters. They include:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Specify the Windows OS updates on which to run the report: You can specify filter criteria such as Category, Severity, Superseded, and Release Date for Windows OS updates, or you can select specific Windows OS updates to run the report.

     Click **Next**.
4. Based on your selection in Step 3, either the Define Updates Filter window is displayed or the Select Updates window is displayed. When the Define Updates Filter window is displayed, go to Step 6. When the Select Updates window is displayed, go to Step 7.
5. Make your selections in the Define Updates Filter screen. They include:
   - Category: Includes Application, Critical Updates, Definition Updates, Drivers, Service Packs, Security Updates, Tools, Update Rollups, and WSUS Infrastructure Updates. You can select either All available updates under all category or Selected categories only. Use the Control key on the keyboard to select multiple items in the list under Selected category only.
   - Severity: Includes Critical, Important, Moderate, Low, and Default. You can select either All updates with any severity or Selected severities only. Use the Ctrl key on the keyboard to select multiple items in the list under Severity.
   - Superseded: Enables you to select all or just the most recent updates.
   - Release Date: Refers to the date that the update updates were released. You can select the range of release dates to include in your report by filling in the From and To fields. Click **Next**. Go to Step 8.
6. Make your selections in the **Select Updates** window. Under Search, Select All enables you to include a bulletin ID, article ID, and title in your search, or you can select specific fields to narrow your search. Use the Control key on the keyboard to make multiple selections in the list under Available Windows Software Updates.

   Click **Add to Updates List**, and then click **Next** to select the targets.
7. Add the targets by selecting them in the list of Available Items, by clicking **Add to Target List**.

   Click **Next** to display the Summary page.
8. Click **Finish** to run the report.

   The results of the report are displayed under the Report Results list.

Creating a Host Compliance Report

You can run a host compliance report to determine whether the hosts are compliant with security and bug fix incidents. This report displays the number of updates that are applicable to each system, and whether the updates are installed or must be installed to make the system compliant. You can also create an update job based on the results of a Host Compliance Report.

Host Compliance Report for Oracle Solaris or Linux

The Host Compliance Report provides information if your systems are compliant with update incidents.

To Create a Host Compliance Report for Oracle Solaris or Linux

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Host Compliance Report** from the Actions pane.

   The Create Host Compliance Report Wizard is displayed.
4. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Update Level: Select whether you want the compliant status for Security and Bug Fixes or for only Security Updates.
   - Compliance: Select either **Compliant** or **Non-Compliant**.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them in the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to schedule the report.
6. Select a desired schedule to run and generate the report.
7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:
   - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
   - Run and Close: Runs the report and closes the wizard window.

   The report result displays under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating Host Compliance Report for Microsoft Windows

The Host Compliance Report for Windows provides information if your systems are compliant with the Windows updates incidents. This report displays the number of Windows updates that are applicable to each system, and whether the updates are installed or must be installed to make the system compliant. You can also create a Windows update job based on the results of a Host Compliance Report.

To Create a Host Compliance Report for Windows

1. Select **Reports** from the Navigation pane.
2. Select **Windows Host Compliance Report** from the Actions pane. The Windows Host Compliance Report Wizard is displayed.
3. Specify the report parameters. They include:
   - Report Name: A name for the report.
   - Description: A description of the report.
   - Specify the Windows OS updates on which to run the report. You can specify filter criteria such as Category, Severity, Superseded, and Release Date for Windows OS updates, or you can select specific Windows OS updates to run the report.
4. Click **Next**. Based on your selection in Step 3, either the Define Updates Filter window is displayed or the Select Updates window is displayed. When the Define Updates Filter window is displayed, go to Step 5. When the Select Updates window is displayed, proceed to Step 6.
5. Make your selections in the Define Updates Filter screen. They include:
   - Category: Includes Application, Critical Updates, Definition Updates, Drivers, Service Packs, Security Updates, Tools, Update Rollups, and WSUS Infrastructure Updates. You can select either All available updates under all category or Selected categories only. Use the Control key on the keyboard to select multiple items in the list under Selected category only.
   - Severity: Includes Critical, Important, Moderate, Low, and Default. You can select either All updates with any severity or Selected severities only. Use the Ctrl key on the keyboard to select multiple items in the list under Severity.
   - Superseded: Enables you to select all or just the most recent updates.
   - Release Date: Refers to the date that the updates were released. You can select the range of release dates to include in your report by filling in the From and To fields. Click **Next**. Go to Step 8.
6. Make your selections in the Select Updates window. Under Search, Select All enables you to include a bulletin ID, article ID, and title in your search, or you can select specific fields to narrow your search. Use the Control key on the keyboard to make multiple selections in the list under Available Windows Software Updates. Click **Add to Updates List**. Click **Next**.
7. Add the targets by selecting them from the list of Available Items. Click **Add to Target List**. Click **Next** to display the Summary page.
8. Click **Finish** to run the report.

   The results of the report are displayed under Report Results list.

Creating a Common Vulnerability and Exposure (CVE) Report

The CVE report provides information about incidents that are related to specific Common Vulnerability and Exposure Identifiers (CVE IDs) and the systems that must have these incidents installed. CVE IDs are unique, common identifiers for publicly known security vulnerabilities. The updates from a list of vendors are published as common vulnerabilities and security exposure incidents. CVEs are identified by a candidate ID (CAN ID).

The following types of CVE reports are available:

- Non-Compliant: Displays systems that might require action. A Target is non-compliant when no Common Vulnerability and Exposure Identifiers (CVEs) are found on it.
- Compliant: Lists the CVE components installed on a target. The COMPLIANT status in the report status column indicates the target has at least one CVE installed; however, other CVE components might be required.

The COMPLIANT status of the Figure 10–11, "CVE Report Details" indicates that the target has at least one CVE installed.

To Create a CVE Report

Perform the following steps to generate a CVE Report.

1. Expand **Reports** from the Navigation pane.
2. Expand **OS Reports** and then **Solaris/Linux OS Updates** to select **CVE Report**.
3. Select **Create CVE Report** from the Actions Pane. The **Create CVE Report** Wizard is displayed.
4. Define the report parameters and click **Next**.
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Compliance: Select **Compliant** type.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.

Figure 10–8 CVE Report Parameters

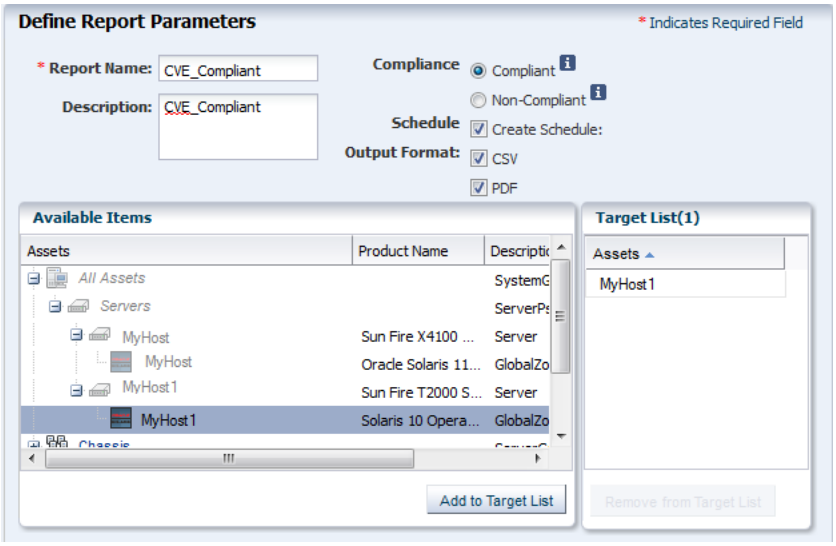

5. Select one or more CAN IDs and click **Add** or **Add All** to select all the available CAN IDs.

Figure 10–9 Selecting CAN IDs in CVE Report

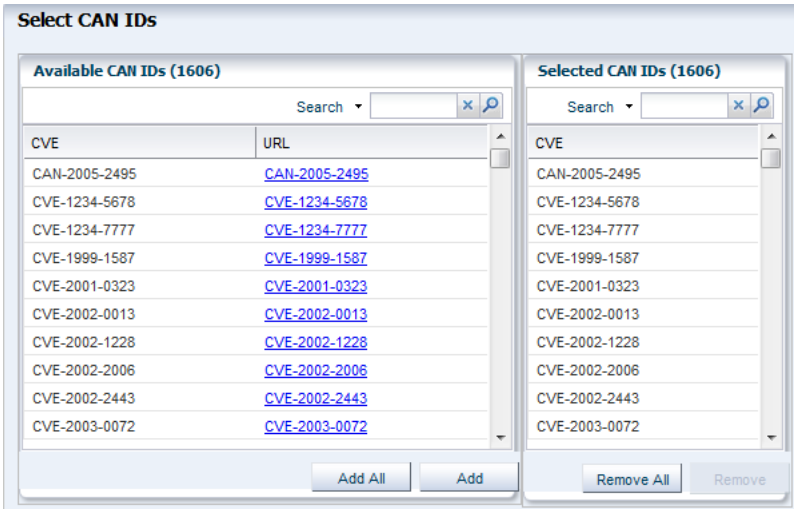

6. Click **Next** to schedule the report.
7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:
   - **Save Template and Close**: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
   - **Run and Close**: Runs the report and closes the wizard window.

   The report result appears in the Report Results in the center pane.

Figure 10–10 CVE Report Result

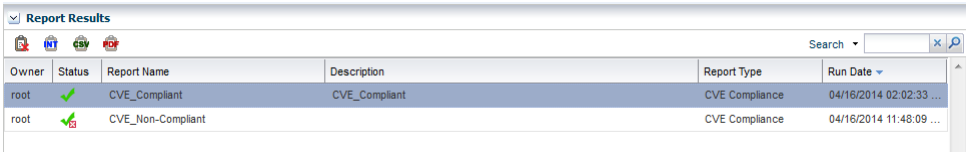

Double-click the selected CVE Report to display the CVE Report Details.

Figure 10–11 CVE Report Details

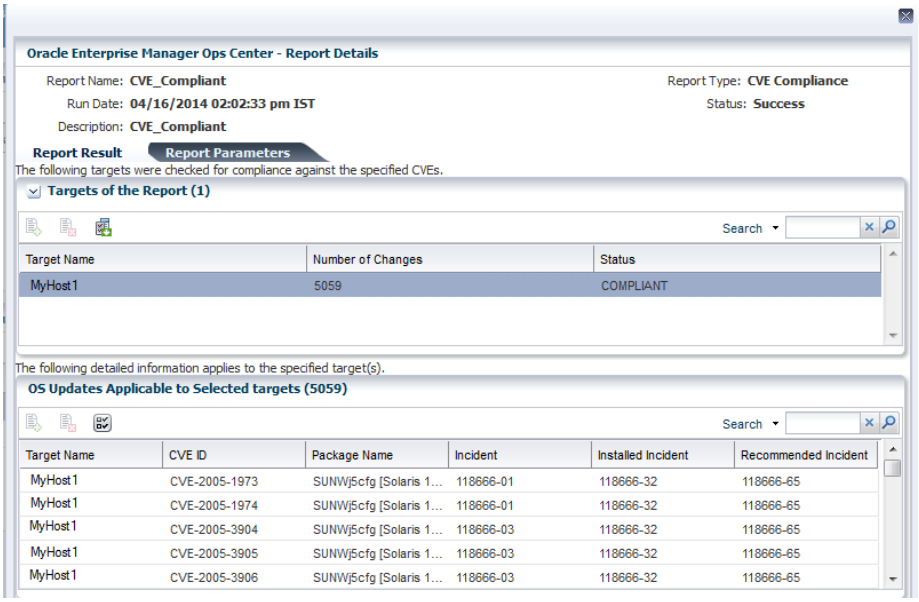

Note: When running a CVE Compliance Report with 'Compliant' selected as the option for Compliance, the Status column will show 'COMPLIANT' if there are changes found for the target.

When running a CVE Compliance Report with 'Compliant' selected as the option for Compliance, the Status column will show 'NONCOMPLIANT' if no changes are found for the target.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating a System Catalog Report

A System Catalog Report lists the current catalog of one or more systems. A system catalog contains a list of operating system software components that are installed on a managed system. Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems.

After an operating system is available and selected, you can view and modify the catalogs, and create historical catalogs. Historical catalogs are snapshots of the system. The software automatically takes a snapshot of the operating system after running a job on the operating system, including when you discover and manage the operating system. A snapshot is stored as a catalog with the time stamp and job details after every update job that you run on a system.

You can create a new catalog at any time and use it to record the state of a system. Catalogs enables us to rollback our system to any previous configuration or to create a profile that is used to apply a consistent configuration throughout our datacenter.

See Chapter 14, "Operating System Updates" for more information on System Catalogs.

To Create a System Catalog Report

1. Select **Reports** from the Navigation pane.

2. Select **System Catalog Report** from the Actions pane.

   The System Catalog Report Wizard is displayed.

3. Define the report parameters, including:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add one or more targets by selecting them in the list of Available Items by clicking **Add to Target List**.
4. Click **Next** to display the Schedule.
5. Select a desired schedule to run and generate the report.
6. Click **Next** to display the Summary.
7. Review the Summary, then click **Run and Close**.

   The report result is displayed under the Report Results in the center pane.

See Viewing a Report Result for more information about viewing a report result and generating a compliance job from the result.

Creating System Information Reports

Create a system information report to obtain the information on assets such as operating systems, servers, chassis, logical domains, global zones, non-global zones, and M-Series servers. The information on assets include details like architecture, type, host id, host name, logical units, version, description and so on.

To Create a System Information Report

1. Select **Reports** from the Navigation pane.
2. Select System Information Reports from the Reports section.
3. Select **Create System Information Report** from the Actions pane. The Create System Information Report Wizard is displayed.
4. Define the report parameters, including:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Schedule: Select Create Schedule to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Targets: Add one or more targets by selecting them in the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to display the Attribute Selection.
6. Select one or more attributes from the list and click **Add** or **Add All** to choose all the attributes.

   Click **Next** to display the Attribute Filters.

7. To set a filter for an attribute, select the attribute and specify its condition. Click the **Add icon** to set filters for other attributes.

   Click **Next** to display the Schedule, when you are finished setting the filters.

8. Select a desired schedule to run and generate the report.

9. Click **Next** to display the Summary.

10. Review the report parameters and select one of the options as required:

    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
    - Run and Close: Runs the report and closes the wizard window.

    The report results are displayed under the Report Results in the center pane.

To view the report, select it from the Reports Results section and then click one of the icons to choose the format: View interactively, View CSV, or View PDF.

See Viewing a Report Result for more information about generating a compliance job from the result.

Creating Oracle Engineered Systems Reports

Using Oracle Enterprise Manager Ops Center, you can discover and manage Oracle Engineered Systems. The Oracle Engineered Systems report is all about viewing the rack setup for each of the rack within the system including the asset details related to the firmware. You can view and access multiple Engineered Systems from a single datacenter through Oracle Enterprise Manager Ops Center. You can also view the Engineered System's assets, incidents reported from the Engineered Systems, and Service Requests. You can generate reports for all datacenter assets, including Engineered Systems. This also displays the number of assets in the rack by their types, such as compute nodes, switches, and storage nodes in the system.

To Create an Oracle Engineered Systems Report

You can generate and view the report for multiple engineered systems. Perform the following steps to create an Oracle Engineered Systems Report.

1. On the Navigation pane, click **Reports**.

Figure 10–12 Create Oracle Engineered System Reports

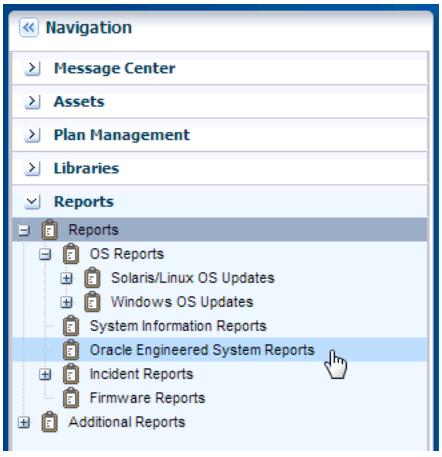

2. Click **Oracle Engineered Systems Reports**.

3. On the Actions pane, click **Create Oracle Engineered Systems Report**.
4. In the **Define Report Parameters** wizard, enter a name and description for the report.

   The Schedule and Output Format are checked by default.

Figure 10–13 Oracle Engineered Systems Report Parameters

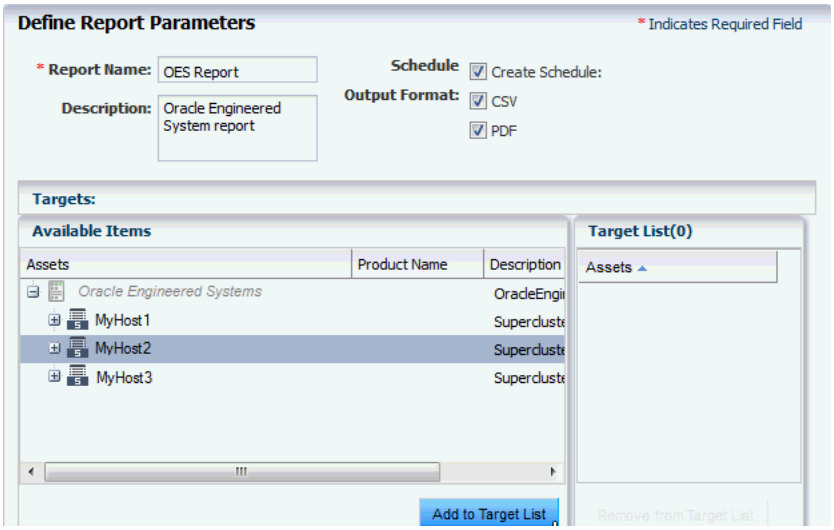

   - Select **Create Schedule** if you want to run the report later or on a recurring schedule.
   - Select the output formats of the result that will be generated for the report.
5. In the Targets section, select the asset for which you want to run the report and click **Add to Target List**.
6. Click **Next**. The Schedule wizard is displayed.
7. Select a schedule for the report. You can schedule the report to run on the following instances:
   - Now: Runs the report immediately.
   - At a later date/time: Select a date and time to generate the report.
   - On a Recurring Schedule: Select the month and day when you want to generate the report. Select the Start Time, End Time and Number of Hours between runs. This is to set the number of times the report is generated between the specified start and end time. For example, if you set the start time at 6.00 a.m, end time at 12.00 a.m, and the number of hours between runs as 2, then the report is run at 6.00 a.m, 8.00 a.m, 10.00 a.m, and 12.00 a.m.
8. Click **Next**. The Summary wizard is displayed.
9. Verify the report parameters and click one of the options as required:
   - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
   - Run and Close: Runs the report and closes the wizard window.

Figure 10–14 Oracle Engineered Systems Report Results

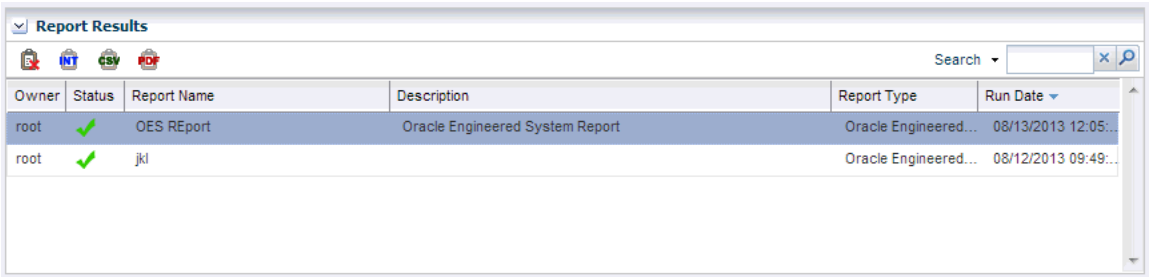

Creating Incident Reports

You can create incident reports to obtain information about incidents.

The following types of incident reports are available:

- The Incident Summary Report: This is a historical report that summarizes information about all alerts and incidents for a specified category, such as alarm state, alarm owner, asset type, date range, severity levels, and affected asset groups.
- The Incident Detail Report: This contains detailed information about one or more incidents. In addition to a summary, the report includes an audit trail consisting of state-change annotations, alert annotations, suggested-fix annotations, comment annotations, operation annotations.

  Each incident has four pages, after the summary page. They are:

  - Details of the incident
  - Suggested actions, if any
  - Alerts History
  - Any annotations that are associated with the incident

When you create a report, you can save the report as a template, or you can generate the report. After a report is created, you can view the report, re-run the report to get updated information, or save it as a template.

To Create an Incident Summary Report

1. Select **Reports** in the Navigation pane.
2. Select Incident Reports, and then click Incident Summary Report.
3. Select **Create Incident Summary Report** from the Actions pane.
4. Define the report parameters:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Schedule: Select **Create Schedule** to schedule the report.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Select Assets: Select either all assets or specific assets.
5. Click **Next** to specify the Incident Parameters.
6. Define the incident parameters by providing the following information:

- To create a historical report of all incidents and the date that each incident was detected, select **All Creation Dates**.
- To create a summary report for incidents detected during a specific date range, select **Range of Creation Dates**, then enter the beginning date in the From field and the ending date in the To field.
- To filter by severity level, owner, or state, highlight the fields to include in the report. Use Ctrl+Enter to select multiple options.
- To filter by one or more criteria, add the criterion in the Description Contains field. Use a comma-delimited list for multiple criteria. For example, FileSystemUsage or FileSystemUsage, SwapUsage.

7. Click **Next** to display the Schedule.
8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.
10. Review the summary and select one of the options as required:
    - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
    - Run and Close: Runs the report and closes the wizard window.

    The report results are displayed under the Report Results in the center pane.

To Create an Incident Detail Report

1. Select **Reports** in the Navigation pane.
2. Select Incident Reports, and then click Incident Detail Report.
3. Select **Create Incident Detail Report** from the Actions pane.
4. Define the report parameters:
    - Report Name: The name of the report.
    - Description: A description of the report.
    - Output Format: Select the output format of the report result. CSV and PDF formats are available.
5. Click **Next** to display the Summary.
6. Review the summary and select **Save Template** and **Close**. This saves the report as a template and closes the wizard. You can use the report template to generate the report later.

   The report results are displayed under the Report Results in the center pane.

> **Note:** The Incident Compliance Report refers to the Microsoft Windows operating system incidents, while Incident Summary Report and Incident Detail Report refers to alerts and alarms that are raised by Oracle Enterprise Manager Ops Center incidents.

Creating a Firmware Report

Firmware Compliance Reports enables us to maintain consistent firmware versions across the datacenter. The Firmware Report feature compares the firmware images specified in a firmware profile to the firmware images installed on hardware assets.

The report indicates whether the firmware on the asset complies with the profile's specifications. You can update the firmware on any non-compliant asset by clicking the Make Targets Compliant button in the Interactive report.

To Create a Firmware Report

1. Select **Reports** in the Navigation pane.
2. Select **Firmware Reports**.
3. Select **Create Firmware Report** from the Actions pane. The Create Firmware Report Wizard is displayed.
4. Define the report parameters, including:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Schedule: If you do not plan to create this report routinely, deselect the **Create Schedule** option.
   - Output Format: Select the output format of the report result. CSV and PDF formats are available.
   - Profile: Select the firmware profile for a service processor or for a storage component, such as a RAID controller, expander, or disk.
5. Click **Next** to display the Select Targets.
6. Select the targets you want to test against the profile. Select the asset from the Available Items hierarchy and click **Add to Target List**. When you have selected all the targets, click **Next** to display the Schedule page.
7. Select a desired schedule to run and generate the report.
8. Click **Next** to display the Summary.
9. Review the summary and click **Run and Close** to create the report job.

   The report job starts at the time you specified and compares the values in the profile to the existing values on the targets you selected. The report displays whether a target asset is compliant, not compliant, or not applicable:

- A compliant asset has the firmware images specified in the profile.
- A non-compliant asset does not have the same firmware images as specified in the profile.
- A non-applicable asset indicates that a firmware image in the profile does not match the model of service processor in the asset. This condition can occur when either the profile does not recognize the model that the service processor is reporting or the profile includes firmware images that are not designed for the service processor.

  a. Compare the model of the service processor displayed in the asset's Summary tab with the model of the service processor included in the profile. If they are different, add the name in the profile to the asset's data.

     See *Oracle Enterprise Manager Ops Center Administration Guide* for information about adding a product alias.

  b. When the firmware profile was created, only images that matched the service processor could be included. However, if the service processor did not report all the firmware types it supported, an image that did not match the service processor could have been included in the profile. To update the software with

all the service processor's supported firmware types, use the **Refresh** action to update the information about the service processor. When the job is completed, view the service processor's **Summary** tab to see all firmware types.

c. Repeat the procedure to create a firmware report.

Creating Additional Operating System Reports

Use Additional Operating System Reports to obtain information from Service Pack Compliance Report, Distribution Update Report, and Package Compliance Report. You can export report results to CSV format.

Creating a Distribution Update Report

A Distribution Update Report provides a mapping between selected updates, and CVEs and selected distributions to find out whether the updates are installed. This report determines whether a specific distribution like SOLARIS10_SPARC has been updated with specific updates, or CVEs.

To Create a Distribution Update Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Distribution Update Report** from the Actions pane. The Distribution Update Report Wizard is displayed.
4. Define the report parameters. They include:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Category: Select the required categories.
   - Type: Select required types. Package and Update types are available.
   - Released: Mention the start date and the end date.
5. Click **Next** to display the Distributions.
6. Select the required Distribution and click **Add** or **Add All** to add distributions. Click **Remove** or **Remove All** when you do not require any distribution.
7. Click **Next** to select the updates.
8. Click **Next** to select the packages.
9. Click **Next** to select the CVEs. Click **Add** or **Add All** to add CVEs and **Remove** or **Remove All** when you do not require any CVE.
10. Click **Next** to display the Summary.
11. Review the Summary and select one of the options as required:
    - Save Report: Saves the report as a template and closes the wizard.
    - Run Report: Runs the report and closes the wizard window.

    The report result is displayed under the Report Results in the center pane.

Creating a Service Pack Compliance Report

A Service Pack Compliance Report provides information on updates created by the publication and release of a service pack by a vendor. This report enables you to determine whether the target system has the latest service package installed that is provided by the vendors.

To Create a Service Pack Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Service Pack Compliance Report Creation** Wizard from the Actions pane. The Service Pack Compliance Report Creation wizard is displayed.
4. Define the report parameters. They include:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Status: Select the compliant or non compliant status.
   - Services: Select the required services.
5. Click **Next** to select the Targets. Add the targets by selecting them from the list of Available Items by clicking **Add to Target List.**
6. Click **Next** to display the summary.
7. Review the Summary and select one of the options as required:
   - Save Report: Saves the report as a template and closes the wizard.
   - Run Report: Runs the report and closes the wizard window.

   The report result is displayed under the Report Results in the center pane.

Creating a Package Compliance Report

A Package Compliance Report provides a mapping between the selected packages and the selected target systems to find out the installed packages.

To Create a Package Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Package Compliance Report Creation** Wizard from the Actions pane. The Package Compliance Report Creation Wizard is displayed.
4. Define the report parameters. They include:
   - Report Name: The name of the report.
   - Description: A description of the report.
   - Status: Select the compliant or non compliant status.
   - Level: Select security updates or security and bug updates.
5. Click **Next** to select the Targets. Add the targets by selecting them from the list of Available Items by clicking **Add to Target List.**
6. Click **Next** to select the Packages.
7. Click **Next** to display the summary.

8. Review the Summary and select one of the options as required:
   - Save Report: Saves the report as a template and closes the wizard.
   - Run Report: Runs the report and closes the wizard window.

   The report result is displayed under the Report Results in the center pane.

Related Resources for Reports

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- Oracle Enterprise Manager Ops Center Documentation Library at http://docs.oracle.com/cd/E40871_01/index.htm.
- For an example, see *Creating System Catalog Reports* in the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.

11 Hardware

This chapter describes hardware management features that are available in Oracle Enterprise Manager Ops Center:

Introduction to Managing Hardware Assets

Oracle Enterprise Manager Ops Center provides comprehensive lifecycle management for the hardware assets in your data center. The hardware assets can be handled individually or as a group. After the discovery and management of the hardware assets, as described in Asset Management, you can configure and then monitor them to gather the information to maintain them.

Configuring Hardware Assets

Use these deployment plans:

- Install Server
- Update BIOS Configuration
- Configure Service Processor
- Configure Server Hardware and Install OS

- Configure RAID
- Configure M-Series Hardware, Create and Install Domain
- Configure and Install Dynamic System Domain

All the plans are based on hardware resource profiles. See Hardware Resource Profiles

Monitoring Hardware Assets

As soon as a hardware asset is managed, Oracle Enterprise Manager Ops Center starts to monitor it, according to the asset type's monitoring profile. The center pane displays information for a selected asset in a series of tabbed windows. The tabs and the type of information is specific for the asset type but, in general, Oracle Enterprise Manager Ops Center reports the following:

- Health status
- Power state
- Power usage
- Hardware variables and connectivity

You can change the monitoring thresholds in the standard profile to specify the conditions that generate an alert. You can also create custom profiles with different rule sets and alert parameters and apply the profile to a specific system, a group of homogeneous systems, or a group that you define.

For more information about monitoring policies and rules, see Chapter 4.

For information about monitoring power consumption, see Energy Tab.

Maintaining Hardware Assets

Based on your observations, you can control your hardware assets and do the following actions:

- Update management credentials. See Chapter 2, "Asset Management"
- Setting and Changing the Power Policy
- Power systems on and off
- Place in maintenance mode
- Reset a server
- Get access to the serial console
- Use locator lights to identify a specific asset
- Check firmware compliance and update firmware

Use these deployment plans or customize them:

- Updating Firmware
- Update Storage Appliances
- Update BIOS Configuration

Roles for Hardware Management

Table 11–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See

the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 11–1 Hardware Roles and Permissions

Task	Role
Configure and Deploy Server	Server Deploy Admin
Install Server	Server Deploy Admin
Configure RAID	Server Deploy Admin
Update Management Credentials	Security Admin
Importing and uploading firmware images	Storage Admin
Edit Attributes	Asset Admin
Power On, Power Off, Power on with Net Boot	Asset Admin
Set Power Policy	Asset Admin
Reset Servers, Reset Service Processors, Refresh	Asset Admin
Locator Light On/Off,	Asset Admin
Snapshot BIOS Configuration, Update BIOS Configuration	Asset Admin
Update Firmware	Update Admin
Simulate a firmware update	Update Admin Update Sim Admin
Launch LOM Controller	Asset Admin
Edit Tags	Asset Admin

Actions for Hardware Management

After you manage your assets, you can perform the following actions:

- Use a resource profile to configure a hardware asset
- View utilization of systems
- Modify energy consumption
- Reset a server
- Power systems on and off, including a forced power off
- Place in maintenance mode
- Get access to the serial console
- Use locator lights to identify a specific asset
- Check firmware compliance
- Update firmware
- Use a provisioning profile to update firmware

Location of Hardware Information in the User Interface

Hardware assets are visible in the All Assets section of the user interface. For assets with a service processor, both the service processor and the system are included, as

shown in Figure 11–1.

Figure 11–1 Managed System

Each type of hardware is also added to the appropriate group in the Resource Management view, as shown in Figure 11–2.

Figure 11–2 Example of Resource Management View

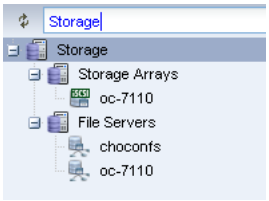

Table 11–2 shows where to find information.

Table 11–2 Location of Hardware Information in the BUI

To See	Location
All servers	Expand Servers in the Assets pane. The center pane lists up to 50 servers.
Details for a specific server	Expand Servers in the Assets pane. Then select one of the servers. The center pane includes a series of tabbed displays.
Monitoring Rules for a specific server	Expand Servers in the Assets pane. Then select one of the servers. Then select the Monitoring tab.
Monitoring rules for a hardware type	Expand Plans and then Operational Plans and then Monitoring Policies. Select a type. The center pane shows all of the rules and thresholds.
Current firmware version	Expand Servers in the Assets pane. Then select one of the servers. Then select either the Summary tab or the Hardware tab. Each one includes the Firmware table.
A report of all attributes	Expand Reports and then select System Information Report.

Profiles for Hardware Management

The work flow of an asset deployment can be captured and enacted in a repeatable fashion using plans and the profiles included in the plans. Oracle Enterprise Manager Ops Center provides default profiles for configuring hardware asset types consistently. You can use the default profiles, make copies of the profile to edit, or create new ones.

Deployment proceeds from configuring the hardware, installing the correct firmware, provisioning the OS, and applying the required updates. The deployment is not only for servers with an OS installed but also for chassis, racks, power distribution units, and M-Series servers.

Hardware Resource Profiles

Use hardware resource profiles in a deployment plan to configure, install, or update systems. The Oracle Enterprise Manager Ops Center software provides the following hardware resource profiles:

- Configuring the Service Processor
- Configuring a RAID Controller
- Configuring a Dynamic System Domain
- Configuring a Rack and Placing Components

You can use the default profiles, make copies of the profile to edit, or create new ones. Use the profiles in a deployment plan to configure your hardware assets.

Firmware Provisioning Profiles

Hardware depends on firmware to perform operations. Part of monitoring and managing a hardware asset is to make sure that it has the appropriate version of firmware. Oracle Enterprise Manager Ops Center uses firmware profiles to provision, or update, firmware on each type of asset.

A firmware profile is a set of actions and values that define how to provision one or more assets and specifies one or more firmware images. A firmware profile updates existing firmware assets completely and consistently. When you apply a deployment plan that contains a firmware profile, Oracle Enterprise Manager Ops Center compares the versions of each firmware image specified in the profile with the versions of the existing firmware on the asset and then takes the action you specify in the profile.

See Firmware Compliance Reports and Firmware Provisioning.

Configuring the Service Processor

The **Declare Unconfigured Asset** action includes the service processor in the Oracle Enterprise Manager Ops Center environment. You can configure only unconfigured service processors, that is, a processor in its factory default state. Use a deployment plan to configure the service processor:

- Configure Service Processor
- Update BIOS Configuration
- Configure and Deploy Server
- Install Server

An alternative to specifying the configuration in the profile, is to duplicate an existing configuration. You can create a snapshot of the BIOS of a working service processor, which creates also creates profile. You then apply the profile to an unconfigured service processor.

Creating a Service Processor Configuration Profile and Plan

1. Expand Plans in the Navigation pane and the select Profiles and Policies.
2. Expand Service Processors.
3. Click **Create Profile**. The first step of the wizard is displayed, as shown in Figure 11–3. Depending on the subtype and target you select, more steps are added.

Figure 11–3 Create Profile - Service Processor

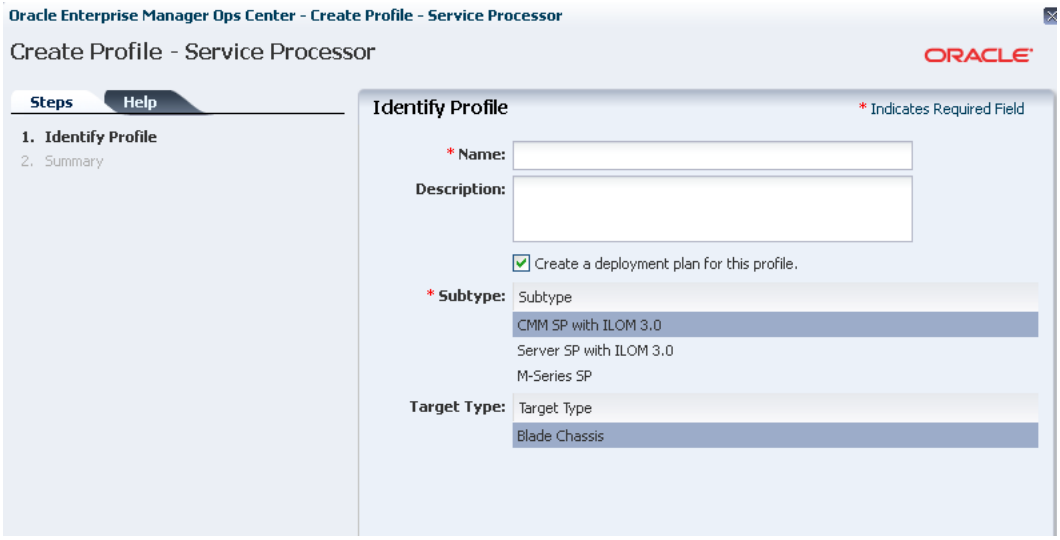

4. Complete the specification of the service processor and click **Finish**.

The new profile and plan are available from the Assets pane.

Creating a BIOS Configuration Profile and Plan

1. Expand Plans in the Navigation pane and the select Plans.
2. Expand Update BIOS Configurations.
3. Click **Create Plan from Template**. Figure 11–4 shows the first step of the wizard, which includes the Update BIOS profile.

Figure 11–4 Create Plan - BIOS

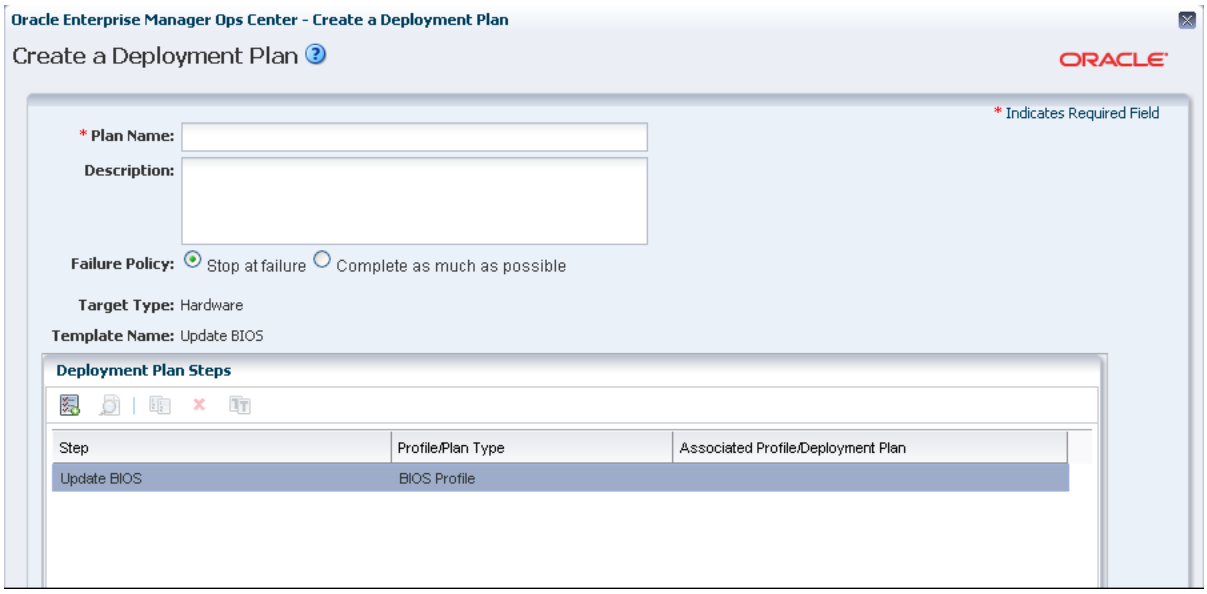

4. Click the Update BIOS profile.

5. Click the **Create Profile** icon. Figure 11–5 shows the window for creating the profile.

Figure 11–5 Create Profile-BIOS Configuration

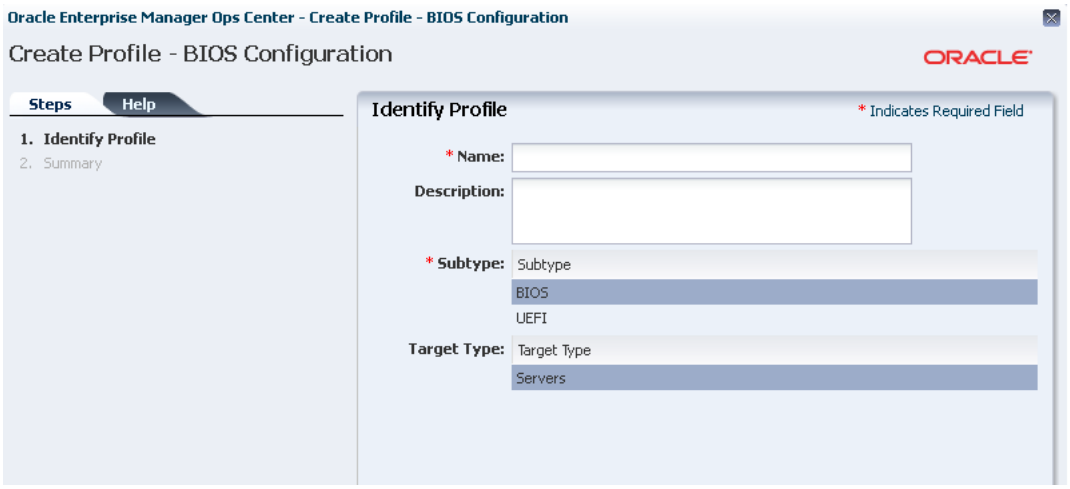

6. Enter a name and description and then select the type of BIOS and the type of target.
7. Click **Next** to review and then click **Finish** to submit the job.
8. When the job is completed, return to the Create a Deployment Plan window.
9. Specify a name for the plan and select the new profile.
10. Click **Save** to submit the job.

The new profile and plan are available from the Assets pane.

Creating a Snapshot of a Service Processor Configuration

1. Expand Assets in the Navigation pane and select Servers.
2. Select the configured server.
3. Click the **Hardware** tab in the center pane.
4. In the Component Navigation section, select **Service Processor**. The Service Processor Configuration section displays the configuration attributes and values.
5. In the Service Processor Snapshots section, click the **Create Snapshot** icon to display the window shown in Figure 11–6.

Figure 11–6 Create Service Processor Configuration Snapshot

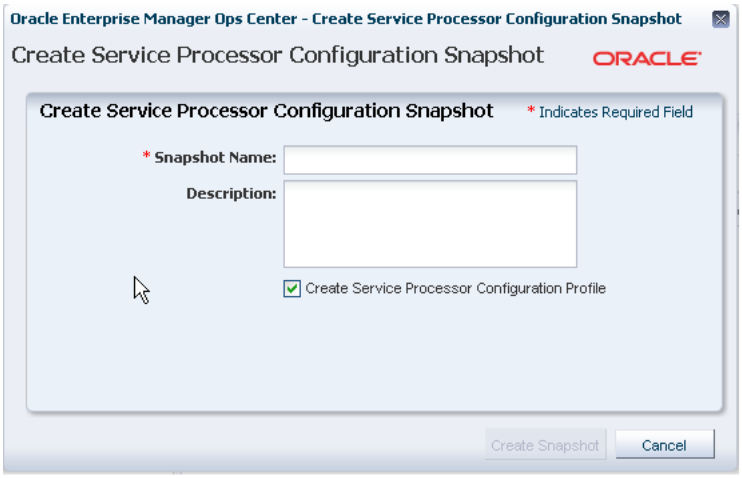

Note: An alternative method of displaying this window if you are familiar with the current configuration is to select the server and then choose **Snapshot BIOS Configuration** in the Actions pane.

6. Specify a name for this snapshot and accept the default of creating a profile from the snapshot.
7. Click **Create Snapshot** to submit the job.
8. When the job completes, the snapshot is located in the EC local library.

Applying a Snapshot of a Server Processor Configuration

This procedure applies a snapshot of one server's service processor to another, unconfigured service processor. Use the **Update BIOS Configuration** action or use the snapshot profile as a step in the Install Server deployment plan.

View a Service Processor Snapshot

To verify the snapshot configuration before you attempt to configure a new server, view the snapshot.

1. Expand Libraries in the Navigation pane and then select the current EC local library.
2. In the center pane, scroll down to the BIOS Configuration section of the library.
3. Select the snapshot and then select the **View Snapshot** icon.
4. You can apply the snapshot directly to a server or you can use a deployment plan.

Apply a Service Processor Snapshot to a Server

1. Expand Assets in the Navigation pane and then select Servers.
2. Select the unconfigured server.
3. Click **Update Service Processor** in the Action pane. A one-step wizard is displayed, including the list of snapshots in the EC local library that are appropriate for the selected server.
4. Select the snapshot from the list.

5. Click **Next** to review the Summary and then click **Finish** to submit the job.

When the service processor is configured, the server reboots.

Apply a Service Processor Snapshot Profile to a Server

1. Expand Plans in the Navigation pane and then select Profiles.
2. Expand Update BIOS Configuration and select the profile for the snapshot.
3. Review the profile in the center pane. To modify the profile, click the **Edit Profile** icon.
4. Create a plan that includes this profile as its only step or edit an existing deployment plan to include this profile.

Configuring a RAID Controller

Oracle Enterprise Manager Ops Center provides a default profile for configuring RAID controllers. You can configure and update the hardware devices that can be done only through the host OS. This profile provisions a reduced OS image along with the management pack on the target to configure RAID.

Caution: When you reconfigure an existing RAID controller, all the data on the disk is lost.

Configuring a Dynamic System Domain

You can create domains on a M-series server.

- Verify that a user account with `platadm` privilege exists on the XSCF processor and is included in the profile. Disable the audit policy for this user account, using the following command:

  ```
  setaudit -a opsadm=disable
  ```

- Verify that you can use `ssh` to log in to the XSCF processor.

To configure a Dynamic System Domain, you apply a deployment plan to an M-Series server, as described in Chapter 8. The plan includes the profile for the Dynamic System Domain.

Configuring a Rack and Placing Components

The rack asset is a group that includes the other managed assets installed in the physical rack, such as servers or compute nodes, storage arrays or appliances, switches, and power distribution units (PDU). After you create the rack, you put assets in the rack so that Oracle Enterprise Manager Ops Center can present them and manage them as if it were a physical asset.

Creating a Rack

1. Expand Assets in the Navigation pane and select Racks.
2. Click **Create Rack** in the Action pane.
3. Enter the name for the rack.
4. For a rack containing an Oracle Engineered System, specify the serial number of the rack.

5. Enter the total number of slots in the rack. A full rack has 42 slots. Default to 42.
6. You have the option to place assets into the rack immediately after this procedure ends or to perform this task at a later time.
7. Enter a description for this rack.
8. Add semantic tags that are appropriate for this rack.
9. Click **Create Rack** to submit the job.

Use the **Place/Remove Assets in Rack** action and the **Place/Remove PDU in Rack** action to specify each asset in the physical rack and its location in the rack so that Oracle Enterprise Manager Ops Center can represent the physical rack accurately to remote users.

Placing Assets in a Rack

The rack asset is a group that includes the managed assets installed in the physical rack, such as servers or compute nodes, storage arrays or appliances, switches, and power distribution units (PDU).

Use the **Place/Remove Assets in Rack** action and the **Place/Remove PDU in Rack** action to specify each asset in the physical rack and its location in the rack so that Oracle Enterprise Manager Ops Center can represent the physical rack accurately to users who cannot examine the rack. Update the type and location of the assets in the rack asset when the configuration of the physical rack changes.

Placing a Power Distribution Unit in a Rack

1. Expand Assets in the Navigation pane and select Racks from the Resource Management Views.
2. Select the rack.
3. Click **Place/Remove Assets in Rack** in the Actions pane or click the **Details** tab to navigate to the icon. The Place Assets in *name* Rack window opens, which displays a list of all managed assets in the physical rack. You can filter the list by type or attribute.
4. Click an asset and then enter its slot number.
5. Click **Place Asset in Rack**. The Assets in the Rack pane is updated to show the location.
6. Continue to select assets and place them in the rack.
7. To change an asset's slot, select it and click **Edit Placement**. Enter a new slot number.
8. To remove an asset, select it and click **Remove Asset**. The asset is deleted from the Assets in the Rack pane and is available to be placed.
9. When you are satisfied with the configuration displayed in the Assets in the Rack pane, click **Submit**.

To add PDUs to the rack asset, use the **Place/Remove PDU**s action in the same way. Select the rack and the action. The assets you have placed in the rack are shown in the Assets in Rack pane. Add each PDU.

Hardware Monitoring

The Oracle Hardware Management Agents use Simple Network Management Protocol (SNMP) to monitor your Oracle hardware and storage devices. The software uses the

Intelligent Platform Management Interface (IPMI) protocol to accesses the Oracle ILOM service processors.

A monitoring policy is a set of rules applied to an asset. If a status changes or a threshold is crossed, an alert is created. Oracle Enterprise Manager Ops Center provides default policies for each asset type. You can create new policies or modify existing policies.

Hardware Status

If a hardware asset can report a value for a hardware variable, Oracle Enterprise Manager Ops Center reports its current state and compares it to the threshold value.

- Good – The hardware asset is working properly.
- Unknown – Oracle Enterprise Manager Ops Center is unable to retrieve information from the sensor. The hardware asset is connected but is not reporting information.
- Unreachable – The hardware asset cannot be contacted. This state indicates a network problem.
- Warning Failure – Oracle Enterprise Manager Ops Center has detected a potential or impending fault condition. Take action to prevent the problem.
- Critical Failure – A fault condition has occurred. Take corrective action.
- Nonrecoverable Failure – The hardware asset has failed. Recovery is not possible.
- Faulted – The hardware asset reports a fault. Contact service personnel to repair.

Groups of Hardware Assets

Oracle Enterprise Manager Ops Center monitors hardware assets according to the monitoring profile for that type of asset. To see the default profile for monitoring a hardware type, see Hardware Monitoring.

To change a threshold value, see Editing Monitoring Rules in Chapter 4.

Connectivity Status

Connectivity is the network interface of the system. You can view information about a hardware asset's Network Interface Card (NIC).

Service Processor Details

Use the Hardware tab to view information about each component of the system:

- Name and SNMP Community
- Whether Auto DNS through DHCP is in use, and any DNS Servers
- Search Path, if any
- Time Zone
- Whether an NTP server is in use an its identifier

RAID Controller Details

- RAID volume name
- RAID level

- Number of disks
- Stripe zone
- RAID Controller ID

Oracle ZFS Storage Appliance Details

The Oracle ZFS Storage Appliance support both file storage and application use.

Dashboard Tab

The Dashboard tab reports the following hardware information:

- Name
- Description
- Current Alert Status
- Model
- Serial Number
- Management IP
- Memory
- Power
- Locator Light
- Appliance Kit Version
- Running Time
- Processor

Hardware Tab

The Hardware tab displays the appliance's firmware version and the following information for each component:

- CPU: Name, Model, Architecture, Speed, Manufacturer
- Memory: Name, Type, Size in bytes, Manufacturer, Part number, Serial number
- Network Adapters: Name or each, MAC Address, Description, Manufacturer, Part number, Serial number
- Disks: Name, Size in bytes, Manufacturer, Part number, Serial number
- Power Supply: Name, Manufacturer, Part number, Serial Number
- Fan Tray: Name, Manufacturer, Part number, Serial number

See Oracle ZFS Storage Appliance in Chapter 2 for discovery requirements.

ALOM and ILOM Servers Details

This version of the product software discovers and manages servers that use the ALOM or ILOM service processors. This is a description of the general details displayed for these servers.

Summary Tab

For server hardware, the Summary tab displays:

- Server Name
- Description
- Current Alert Status
- Model
- Serial Number
- Management Interface IP
- MAC Address
- Processor
- Memory
- Power state
  - On – The server is powered on and running.
  - Standby – The server is powered off but responds to commands.
  - Unknown – An error occurred while attempting to retrieve the power status of the hardware. The server is connected but is not returning any information on power status.
  - Unreachable –The server cannot be contacted for information about its power state. This indicates a network problem or that the server is in standby mode.
- Locator Lights state
- A table with the available Tags
- A table with the Firmware information

Hardware Tab for ALOM Servers

ALOM servers display summary information and a table with all the firmware installed.

Hardware Tab for ILOM Servers

ILOM servers include more information in the component navigation pane of the Hardware tab:

- System: Description, type, and version of all firmware installed except for disk firmware. See the Disk tab for firmware version.
- Processors or CPU: Architecture, number of Installed CPUs, Actual Power Consumption, Summary Description, number of Max CPUs, and Status. The Processors table displays for each CPU: Name, Model, Speed, Manufacturer, and Status.
- Memory: number of Installed DIMMs, Installed Size in bytes, number of Max DIMMs, Status, and Actual Power Consumption. The Memory table displays for each DIMM: Name, Size in bytes, Manufacturer, Part number, Serial number, and Status is displayed on the Memory table.
- Power or Power Supply: number of Installed Power Supplies, Actual Power Consumption in watts, Status, number of Max Power Supplies, and Max Permitted Power in watts. The Power Supplies table displays for each power supply: Name, Manufacturer, Part number, Serial number, and Status.
- Cooling or Fan Tray: number of Installed Chassis Fans, number of Installed PSU Fans, Inlet Temperature, Status, number of Max Chassis Fans, number of Max PSU

Fans, and Exhaust Temperature. The Fans table displays for each item: Name, RPM (%), and Status. Servers with older versions of ILOM list a Fan Tray group instead including the identifier of each available fan and its speed (RPM).

- Storage or Disk: Installed Disk Size in bytes, number of Installed Disks, Status, Logical Volumes, and number of Max Disks. The Disks table displays for each item: Name, Presence Status, and Status.
- Networking or Network Adapter: number of Installed Ethernet Nics, and Status. The Network Adapters table displays for each item: Name, MAC Address, Description, and Status.
- PCI Devices: a table listing the PCI Devices. The table displays for each item: Name, Card Type, Description, Vendor ID, Device ID, and Part Number. This information is not available for servers with older versions of ILOM.
- Service Processor: Name, Auto DNS Via DHCP, Search Path, Use NTP Server, NTP Server 2, SNMP Community, DNS Servers, Time Zone, and NTP Server 1. The Service Processor Snapshots displays for each item: Name, Creation Date, Description, and Created By.

M-Series Servers Details

The hardware resources in a SPARC Enterprise M-Series Server are divided into one or more logical units, called dynamic system domains. Oracle Enterprise Manager Ops Center can monitor each domain, in addition to the server hardware.

Dashboard Tab

For an M-Series server, the Dashboard tab displays:

- Number of dynamic system Domains it is supporting
- Model
- Serial Number
- Description
- Support contract
- XCP Firmware Version
- OBP Firmware Version
- XSCF Firmware Version
- Hypervisor Firmware Version
- Operator Panel Switch Status: Locked
- Current Alert Status

Summary Tab

The Summary tab adds details to the information in the Dashboard tab. For the Power status, the reported status is for the server's domains. When any domain is powered on, the status is reported as powered on. When all domains are powered off, the Summary tab shows a status of Powered Off; the M-Series server itself remains powered on.

You can find the following information on the Summary tab:

- Name
- Model

- Serial Number
- Management IP
- MAC Address
- Current Alert Status
- Power
- Locator Light
- Notification
- All firmware versions including Description, Type and Version.
- The table Domain displays for each item: Name, Model, Health, Power, Locator Light, Notification

Hardware Tab

The Hardware tab shows the state of the server or, if a Dynamic System Domain is selected, the state of that domain. At the System level, the Hardware tab includes the following information:

- Model
- Serial Number
- State
- Power
- Locator Light
- Notification
- Operator Panel Switch State

For the M-5000 server, the System level of the Hardware tab also includes:

- The Unallocated Resources table lists all the physical system boards and their status: PSD ID, Assignment Status, Power Status, Connection Status, Diagnostics Status, and Operational Status
- The Allocated Resources table lists all domains that are using the physical system boards and their status: Domain ID, PSB ID, XSB ID, LSB ID, Assignment Status, Power Status, Connection Status, Diagnostics Status, and Operational Status
- The Dynamic System Domain table lists all the domains and their details: Domain ID, MAC Address, Autoboot Policy, Secure Mode Policy, CPU Mode, Diagnostics Level, Domain Degradation Policy, and Operational Status

For Oracle SPARC M5-32 and M6-32 servers, the System level of the Hardware tab also includes:

- System Type
- Part Number
- System Identifier
- Management IP
- Management MAC Address
- Actual Power Consumption
- Status

- Data Source
- The Subsystem Status table displaying a summary including: name of Subsystem, Status, and Inventory.
- The Configured Dynamic System Domains table listing all the domains and their details: Domain ID, Domain Name, Priv MAC address, Auto Boot Policy, ILOM IP, Keyswitch State, and Operational Status
- The Unconfigured Dynamic System Domains table listing the same information as the table above except for Domain Name.
- Allocated Resources. No data is displayed on this table.
- Unallocated Resources. No data is displayed on this table.
- The Firmware table displaying for each item: Description, Type, and Version.

Oracle SPARC M5-32 and M6-32 servers include an ILOM 3.2 service processor. For more information about ILOM servers see ALOM and ILOM Servers Details.

Note: M5 and M6 servers are supported, but some features have additional limitations. For more information see the `Target Servers` section of the *Certified Systems Matrix* document in the Oracle Enterprise Manager Ops Center document library.

Component Navigation Pane of Hardware Tab

Use the component navigation pane in the Hardware tab to view information about each component of the system:

- CPU: Name, Architecture, Type, Manufacturer, Speed, Core Count, Thread Count, Serial Number, Part Number, Version, Status For Sensors: Name, Description, Type, and Value.
- Memory: Name, Type, Size in bytes, Serial number, Part number, Status For Sensors: Name, Description, Type, and Value.
- Board: Name, Serial number, Part number, Memory mirrored, Version, and Status.
- Power Supply: Name, Serial number, Part Number, Status For Sensors: Name, Description, Type, and Value.
- Board: Name, XSB Mode, Memory Mirrored, Serial Number, Part Number, Version, Status For Sensors: Name, Description, Type, and Value.
- IO Unit: Name, Serial Number, Part Number Version, Status For Sensors: Name, Description, Type, and Value.
- XSCF: Name, Host Name, Serial Number, Part Number, Version, and Status.
- Fan Tray: Name, Manufacturer, Part number, and Serial number.
- Fans: Name, Speed For Sensors: Name, Description, Type, and Value.

Oracle Enterprise Manager Ops Center monitors the voltage for the Board and IO Unit components and the speed for the Fan components. The Monitoring tab shows the actual value and the threshold values.

See Discovering a SPARC Enterprise M-Series Server in Chapter 2 for discovery requirements.

See ALOM and ILOM Servers Details for more information about the specific data available in the component navigation pane for Oracle SPARC M5-32 and M6-32 servers.

Switch Details

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches and Sun Datacenter InfiniBand switches. These switches reside in the system or blade system and provide the switch fabric. Cisco Catalyst switches are also supported.

For more information about Oracle Datacenter and Ethernet switches, see: http://www.oracle.com/technetwork/documentation/oracle-net-sec-hw-190016.html#legacysecapp.

Summary Tab

Oracle Enterprise Manager Ops Center reports hardware information on the Summary tab:

- Name
- Model
- Port count
- Serial number
- Management Interface IP
- MAC Address
- Fabric Manager: true or false
- Fabric Manager Address
- Power state
- Locator lights state
- Notification state
- Current Alert Status
- Firmware types and versions

Hardware Tab

At the System level, the Hardware tab includes:

- Model
- Server Name
- Serial Number
- State
- Power
- Firmware versions
- Sensors: temperature and voltage

You change the display to show information about each component of the switch:

- Network Adaptors: Name or each, MAC Address, IP Address, Description

- Power Supply: Name, Manufacturer, Part number, Serial Number For Sensors: Description, Type, Status
- Fan Sensors: Description, Type, Value, Status, Warning Threshold (Lower), Warning Threshold (Upper), Critical Threshold (Lower), Critical Threshold (Upper), Non-Recoverable Threshold (Lower), Non-Recoverable Threshold (Upper)

The Actions pane displays the set of available actions to manage a switch. It includes the **Launch Switch UI** for accessing and managing the switch directly from its Web UI. For the Cisco Catalyst switch, the Launch Switch UI action will be disabled if the HTTP server is not enabled in the switch.

Rack Details

The rack's Dashboard displays the following information:

- Name
- Rack ID
- Description
- Number of Slots
- Tags
- Support

The Details tab shows the configuration of each slot in the rack. For each component in the rack, this tab displays the position, name, description, type, model and health status. The Power Distribution Units are included and their current status is displayed.

The Firmware tab displays the name, description, and version of the firmware for each component and the slot for each component:

- Compute Nodes
- Switches
- Storage Appliances
- Power Distribution Units

The rack's Charts tab displays the following plots, as described in Charts Tab:

- Aggregate Power Usage
- Power Usage
- Average Fan Speed

The rack's Energy tab displays the information described in Energy Tab.

PDU Details

The Dashboard tab shows:

- Name and Description
- Model and serial number
- Management IP address and MAC address, if any

The Details tab reports the same information and adds whether SNMP and HTTP is enabled and the version of the firmware.

Oracle Solaris Cluster Details

The Dashboard tab shows:

- Name, description, and ID
- Number of possible quorum votes and the current quorum votes

The Network tab shows the public and private interconnects used by the cluster.

The Quorum tab shows the status of each member of the quorum and the number of quorum votes for each member.

Monitoring Power Utilization

Input power is the power pulled into a power supply from an external resource. The power consumption of a hardware asset is the sum of the input power consumed by each power supply of the asset. Output power is the amount of power provided from the power supply to the system components, measured at the power supply output. Input power is calculated from output power by applying an efficiency function to the output power from each power supply.

Calculating power compensation for the blades is difficult because the power supplies are shared. Each blade gives a report based on the power consumption of the local components, but this is not an accurate power consumption value for an individual blade.

To measure the input power, the interfaces must be exposed and the service processors must be able to retrieve and report data with one-minute accuracy. Servers that can report power usage have a Charts tab.

You can see current power usage and change the display of power graphs using the controls on the Energy tab and the Charts tab.

Energy Tab

The asset's Energy tab reports power consumption as the current value and for a period of time, as well as attributes of the fan and power supplies.

General Information

The following general information is displayed in the energy tab:

- Power consumption in watts.
- System load for an OS.
- Power policy.
- Utilization percentage for an Oracle VM Server for SPARC.
- Inlet and outlet temperature reporting the incoming and outcoming air temperature.
- Cost per kilowatt-hour in the selected currency.
- The currency units used to compute cost. The price per currency unit is set by the **Edit Energy Cost** action in the Administration section of the Navigation pane. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information.
- The total power cost in the selected currency. The period of time used to compute the cost is determined by the value selected in the Display list.

Energy Performance Charts

The data over time is represented in the following charts:

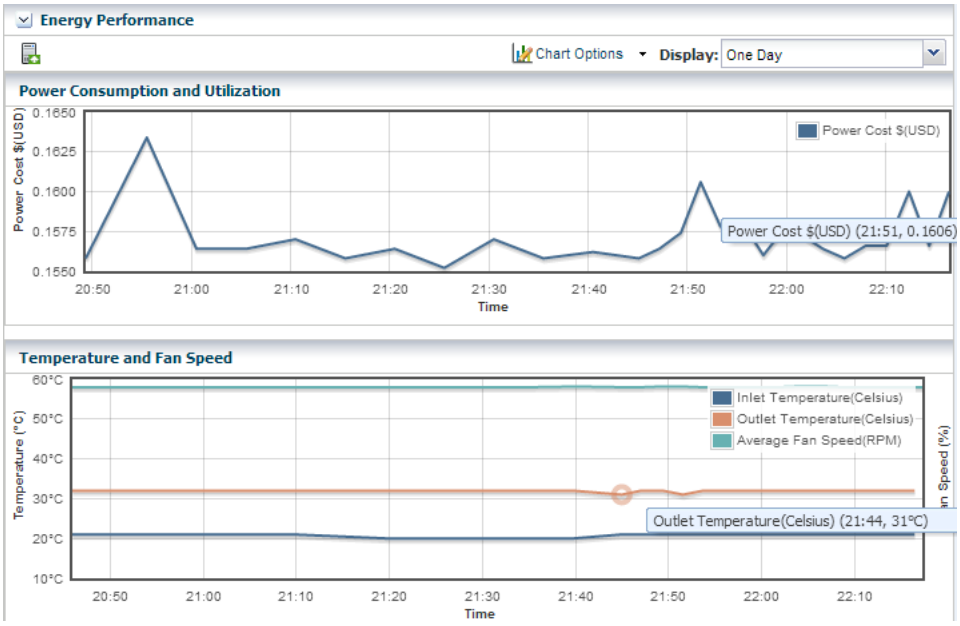

- Power Consumption and Utilization. By default, the chart shows the power consumed in the last day in watts. If the server is shut down, the chart shows any existing historical data.
- Temperature and Fan Speed. By default, the chart shows the incoming air temperature and the outgoing air temperature in Celsius, and the average fan speed in RPM. Click any point in the chart to see the data for that point in time.

> **Note:** Servers that use the ALOM and ILOM data model report fan speed in RPMs. Servers with SDM enabled ILOM –such as the M5-32 and M6-32 servers– report fan speed as a percentage of the maximum speed of the fan.

You can set the units for the power consumption and utilization chart's power axis using the Chart Options list. You can select watts or cost.

You can use the Display list to set the period of time that the charts display to one of the following values:

Table 11–3 Values of the Display List of the Energy Performance Charts

Value	Charts' Points Sample Rate
Live	Five minutes
One hour	Five minutes
One day	Five minutes
Five days	Five minutes
Three weeks	One hour
Six weeks	12 hours
Six months	One day

To make a graph with the minimum of two points, a hardware asset must have been managed for at least 10 minutes to view a one-hour graph and for at least two days to view the six-months graph.

The data for these time periods is stored separately. For example, if a server has been managed for two hours and you select the six weeks view, the graph cannot be displayed because only one point of data of that type has been stored; the second point has not yet occurred. If you then select the one day view, the graph can display 24 points of data (120 minutes at 5-minute intervals). However, the graph displays these points over a 24-hour period and not over the actual two-hour period. For the most accurate representation of the data, choose a time period that is less than or equal to the time that the hardware asset has been managed.

You can export the data for either the current view or all available data to a file in either CSV or XML format. Use the Export Chart Data toolbar icon to choose options for exporting the data.

If the graph is blank, one of the following conditions has occurred:

- The server does not have the appropriate ILOM version.
- The server has not been discovered through the ILOM driver.
- The server is unreachable.

Power Supply and Fan Information Tables

The Power Supply table lists: the power supply number, manufacturer, and part and serial numbers.

The Fan Information table lists: the fan number, and fan speed as a percentage or in RPM.

Charts Tab

The Chart tab provides more ways to display the power utilization data. You can change the graphed data to a bar chart or an area chart. You can also export the data for either the current view or all available data to a file in either CSV or XML format. Use the Export Chart Data button to choose options for exporting the data.

For groups and virtual pools, the following options are available:

- Select Order: The five highest or five lowest historical power utilization.
- Select Resource: Select the Power or Aggregate Power option for a homogeneous or heterogeneous group of servers.
  - The Power option displays power utilization for the five highest or lowest power consumers in the group or virtual pool.
  - The Aggregate Power option displays the power utilization, using the sum of all members that report power consumption. The number of systems in the aggregate is included. For heterogeneous group, the Chart tab includes a table of all systems in the group and their various power attributes for the selected time period. From this table, you can power off and power on selected servers to conserve power.

Maintaining Hardware Assets

- Update management credentials, see Updating Management Credentials in Chapter 2

- Set power policy
- Power systems on and off
- Place in maintenance mode
- Reset a server
- Get access to the serial console
- Enable and disable ports
- Use locator lights to identify a specific asset
- Check firmware compliance and update firmware. See Firmware Provisioning

Setting and Changing the Power Policy

The power policy allows you to set an asset in one of three different modes:

- Performance: Unused components are put into a slower speed or sleep state and power savings features with insignificant performance impact are enabled.
- Elastic: Components are brought in to or out of a slower speed or a sleep state to match the system's utilization of those components.
- Disabled: All components run at full speed or capacity. This option is available for some models and ILOM version.

View an Asset's Power Policy

1. Expand the Assets in the Navigation pane.
2. Select a hardware asset.
3. Click the **Energy** tab in the center pane.

Changing an Asset's Power Policy

1. Expand the Assets in the Navigation pane.
2. Select a hardware asset.
3. Click **Set Power Policy** in the center pane to display the window shown in Figure 11–7.

Figure 11–7 Set Power Policy

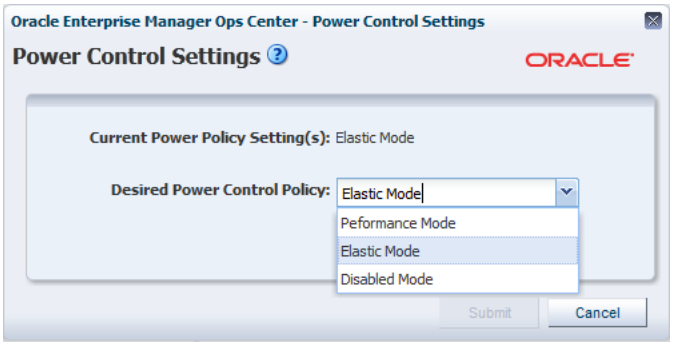

4. The current power policy is displayed. Choose the alternative policy.
5. Click **Submit** to create a job that sets the power policy.

Replacing a Failed Power Distribution Unit in a Rack

1. Expand Assets in the Navigation pane and then expand Racks.
2. Navigate to the rack type and expand to show its components. Expand Power Distribution Units.
3. Select the failed power distribution unit.
4. Click **Place in Maintenance Mode** in the Action pane.
5. Go to the rack's location and remove the failed PDU.
6. Install the new PDU and connect it to the network, according to the procedures in the Power Distribution Units User's Guide at http://docs.oracle.com/cd/E19844-01/index.html.
7. Verify that the new PDU has the same IP address as the failed PDU.
8. Use the PDU's web interface to configure the administrator user account, according to the procedure in the Power Distribution Units User's Guide.
9. Return to the Oracle Enterprise Manager Ops Center user interface and navigate to the same rack and the new PDU. Select the PDU.
10. Click **Remove From Maintenance Mode** in the Action pane.
11. To update the credentials for the administrator user, click **Update Management Credentials** in the Action pane. The wizard opens.
12. At the Management Type step, select HTTP credentials and then select the Create a new set of credentials option. Click **Next.**
13. Enter new credentials for the administrator user.
14. Click **Apply** to update credentials.
15. To create new SNMP credentials, click **Update Management Credentials** in the Actions pane again.
16. Select SNMP credentials and Create a new set of credentials.

    Specify a community string that is different than the previous community string.
17. Press Apply to update credentials.

Installing and Upgrading Oracle Solaris Cluster

You use cluster profiles in a deployment plan to perform the following operations:

- Install Oracle Solaris Cluster software
- Upgrade Oracle Solaris Cluster software

See Appendix A for instructions on obtaining the current profiles and supporting scripts.

Firmware Provisioning

The Oracle Enterprise Manager Ops Center provisioning feature installs firmware on the managed hardware assets. You initiate the installations from the UI, rather than from the asset itself.

Oracle Enterprise Manager Ops Center provides default profiles for configuring firmware for servers and for disk storage. An alternate procedure is to create a Firmware Report. You then use the report results to update the firmware.

The benefit of using a profile to install firmware is that the firmware is installed consistently, no matter how many assets you provision. The benefit of using the Firmware Compliance Report is to identify the firmware on a specific asset or set of assets.

Firmware Profiles

The general procedure for provisioning firmware has the following steps:

1. Import a file with the firmware and the associated metadata into a software library, according to the procedure in *Keeping Your Firmware Up-to-Date* in the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.
2. Create a firmware profile, based on one or more firmware images, according to the procedures in *Keeping Your Firmware Up-to-Date* in the Operate How To Library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.
3. Shut down the server gracefully. Most firmware requires that the server is not running when the firmware is updated. Most firmware images include a power-off command for a running server, which causes a hard shutdown of the server.
4. Apply the firmware profile.

Firmware Compliance Reports

The Firmware Report feature compares the firmware images specified in a firmware profile to the firmware images installed on one or more hardware assets.

The report shows whether a target asset is compliant, not compliant, or not applicable:

- A compliant asset has the firmware images specified in the profile.
- A non-compliant asset does not have the same firmware images as specified in the profile. Update the firmware by either clicking the **Make Targets Compliant** button in the Interactive report or using the procedure in Updating Firmware.
- A non-applicable asset indicates that a firmware image in the profile does not match the model of service processor in the asset. This condition can occur when either the profile does not recognize the model of the service processor or the profile includes firmware images that are not designed for the service processor.
    1. Compare the model of the service processor displayed in the asset's Summary tab with the model of the service processor included in the profile. If they are different, add the name in the profile to the asset's data.
    2. When the firmware profile was created, only images that matched the service processor could be included. However, if the service processor did not report all the firmware types it supported, an image that did not match the service processor could have been included in the profile. To update the Oracle Enterprise Manager Ops Center software with all the service processor's supported firmware types, use the **Refresh** action. When the job is completed, view the service processor's Summary tab to see all firmware types.
    3. Create a new firmware compliance report.

See Creating a Firmware Report for the procedure to create the report.

Updating Firmware

To update the firmware on one or many assets, you use a deployment plan to apply a firmware profile for the type of asset. For a server, the profile updates the firmware on a service processor, and restarts the service processor and operating system. For storage components, profiles update firmware on a RAID controller, an expander, or disk.

To see the deployment plans that update firmware, expand the Deployment Plans section of the Navigation pane and then click **Firmware**. A list of existing plans and profiles is displayed.

To update the firmware of one asset, an alternative to a deployment plan is to use the **Update Firmware** action. Select the asset from the Asset section of the Navigation pane and then click **Update Firmware** in the Actions pane.

Before You Begin

The software library must contain the images that provision the firmware. Perform the Uploading a Firmware Image procedure in the *Keeping Your Firmware Up-to-Date* document in the *Operate How To library* at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.

If you are updating the firmware on a server, shut down the server before you update the firmware. A firmware update to a server's service processor usually requires that the server is not running. If you start to update the firmware on a running server's service processor, the procedure performs a hard shutdown of the server.

Requirements for ALOM Service Processors and M-Series Servers

The firmware provisioning process for M-Series servers and servers that have ALOM service processors relies on a temporary account that performs an FTP operation.

> **Note:** Advanced Lights Out Management (ALOM) is a Sun Microsystems standard for servers such as: SunFire V125/V210/V215/V240/V245/V250/V440/T1000/T2000, Sun Netra 210/240/440, and SunBlade T6300.

If your site does not allow a temporary account, use the following procedure to prepare for the provisioning operation:

1. On the Enterprise Controller, open the `/var/opt/sun/xvm/hal.properties` in an editor.
2. Add the following properties to the file:

   ```
   ftp.user.name=username
   ftp.user.password=password
   ```
3. Restrict access to the file to root user:

   ```
   chmod 600 /var/opt/sun/xvm/hal.properties
   ```
4. On the Proxy Controller that provisions the firmware, enable the `ftp` service on Oracle Solaris or the `vsftpd` service for Oracle Linux systems.

You can now apply the firmware provisioning deployment plan. The FTP operation retrieves the credentials from the file.

If a network failure occurs while updating the firmware, repeat the firmware update procedure. If you do not repeat the procedure, the firmware inventory list might be incomplete.

Option for Deferring the Stop and Restart of the Operating System and Server

The procedure is available for Oracle Solaris 10 Update 10 operating systems running on servers with the ILOM x86 (3.0 and higher) service processor.

When you update the firmware of a service processor, the procedure stops the operating system and the server before the update and restarts them after the update so that the new BIOS takes effect. If you prefer to stop and restart at a convenient time, keeping the current BIOS in effect, use the following procedure to change the action of both Oracle Enterprise Manager Ops Center and the firmware's metadata:

1. On the Enterprise Controller, open the `/var/opt/sun/xvm/hal.properties` in an editor.
2. Add the following property to the file:

   ```
   ilom.fwp.skipAutoReboot=true
   ```
3. On the Proxy Controller that provisions the firmware, enable the `ftp` service on Oracle Solaris or the `vsftpd` service for Oracle Linux systems.
4. At a later time, reboot the servers.

Launching LOM and XSCF Browser User Interfaces

When you select a server on the Assets pane, the **Launch LOM Controller** link is displayed on the Actions pane. This functionality launches the Browser User Interface (BUI) for servers with a Lights Out Management (LOM) port.

The **Launch SP Controller** link is only available to M-Series servers on the Actions pane for launching the specific BUI for the XSCF controller. The BUI is disabled by default on M-Series servers for security reasons and must be manually enabled before attempting to use this functionality.

The BUI for XSCF runs on the HTTPS protocol and can be enabled using the following command with a user with `platadm` privileges:

```
XSCF> sethttps -c enable
```

Related Resources for Hardware Management

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- For end-to-end examples see the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm and the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.
- For more information about Oracle Datacenter and Ethernet switches, see: http://www.oracle.com/technetwork/documentation/oracle-net-sec-hw-190016.html#legacysecapp.
- See the Power Distribution Units User's Guide at http://docs.oracle.com/cd/E19844-01/index.html.
- For information about Oracle SPARC servers, including SPARC T5, SPARC M5-32, and SPARC M6-32 servers, see SPARC Systems at http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html.

- See Systems Management and Diagnostics at http://www.oracle.com/technetwork/documentation/sys-mgmt-networking-190072.html for information about ILOM configurations.

12

Operating System Management

This chapter provides an overview of the operating system (OS) management features that are available in Oracle Enterprise Manager Ops Center.

The following information is included:

- Introduction to Operating System Management
- Roles for Operating System Management
- Actions for Operating System Management
- Location of Operating System Management Information in the User Interface
- Operating System Profiles
- Using Agent Management for Operating Systems
- Monitoring Operating Systems
- Using Analytics
- Overview of Oracle Solaris Boot Environments
- Overview of Oracle Solaris 11 Boot Environments
- Overview of Oracle Solaris 10 Boot Environments
- Related Resources for Operating System Management

See Chapter 13, "Operating System Provisioning" for how to install, or provision operating systems. See Chapter 14, "Operating System Updates" for information about patching and updating your operating systems.

Introduction to Operating System Management

Oracle Enterprise Manager Ops Center provides comprehensive lifecycle management for Oracle Solaris and Linux operating systems in your datacenter. The software also enables you to patch, or update, Microsoft Windows operating systems.

The discovery feature makes adding operating systems and other assets quick and easy. After the operating systems are added, they are considered managed and you can begin using the monitoring, analytics, OS provisioning, and update features to gather information and perform tasks.

Your managed operating systems are visible in the All Assets section of the user interface. The operating system appears under the service processor and hardware, as shown in Figure 12–1.

Figure 12–1 Managed System and Operating System

The operating system also appears in the appropriate platform-specific group in the Operating Systems view, as shown in Figure 12–2.

Figure 12–2 Operating Systems View

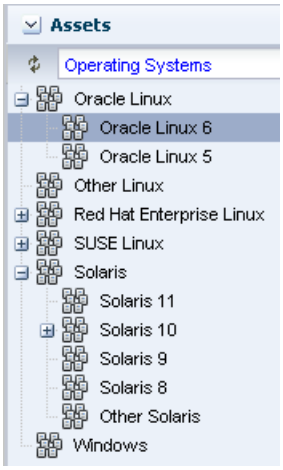

You can create user-defined groups and subgroups to refine your administration tasks. For example, you might want to create groups for Critical Systems, Training, Region 1 and Region 2. Groups are useful when you want to organize the systems to apply different monitoring standards, implement update requirements, or job scheduling times. You can create rules for your groups to automatically add existing and newly managed operating systems to the correct group or subgroup. See Chapter 2, "Asset Management" for details on using the discovery feature to add assets and for more information on creating user-defined groups.

Two types of OS management are available: agent managed and agentlessly managed. In some cases, the features and actions that you can perform on an operating system are determined by the management type. See Using Agent Management for Operating Systems for more information about the features available for agent and agentlessly managed operating systems.

The following features are available for operating systems:

- Monitoring: A series of monitoring rules and parameters monitor your managed assets. Alerts and incidents are raised for components that are not performing as expected.
- Performance: Analytics provides you with a detailed view into OS performance.
  - System resource graphs, processes information, and a view of the top consumers
  - Resource usage of virtualized OS instances
- Provisioning: Install Oracle Solaris or Linux operating systems onto your systems, making it easy to install one system or many servers simultaneously.

Note: The Enterprise Controller must be installed on an Oracle Solaris operating system in order to perform the following tasks:

- Provision Oracle Solaris 10 using JET customization.
- Provision Oracle Solaris 11 (requires that both the Enterprise Controller and Proxy Controller are running on an Oracle Solaris 11 operating system.
- Provision Oracle VM Server for SPARC (control domain).

- Manage Oracle Solaris Boot Environments: Create and manage Oracle Solaris boot environments in a repeatable and consistent manner from a single user console.
- OS Updates: Apply update packages to keep your operating systems up-to-date. See Chapter 14, "Operating System Updates" for information.
- Reports and Snapshots: A variety of reports are available for your operating systems. See Chapter 10, "Reports" for OS reports. See Chapter 14, "Operating System Updates" for information on how to use the System Catalogs to maintain snapshots of your operating system.

The monitoring feature provides extensive monitoring capabilities that are enabled as soon as you begin managing an operating system. A series of three escalating status levels notifies you when something is not operating as expected. The first level is informational, then warning, and finally a critical status.A set of default monitoring attributes and alert triggers are included with the software. You can tune the monitoring thresholds and triggers to define what you want to generate an alert and when. You can create custom monitoring rule sets and alert parameters and apply the customized monitoring rules to a specific operating system, a group of homogeneous operating systems, or a group that you define, such as critical systems or regional systems.

The Analytics feature provides extensive information about a specific operating system in one location so that you can maximize performance and utilization. The Analytics information includes process details, defined monitoring thresholds for operating systems, metrics and historical information on the top consumers, and a extensive list of metrics data. The Summary contains details on the top five consumers for CPU, memory, network, and I/O utilization. Use the graphical representation to quickly view utilization trends and high resource consumers. You can drill down to get detailed utilization and process information, and kill a process that is consuming too many resources.

Roles for Operating System Management

Table 12–1 lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 12–1 OS Management Roles and Permissions

Task	Role
Reboot an OS	Asset Admin

Table 12–1 (Cont.) OS Management Roles and Permissions

Task	Role
Charts and Utilization	Asset Admin
	Cloud Admin
Analytics	Read
	Asset Admin
Kill action in Analytics	Operating System Management
Update Management Credentials	Security Admin
Any Actions related to changing credentials	Security Admin
Import image	Storage Admin
Upload image	Storage Admin
Unconfigure, SCCM Configuration	Oracle Enterprise Manager Ops Center Admin
Reboot, upgrade Agent Controller	Asset Admin
Edit Tags	Asset Admin
Edit Attributes	Asset Admin
View Boot Environment	Read
	Asset Admin
	Update Admin
Create Boot Environment	Asset Admin
	Update Admin
Update an Alternate Boot Environment	Update Admin
Activate and Reboot a Boot Environment	Asset Admin
	Update Admin
Synchronize Boot Environments	Asset Admin
	Update Admin
Delete an Alternate Boot Environment	Asset Admin
	Update Admin
Monitor Boot Environment Attributes	Asset Admin

Actions for Operating System Management

You can manage an operating system in one of two modes: agent managed or agentless managed. The management mode determines the features that are enabled for your operating system.

Agent managed is the more robust management mode because the Agent Controller enables a greater level of communication with the Proxy Controller and Enterprise Controller than the agentless managed operating systems. You can use the features and perform the actions described in this chapter with an agentless managed operating system, but OS update functionality requires an agent managed operating system. You can manage your operating systems by installing an Agent Controller on the OS or by using SSH to perform tasks. See Using Agent Management for Operating Systems for more details on managing operating systems with or without an Agent Controller.

See Chapter 2, "Asset Management" for details on using the discovery feature to add assets and for more information on creating user-defined groups.

After you manage your assets, you can perform the following actions:

- Monitor your physical and virtual operating systems.
- View OS utilization for Oracle Solaris and Linux operating systems.
- Manage Oracle Solaris boot environments.
- Update or patch operating systems. See Chapter 14, "Operating System Updates".
- Provision or upgrade Oracle Solaris and Linux operating systems. See Chapter 13, "Operating System Provisioning".

> **Note:** To perform actions, such as OS provisioning, updating, and managing boot environments, on an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must be installed on an Oracle Solaris 11 operating system.

Location of Operating System Management Information in the User Interface

When you manage a physical or virtual operating system, it appears under the associated server in the All Assets view, as shown in Figure 12–3 and it appears in a special pre-defined operating system group, as shown in Figure 12–4.

Figure 12–3 All Assets View

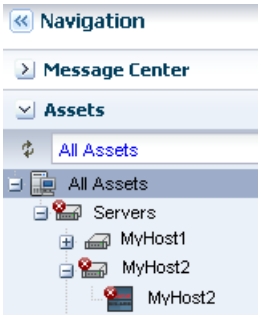

Operating systems are automatically added to a homogenous group of operating systems. The group contains directories for each release. In Figure 12–4, the list of Oracle Solaris 10 operating systems includes physical operating systems, virtual hosts, and zones. You can add user-defined groups and create rules to automatically add newly discovered assets to that group, or you can manually add them at any time. See Using Groups for details.

Figure 12–4 Operating Systems View

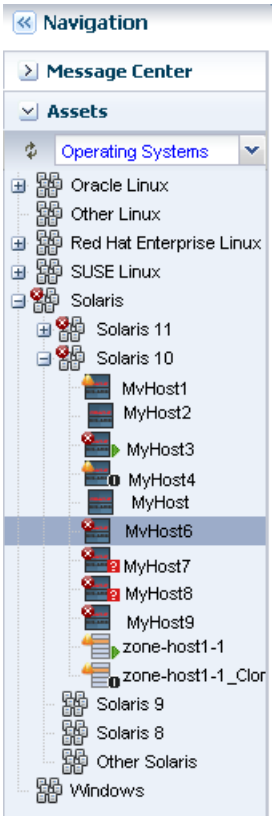

To view information about a specific OS, select the OS from the Assets pane. OS details appear in the tabs across the center pane.

Table 12–2 Location of Operating System Information in the UI

To Display	Select
Managed operating systems	Expand the Assets pane. Each operating system appears in the Assets Navigation tree under the system on which it is installed. To only view operating systems, click the drop-down next to All Assets and select Operating Systems. The systems are grouped by platform.
Operating system details for a specific operating system	Select an operating system in the Assets pane, then click the Summary tab.
Operating system details for a group of operating systems	Select an operating system in the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Summary tab.
Unresolved incidents and alerts for a specific operating system	Select an operating system in the Assets pane, then click the Incidents tab. The details are in the Unresolved Incidents and Alert subtabs.
Unresolved incidents and alerts for a group of operating systems	Expand the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Incidents tab. The details are in the Unresolved Incidents and Alert subtabs.
Monitoring Rules for a specific operating system	Select an operating system in the Assets pane, then click the Monitoring tab.

Table 12–2 (Cont.) Location of Operating System Information in the UI

To Display	Select
Monitoring Rules for a group of operating systems	Select an operating system in the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Monitoring tab.
Analytics	Select an operating system in the Assets pane, then click the Analytics tab.
Boot environments	Select an operating system in the Assets pane, then click the Boot Environments tab.

Viewing Operating Systems

You can view all managed assets in the Assets section of the Navigation pane, or you can filter to display a specific group of assets. Figure 12–5 shows the asset filter menu where you can choose the Operating System group or user-defined groups.

Figure 12–5 Asset Filter Menu

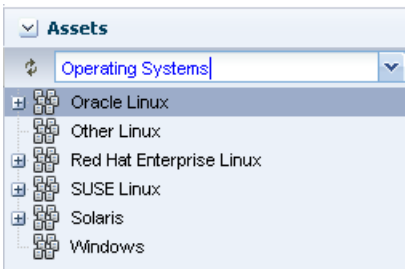

See Using Groups for information about creating user defined groups.

Displaying Operating System Details

Details about a specific operating system appear in a number of tabs in the center pane of the UI. The tabs and contents might vary based on the type of operating system selected and whether the operating system is agent or agentlessly managed. When a tab is not applicable, it will not appear in the UI. For example, the Libraries tab does not appear when there are no libraries associated with the operating system.

To view the details in the UI, perform the following steps:

1. Expand **Assets** in the Navigation pane.
   - To display All Assets, expand **Servers**.
   - To display the Operating System groups, select **Operating Systems** from the menu.
   - To display user-defined groups, select **All User Defined Groups** from the menu, then select the group.
2. Select an operating system in the Navigation pane.

   The following tabs appear in the center pane:

   - Dashboard: Displays the summary of the selected operating system, including the name, server name, whether the operating system is agent managed, the number of unassigned incidents, the current alert status, and the OS release. The dashboard also includes the operating system's membership graph, the status of incidents, and compliance reports, if any.

- Summary: Displays the name of the selected Oracle Solaris operating system, description, server name, operating system version, number of CPU threads, active boot environment, the state, the length of running time, and zone patching information. Tables display total CPU utilization and total ZPool Utilization for Oracle Solaris operating systems.
- Libraries: Displays library details when a library is associated with an operating system. See Chapter 5, "Software Libraries" for more information.
- Storage: Displays logical unit (LUN) details and iSCSI targets, when applicable. LUN details include the initiator, MPxIO details, the GUID, type, size, status, vendor, and product. The iSCSI details include the initiator, iSCSI target address, address type, port, and IP address. You can discover iSCSI addresses from this tab. See Block Storage for more information.
- Analytics: Displays utilization and metrics for the selected operating system and zones, when applicable. See Using Analytics for more information.
- Connectivity: Displays Linux OS network interface of the system, including network connectivity and aggregated links. See Network Connectivity for more information.
- Networks: Displays Oracle Solaris network connectivity, IPMP groups, and aggregated links. For an Oracle Solaris 11 OS, the tab also provides bandwidth management. See Network Connectivity and Bandwidth Management.
- Incidents: Displays all incidents reported from the selected operating system. See Chapter 9, "Incidents" for more information.
- Monitoring: Displays the alert monitoring rules and service dependencies of the selected operating system. Chapter 4, "Monitoring Rules and Policies" for more information.
- Terminal: Enables you to establish an SSH connection to the terminal window.
- Boot Environments: Displays Oracle Solaris boot environment details, including all available boot environments, the active and enabled boot environment, the size, and the creation or synchronization date. For a selected boot environment, you can view snapshot details, file system details, and any associated zone boot environments. This tab is only available for Oracle Solaris operating systems. See Overview of Oracle Solaris Boot Environments for more information.
- Jobs: Displays current and historical job information for that operating system. See Chapter 3, "Jobs" for more information.
- Service Request: Displays the service request for the selected operating system. See Using Oracle Services and Service Requests for more information.
- Configuration: Displays access points, or resources, that are associated with the operating system. See Using Access Points for more information.

Operating System Profiles

The following categories of OS profiles are available:

- OS Provisioning
- Monitoring Profile
- Oracle Solaris Zone
- Boot Environments

Using Agent Management for Operating Systems

An agent managed operating system has an Agent Controller or a specialized virtualization, or VC, Agent Controller installed to gather information for the Enterprise Controller. The VC Agent Controller is for virtualization technology, such as zones and logical domains.

When you install the Agent Controller on the operating system, the following actions occur:

- The software registers the Agent Controller with the Enterprise Controller. It takes at least five (5) minutes for the software to register the Agent Controller. After registered, you can update the operating system.
- The software sends you a notification when it has enabled the update function for the operating system.
- The Agent Controller checks the inventory of patches and packages and creates the System Catalog. The catalog lists the patches and packages, and the versions that are currently installed on the operating system.

Virtualization Agent Controllers

In addition to the default Agent Controllers, Oracle Enterprise Manager Ops Center uses specialized virtualization Agent Controllers called VC Agent Controllers for Oracle VM Server and Oracle Solaris Zone assets.

You can install the agent during discovery, or at any time after discovery. You have the following agent management options:

- Oracle VM Server for SPARC Virtualization Controller Agent: Manages the logical domains that are running on the Control Domain. The Oracle VM Server, Control Domain and operating system are reflected in the UI. Using this agent enables full monitoring and management actions for the Oracle VM Server system.
- Zones Virtualization Controller Agent: Manages the zones that are running on the logical domains. The global zone is reflected in the UI. Using this agent enables full zone monitoring and management actions.
- Agentlessly: Limited management functionality is available with this method. Information is gathered by using SSH connection between the logical domains and the Proxy Controller.

For robust management, use the OVM Server for SPARC Virtualization Controller Agent to manage the domains. The agent runs on the Control Domain and monitors the configuration and reflects any changes on the configuration in its copy of the metadata.

Functionality With and Without Agent Controllers

The Agent Controller provides the most robust management features. However, you can manage your assets without using an Agent Controller. To gather information on an agentlessly managed operating system, the Proxy Controller uses SSH to perform certain tasks and periodically check on the operating system.

Some features are not available when the operating system is managed agentlessly. Table 12–3 shows the information that are available for each management type.

Table 12–3 Information Available for Agent Managed and Agentlessly Managed Assets

Tab or Feature	Agent Managed	Agentlessly Managed
Dashboard	Yes	Yes
Summary	Yes	Yes
Libraries	Yes	No
Storage	Yes	No
Utilization	Yes	Yes
Analytics	Yes	Limited
Virtualization Analytics for Oracle VM Server	Yes, if the guest is agent-managed	No
Virtualization Analytics for Oracle Solaris 10 Zones	Yes, if the global zone is agent-managed or if the non-global zone is agent-managed.	No
Networks	Yes	No
Incidents	Yes	Yes
Monitoring	Yes	Yes
Charts	Yes	Yes
Reports	Yes	No See Chapter 10, "Reports"
System Catalogs	Yes	Oracle Solaris 11: Yes Oracle Solaris 8-10, Linux, Windows: No
Terminal	Yes	No
Jobs	Yes	Yes
Configuration	Yes	Yes
OS update	Yes	Oracle Solaris 11, Windows: Yes Oracle Solaris 8-10, Linux: No
OS provisioning	Yes	Yes
Zone management: boot, shutdown, halt, delete, edit zone config, get zone console	Yes	Yes
Zone management: create, clone, migrate, add a filesystem to a zone, remove a filesystem from a zone, attach and detach networks, add storage to a zone	Yes	No

Switching Between Agent Controllers or Agent and Agentless

The current management mode of an operating system and the type of Agent Controller appears on the Dashboard for the operating system, as shown in Figure 12–6.

Figure 12–6 Agent Managed

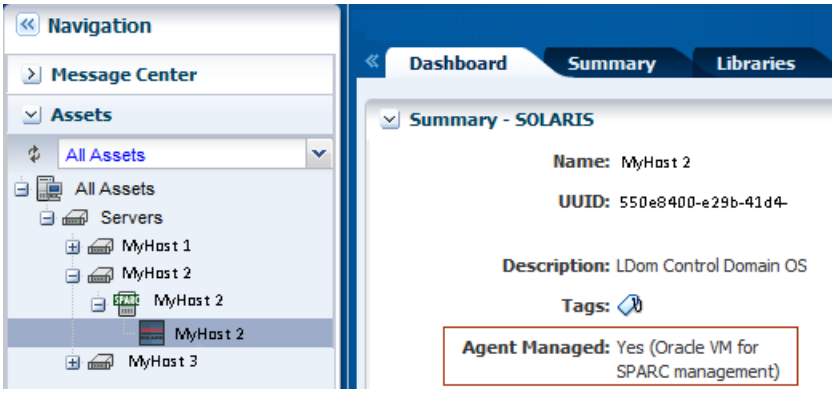

You can use the following methods to change the agent management mode:

- Unmanage the asset, then rediscover the operating system using a profile with the alternative mode.
- Use the Switch Management Access feature.

The Switch Management feature enables you to move back and forth from agentlessly managed to agent managed or to switch the type of agent.

When the operating system is agent managed and you use this action, the software removes the Agent Controller and changes the operating system to agentless. Select or create new credentials for the Proxy Controller to use to obtain information from the asset.

To switch between different types of Agent Controllers, such as changing from a default Agent Controller and a virtualization Agent Controller, you must use Management Access to unmanage and then manage again. When you manage the asset, you are prompted to choose the type of Agent Controller when it is not apparent to Oracle Enterprise Manager Ops Center which agent you want to install.

To Use Switch Management Access

1. Expand **Assets** in the Navigation pane.
2. Expand the Assets tree, then select the operating system.

   The current management status appears in the Dashboard tab.
3. Click **Switch Management Access** in the Actions pane.

Figure 12–7 Switch Management Access

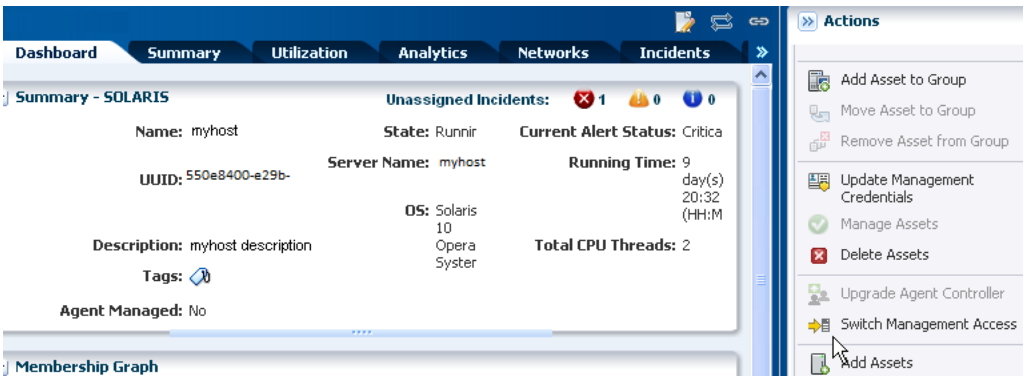

4. Add or select the credentials for the system, then click **Finish**.
   - (Optional) To create a new set of credentials, click **New** and complete the Create Credentials Wizard, then click **OK**.
   - (Optional) To select from a list of existing credentials, click **Select**, highlight the credentials from the list of available credentials, then click **OK**.

Figure 12–8 Switch Management Access Credentials

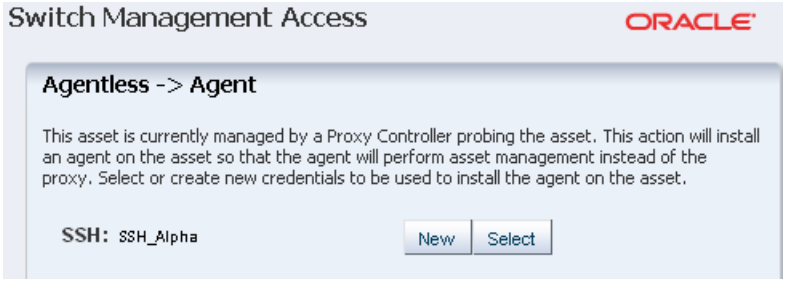

Monitoring Operating Systems

The software monitors the status of your operating systems right after asset discovery. Oracle Enterprise Manager Ops Center uses the default rules and thresholds. Monitoring rules state the values and boundaries for an asset's activity. A monitoring policy is a set of rules. You can change the rules or thresholds to adjust the type and level of monitoring you want. Analytics provides details on operating system activity and utilization.

A monitoring policy defines alert configurations to be performed on one or more managed resources. A policy is for a specific type of resource, such as operating systems. A more specific policy might apply to all Oracle Solaris operating systems. Each monitoring policy contains several alert monitors for a specific type of resource. Alert monitors watch the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

Applying a monitoring policy to all the assets enforces consistency. Each monitoring policy contain rules for threshold levels. Default policies for monitoring hardware, operating systems, and Oracle Solaris Clusters are included in the software. You can use the default policies, but you cannot edit them. To edit or add monitoring rules to a monitoring policy, you must make a copy. Modifying a monitoring rule for a specific asset creates a custom set of monitoring rules for the asset.

See Chapter 4, "Monitoring Rules and Policies" for information about how to change the threshold limits and how to change the deactivate or activate the auto delete policy. See Chapter 9, "Incidents" for details about how incidents are generated, severity badges, how to assign and close an incident.

Using Analytics

The Analytics feature provides a view into operating system and zone performance and status. The charts, reports, and utilization data provide details of an individual eligible operating system or zone. You can use the information to analyze the behavior of an operating system to aid in peak performance and to diagnose and correct incidents.

Monitoring uses the Agent Controller to gather information. When an operating system is not an agent-managed system, or it is Microsoft Windows, the software uses

an SSH connection from a Proxy Controller to perform remote monitoring. Remote monitoring over an SSH connection limits the available metric information. The Summary view, Process view, Historical view and Virtualization Analytics are not available for these operating systems.

The Analytics view provides information for the following agent-managed platforms:

- Linux
- Oracle Solaris 11 and 10 OS
- Oracle Solaris 10 non-global zone, when the global zone or non-global zone is agent-managed
- Oracle Solaris 11 non-global zone, when the global zone is agent-managed

When your operating system is agentlessly-managed, information is not available for zones and less information is available for Linux and Oracle Solaris 10 and 11 operating systems. Table 12–4 shows a list of features and whether the feature is supported on an agent-managed or agentlessly-managed asset.

Table 12–4 Supported Analytics Information

Feature	Supported on Agent Managed OS	Supported on an Agentlessly Managed OS
Summary	Yes	Yes
Virtualization Analytics for Oracle VM Server	Yes, for agent-managed guests	No
Virtualization Analytics for Oracle Solaris 10 Zones	Yes, for an agent-managed global zone or non-global zone	No
Services	Yes	Yes
Processes	Yes, for an agent-managed guest, global zone, or non-global zone	No
Threshold	Yes	Yes
History	Yes, for an agent-managed guest, global zone, or non-global zone	No
Metrics	Yes	Yes

The current management mode appears in the Dashboard Summary page for the operating system. See Using Agent Management for Operating Systems for details and for the steps to switch the management mode.

Displaying Analytics Information

For each operating system, the information appears in the **Analytics** tab in the center pane.

The following diagnostic pages are available for an agent-managed operating system:

- OS Analytics view: Displays analytics details for an operating system.
  - System resource graphs
  - Top-consumers views
  - Processes information

  - View historical data and set thresholds
- Virtualization Analytics view: Displays analytics details for a zone.
  - Virtualized OS instance
  - Breakdown of resource usage of the physical server
  - Running OS instance
  - View historical data and set thresholds

To Display the Analytics View

1. Select an operating system, zone, or guest in the Navigation pane.
2. Click the **Analytics** tab.

To Display the Both the OS Analytics and Zone Analytics Views

1. Select a global zone that has zones in the Navigation pane.
2. Click the **Analytics** tab to see the OS Analytics.
3. Click the **Zones Analytics** tab to view the Summary, Zones, History, and Resource Pool Utilization for all zones of the global zone.
4. See Displaying Virtualization Analytics for information about analytics for virtualized platforms.

The Analytics view contains current and historical information about an operating system's use of resources, including CPU, Network, disk I/O, memory, alert history, and total thread and process counts. Information presented in this section assumes that the operating system is agent managed. See Table 12–4 for the list of supported features for an agentlessly managed OS.

The information appears in a series of charts in the following subtabs:

- Summary: View a high-level overview of the top consumers, by process, for the CPU, memory, network, and disk I/O resources. See Displaying the Analytics Summary for details.
- Processes: View process-specific details. See Displaying the Processes View for details.
- Services: View and monitor Oracle Solaris 10 and 11 SMF services. See Displaying the Services View for details.
- Thresholds: View and edit the threshold limits for a selected monitored attribute. See Displaying Thresholds for details.
- History: View a history for the top consumers. See Displaying Historical Data for details.
- Metrics: View specific details an operating system element, such as the percentage of memory used. See Displaying Metrics for details.
- Charts: You can create a variety of utilization charts, define the coordinates, and export the chart data to CSV or XML output. See Displaying and Creating Charts for details.

You can view analytics information for all supported and managed operating systems, both physical and virtual. The information presented varies, depending on the data available for the OS. For example, when the System Tap is not installed on a Linux operating system, information about the top network and I/O consumers might not be

available. Also, the top consumers for I/O data might not be available with some versions of the Linux kernel.

When an agent-managed operating system is on a virtualization platform, you can view the information in the Analytics view of the virtualization platform. See Displaying Virtualization Analytics for information about analytics for virtualized platforms.

Displaying the Analytics Summary

The Summary view provides details for an individual operating system. The Summary provides a graphical representation and a list of the top five (5) resource usage consumers, by process, for the CPU, memory, network, and disk I/O resources. Click the **View** icon for a row to display process details.

The following details appear in the Summary view:

- CPU Utilization: A graph of the percentage of CPU Utilization over time and a detailed list of the top five (5) CPU processes. The list includes the process identification number (PID), name, and CPU Usage %.
- Memory Utilization: A graph of the top 5 memory processes, including the PID, the process name, and the percentage of memory used.
- Network Utilization: A graph of the top 5 network processes, including the PID, the process name, and the network usage in Kilobytes (KB) per second.
- I/O Utilization: A graph of the top 5 I/O processes, including the PID, the process name, and the I/O usage per second.

Oracle Enterprise Manager Ops Center collects information every five minutes on every managed asset and displays the last hour of data on the Summary tab. Each list has an icon before the PID column for each process. Click an icon to view details of the process. The type of resource determines the available details.

- Process Details: Includes the contract ID, Service FMRI, creator, elapsed time, project name, project ID, and a process tree.
- Thread Information: Includes the thread's light-weight process (LWP) ID, the number of threads (NLWP), the user, the priority, the state (such as sleeping), the percentage of CPU used, the CPU time, the percentage of memory used, how much memory the process has marked for allocation (VSZ memory), the processors, and the command.
- Handles: Includes the port details, such as the family, local address, local port, the protocol, remote address, remote port, device, and node.
- Process Environment: Includes details about the environment, such as the arguments, data model, and flags.
- Memory Information: Includes physical, virtual, and anonymous memory and dirty page details. Details include the virtual address and the number of KB in the virtual mapping size, the resident physical memory, the anonymous memory, and the dirty pages. The lock status, permissions, and mapping name details are displayed, when available.

  See Displaying and Creating Charts to create a chart with different values for the x and y axis of the graph and to export the chart data.

Displaying and Configuring Graphs

When the icon that looks like a wrench appears on a page with graphs, click the icon to configure the graphs or change the graphs that appear on the page. Figure 12–9 shows an example where you can choose the system resource graphs that display on the Summary page.

Figure 12–9 Configure Graphs Menu

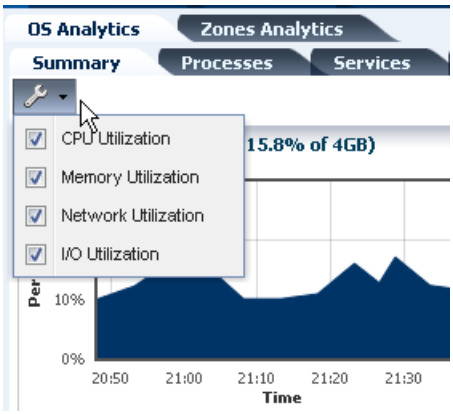

Displaying the Processes View

You can drill down to display process-specific data, based on current data from the operating system. Some data, such as CPU usage, might also be available in the History view.

The following details are available in the Process view:

- Process ID (PID)
- Process Name
- User
- State
- CPU Usage %
- Memory Usage %
- Physical Memory size in MB
- Virtual memory size
- Target

  When you click a process in the Processes table, two icons are enabled in the center pane, one to view more details and the other to kill the process. Figure 12–9 shows an example where you can click check boxes to select or deselect system resource graphs from displaying on the Summary page.

  When the software is configured to work with Oracle Enterprise Manager Cloud Control, information about available Oracle Enterprise Manager Cloud Control targets appears in the Process view.

Displaying the Services View

The Services tab provides a view of Oracle Solaris 10 and 11 Service Management Facility (SMF) services. The Oracle Solaris SMF feature defines the relationships

between applications, or services, and is a method of managing them by providing a framework for startup scripts, `init` run levels, and configuration files. Each service is identified by a Fault Managed Resource Identifier (FMRI).

The Services tab contains SMF service instances, state, dependencies, and severity information. You can drill down to see specific service details, including the configuration, dependencies, and the processes that are in the service contract.

The following actions are available from the Services tab:

- View a list of services currently installed and their states
- View a list of dependencies and dependents for FMRI
- View the relationship between services and processes
- View details about why a service is not available
- Obtain logs for debugging
- Clear faults for FMRI
- Invoke the disable, enable, and restart actions on FMRI
- Read configuration files

Services information is available for agent managed and agentlessly managed Oracle Solaris 10 and 11 operating systems, including global and non-global zones.

Displaying Thresholds

The Thresholds tab contains information about all monitored attributes for the selected operating system, including per-instance attributes such as File System Usage for each file-system on the asset.

To Display Thresholds

1. Display the Analytics view.
2. Click the **Thresholds** tab.
3. Click the **Display** menu to show the available monitored attributes.

Figure 12–10 shows a partial list of the monitored attributes available for display.

Figure 12–10 Thresholds Display Options

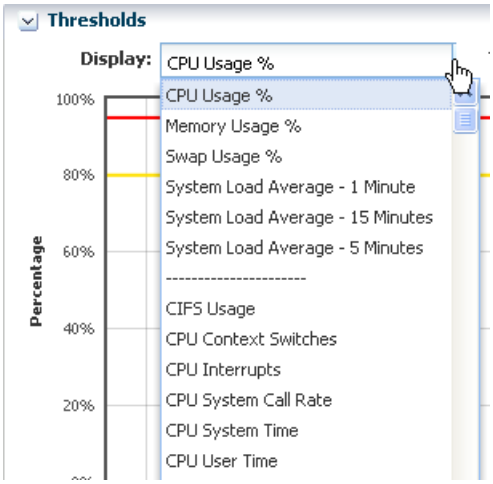

Charts for the historical data of those attributes show the alert monitor threshold levels, if configured, for that attribute. Figure 12–11 shows the chart for the System Load Average - 5 minutes. The time frame selected is Live. Under the chart are the severity levels for this monitored attribute (`SystemLoad.average15MinutePerCpuCore`), the operator, and the threshold values. Instead of editing the thresholds in the Monitoring tab, you change the values on this page. You can either add values in the New Value column or click **Suggest Thresholds** to populate the fields with suggested values. Click **Apply Thresholds** to change the existing thresholds.

Note: When the threshold monitor is modified, the asset is removed from the default monitoring profile and a custom profile created.

Figure 12–11 Thresholds Chart

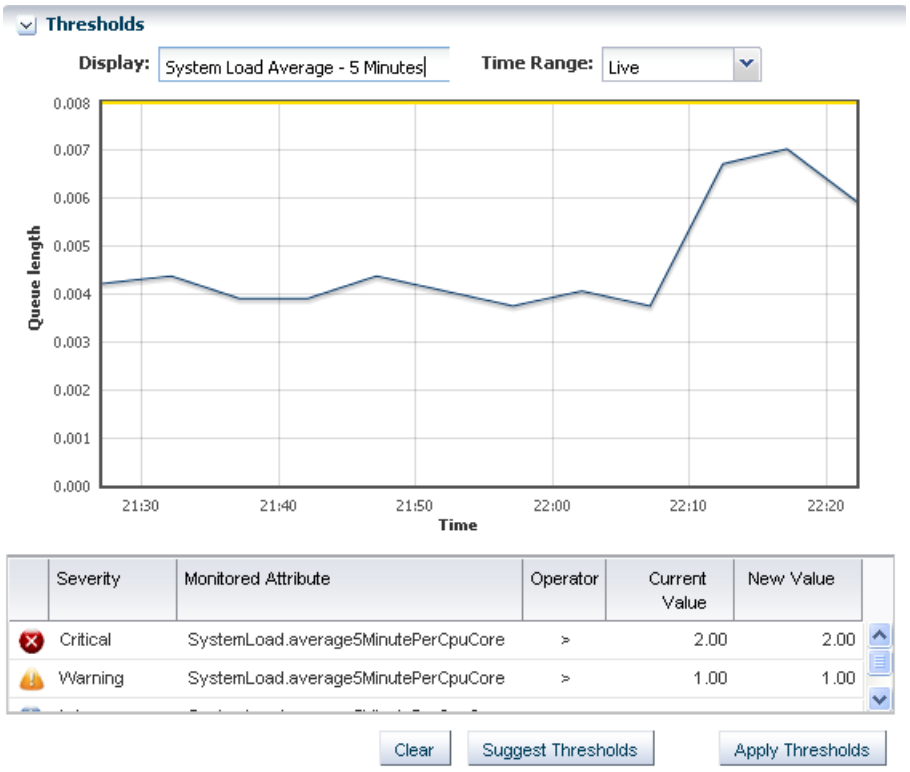

Displaying Historical Data

The History tab contains the history for the top consumers. By default, the current date displays. The time selection box is used to select a time of day for which the top consumer processes are listed. It does not affect the chart, which can only show discrete intervals of 1 hour, 1 day, 5 days, 3 weeks, 6 weeks, and 6 months. The history chart displays the most applicable chart interval for the selected date and time. Figure 12–12 is an example of a Top Consumer History chart.

To Display Historical Data

1. Display the Analytics view.
2. Click the **History** tab.
3. Click the wrench icon to select from a list of options to chart.

Figure 12–12 Top Consumer History

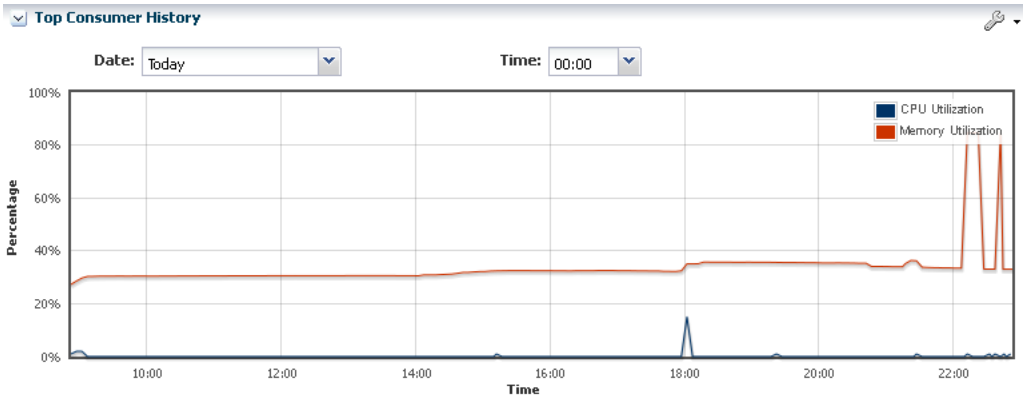

Displaying Metrics

The Metrics tab provides you with specific details about various operating system statistics, such as the percentage of memory used, and view graphs. When monitoring thresholds are enabled for an element, you can reset the thresholds. When an element does not have a monitoring threshold, you can configure a new monitoring threshold.

To Display Metrics

1. Display the Analytics view.
2. Click the **Metrics** tab to see OS-specific components.

   The following are the categories of component details available for an agent-managed Oracle Solaris OS:

   - Load
   - File Systems
   - Networks
   - Users
   - Buffer Activity
   - Disk Usage
   - Paging Activity
   - Message Activity
   - Tables Status
   - TTY Activity
   - Kernel Memory
   - DNLC
   - IPC Message Queue
   - IPC Shared Memory
   - IPC Semaphore Usage
   - CPU Detail
   - File Access
   - Disk Errors

- Memory Utilization

3. Expand a component to view the available elements.

   Click an element to view the details. In Figure 12–13, the Load component is expanded and the elements, such as percentage of memory used, CPU utilization percentage, CPU core count, and system update options are available in the Load list. Memory Used % is selected and the details and a graph are visible.

Figure 12–13 Component Navigation

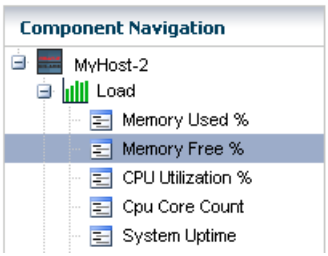

The type of resource determines the available details. Figure 12–14 is an example of the details and graph for the Memory Used % element. This example shows a system in trouble. The Details shows the last known value memory usage at 95.59111 percent and the established monitoring parameters. A warning incident is generated at 80 percent and a critical incident is generated when the used memory reaches 95 percent.

View the bar graph to see how the memory usage has trended over the selected time period. The yellow and red horizontal lines indicate the warning and critical thresholds.

Figure 12–14 Memory Used % Details and Graph

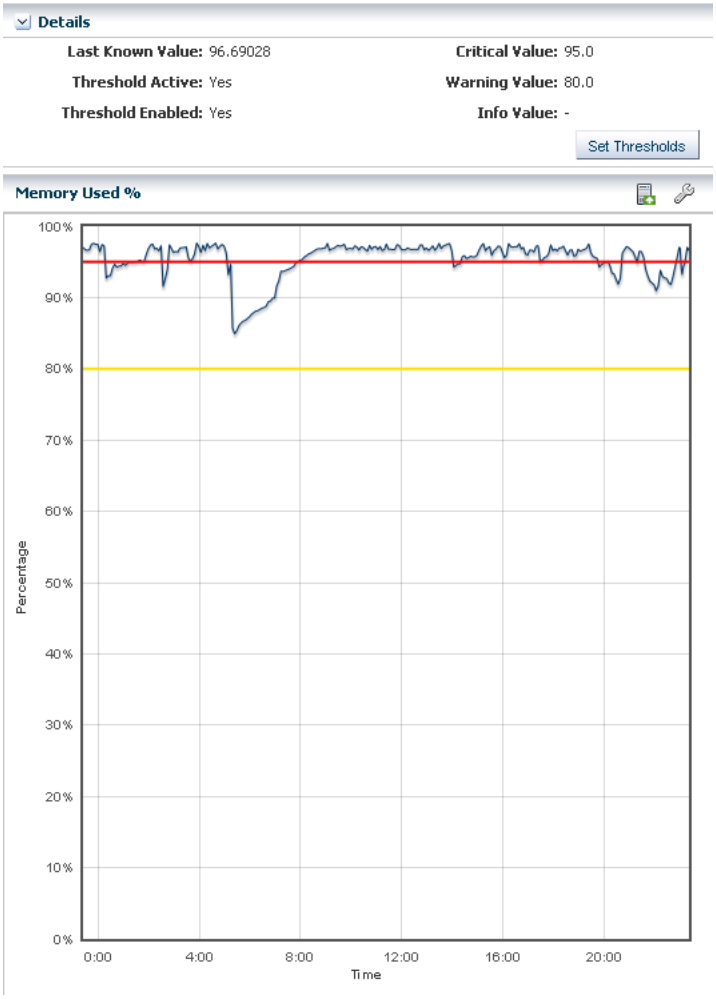

The wrench icon appears above the graph. For this graph, you can click the icon to change the graph style to line, bar, or area. You can also change the definition of the X axis from 1 day to 1 hour, 5 days, 3 weeks, 6 weeks, or 6 months. The export icon enables you to export the graph to file in CSV or XML format.

You have the opportunity to reset the threshold values for the monitoring rule from this page, or to view recommended threshold values based on past performance. When you click the Set Thresholds button, you are taken to the Thresholds tab, where you can enter new threshold values. When the Thresholds tab appears, the drop-down menu contains the name of the monitored attribute from the metrics view, and the associated chart appears. Editing a threshold value creates a custom set of monitoring rules for the asset. See Chapter 4, "Monitoring Rules and Policies" for more information about editing individual monitoring rule parameters and creating a custom monitoring policy.

Displaying and Creating Charts

The Charts tab enables you to modify the following charts to change the type of chart and to see utilization data over longer periods of time, up to six months:

- Power Utilization: Servers, chassis

- CPU Utilization: Operating system, operating system for a virtualization host, virtual host, server pool
- Memory Utilization: Operating system, virtual host, server pool
- Network Utilization: Operating system, operating system for a virtual machine, virtualization host, server pool
- File System Utilization: Oracle Solaris OS and Linux OS
- System Load: Oracle Solaris OS and Linux OS

For the first five days of operation, Oracle Enterprise Manager Ops Center collects data every five minutes. After the fifth day, the reported data is an average, according to the following:

- Five days to three weeks: Average for each hour
- Three weeks to six weeks: Average for each 12-hour period
- Six weeks to six months: Average for each 24-hour period

Oracle Enterprise Manager Ops Center deletes the data after six months.

To Display Charts

1. Display the Analytics view.
2. Click the **Charts** tab.

You can create charts 24 hours after you first manage the asset. The asset must be operating and the Enterprise Controller must be able to get access to the asset. You can use the default format or change the charts to use a line, bar, or area format or to use different time intervals.

To Create Charts

1. Display the Analytics view.
2. Click the **Charts** tab.
3. Select an option from the Chart type menu to display the output as a line, bar, or area chart.
4. Select **Live**, **1 Hour**, **1 Day**, **5 Days**, **3 Weeks**, **6 Weeks**, or **6 Months** from the X Axis (Time) menu to change the time frame for the data.
5. Select **Percentage** or **Unit of Measure** to configure the chart's Y Axis.
6. Select the component and options to plot from the Plot selection.

By default, the graphs show one day of data. Select **Live** in the X Axis menu to change to Live mode, which reports new information every 5 seconds. You can also change the graph to one of the following periods:

- One hour (1H): One point for every 5 minutes
- One day (1D): One point for every 5 minutes
- Five days (5D): One point for every 5 minutes
- Three weeks (3W): One point every hour
- Six weeks (6W): One point every 12 hours
- Six months (6M): One point for every day. To make a graph with the minimum of two points, a system must have been managed for at least 10 minutes to view a one-hour graph and for at least two days to view the six-months graph.

The software stores the data for these time periods separately. For example, when a server is managed for two hours and you select the 6W view, the graph does not display because only one point of data of that type is available; the second point has not yet occurred. When you select the 1D view, the graph displays 24 points of data (120 minutes in 5-minute intervals). However, the graph displays these points over a 24-hour period and not over the actual two-hour period. For the most accurate representation of the data, choose a time period that is less than or equal to the time that the selected server has been managed.

You can export the data for either the current view or all available data to a file in either CSV or XML format. Click **Export Chart Data** to choose options for exporting the data.

When the graph is blank for a server, one of the following conditions has occurred:

- Server does not have the appropriate ILOM version.
- Server has not been discovered through the ILOM driver.
- Server is unreachable.

Displaying Virtualization Analytics

The Virtualization Analytics view displays resource usage of the physical server for each running operating system instance, showing the physical resources consumed by the control domain (global zone or Oracle VM Server), and each non-global zone or guest. Metrics for Oracle VM Server for x86 are available through the Oracle VM Manager.

- Virtualization Analytics Summary
- Virtualization Analytics Zones or Virtualization Analytics Guests
- Virtualization Analytics History

The information is refreshed every 20 seconds for a guest running on a virtualization server container, including global zones, Oracle VM Server for SPARC, and Oracle VM Server for x86.

Customizing the Analytics View

The default display is the resource usage data. You can change the display settings for each view and choose which system resource graphs and details to display. The changes made on a given analytics page affects the views of all analytics screens for the same OS or virtualized OS.

To Customize the Analytics View

1. Select an operating system (global zone) or zone in the Navigation pane.
2. Click the **Analytics** tab.
3. Click the **Settings Menu** icon in the center pane.
4. If you select a global zone that has zone, click the **Zones Analytics** subtab, then click the **Settings Menu** icon in the center pane to change the view for the Zones Analytics page.

Overview of Oracle Solaris Boot Environments

Boot environments are a feature of Oracle Solaris. A boot environment is an instance of a bootable Oracle Solaris image plus additional software packages that are installed

onto the image, and the set of all file systems and devices (disk slices and mount points) that are required to operate an Oracle Solaris OS instance. You can have disk slices, also known as partitions, on the same disk or distributed across several disks.

A dual boot environment consists of a live, or active, boot environment (BE) and one or more inactive alternate boot environments (ABE). A system can have only one active boot environment, which is the booted environment. An alternate boot environment is an inactive environment that is not currently booted. A system can have many inactive boot environments. You can activate an alternate boot environment at any time.

You can use a dual boot environment within Oracle Enterprise Manager Ops Center to manage your Oracle Solaris software. A dual boot environment is often used to manage updates because it can significantly reduce the service outage time that is usually associated with patching. Maintaining multiple boot environments also enables quick and easy rollback to a version before the patches were applied, if needed. The boot environment technology enables you to duplicate a boot environment and perform the following tasks without affecting the currently running system:

- Run an Oracle Solaris software update simulation on the inactive boot environment. You can run the simulation with or without downloading the patches.
- Update your Oracle Solaris OS on the inactive boot environment and test the update before deploying it as your active environment.
- Maintain multiple boot environments with different images. For example, you can create one boot environment that contains all current patches and another that contains only security patches.

Understanding the Differences Between Oracle Solaris 11 and Oracle Solaris 10 Boot Environments

The boot environment feature changed beginning with Oracle Solaris 11, including the supported file systems, file system requirements, zone support, and how boot environments are created.

Oracle Solaris 11 Boot Environments use the `beadm` utility and ZFS file systems to create and manage boot environments. You do not need to create boot environments ahead of time. The software creates them automatically. You can use an agent-managed or agentless-managed operating system with the Oracle Solaris 11 Boot Environments feature.

Oracle Solaris 10 and earlier use the Live Upgrade feature with `lucreate` scripts and ZFS or UDFS file systems to create alternate boot environments. You must use an agent-managed operating system with the Oracle Solaris Live Upgrade feature.

The Boot Environments profile defines the boot environment for the OS update deployment plans. The Oracle Solaris release determines how the profile is used:

- Oracle Solaris 11: The profile indicates the policy that is used when an OS update plan is executed. Use this profile to define the creation policies and as a step in an OS update deployment plan.
- Oracle Solaris 10: The profile defines the `lucreate` script that is used to create your alternate boot environments. Use this profile to specify how to create alternate boot environments for the eligible operating systems in your data center, create alternate boot environments, synchronize and activate boot environments, and as a step in an OS update deployment plan.

The user interface assists you in easily navigating the differences in the boot environments features and provides a unified view whenever possible. See Overview of Oracle Solaris 11 Boot Environments and Overview of Oracle Solaris 10 Boot Environments for detailed version-specific information.

Monitoring Boot Environments

Monitoring rules and thresholds are defined in the Monitoring tab. The following rules apply to boot environments:

- Number of Boot Environments: Defines the number of boot environments in the selected zone.
- Boot Env Usage Percent in a ZPool: Defines the percentage disk utilization on the boot environments in any of the zpools.

Each rule has two defined thresholds:

- Warning: Generates a warning alert and displays the yellow warning badge in the Asset navigation tree. A warning alert contains a suggested action.
- Critical: Generates a critical alert and displays the red critical badge in the Assets tree. By default, a critical alert contains an automated action.

See Monitoring Operating Systems for more information.

To Clear a Boot Environment Incident

Perform the following steps to manually clear the disk space and close the incident:

1. Click an OS in the Assets tree to view the Boot Environment in the center pane, then click the **Incidents** tab.
2. Delete one or more boot environments to clear the disk space.
3. Select one or more incidents, then click the **Close Incidents** icon in the center pane.

See Chapter 4, "Monitoring Rules and Policies" for details on how monitoring rules and policies work in the software and see Chapter 9, "Incidents" for information on how to manage and close an incident.

Viewing Boot Environments

The active boot environment is identified in the Boot Environments tab and in the Summary tab. Inactive boot environments appear in the Boot Environments tab for each operating system. For global zones, the Boot Environments tab displays the relationship between the associated zones and existing boot environments.

Note: If there are no alternate boot environments, the Boot Environments tab does not appear in the center pane.

Boot environment support for Oracle Zones is available beginning with Oracle Solaris 11, enabling boot environments from the non-global zone to appear in the UI. Boot environments for non-global zones appear in the Boot Environment tab, not the Assets tree. Select a non-global zone in the Assets tree, then click the Boot Environments tab in the center pane.

The Summary tab provides the name of the active boot environment, the zpool utilization of all zpools for the selected OS and the amount of zpool utilization

attributed to the boot environments. Details are available in the Boot Environments tab.

1. Click an OS or zone in the Assets tree. Details about the active boot environment appear in the Summary tab.

   Non-bootable snapshots are available beginning with Oracle Solaris 11. When a boot environment has one or more associated non-bootable snapshots, the snapshots appear in the Snapshots subtab, as shown in Figure 12–15.

Figure 12–15 Boot Environment Tab and Snapshots Subtab

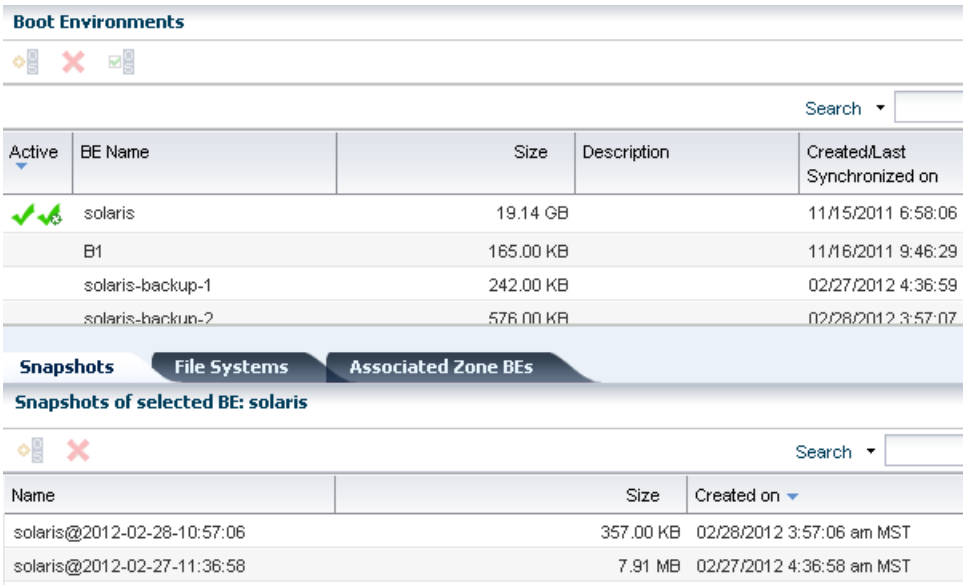

2. Click the **Boot Environments** tab in the center pane to see details about the associated alternate boot environments.

   If the tab is disabled, no alternate boot environments exist for that OS in the zone or non-global zone.

3. To display file system details, click the boot environment in the table, then click the **File Systems** subtab:

   - **Oracle Solaris 11**: Mount point information, boot environment file systems, zone boot environment information, and non-global zones in the selected boot environment. When a global zone boot environment is selected, the zone boot environment details are listed separately, a zone per row. You can expand all rows to display a complete view of all zone boot environments under the selected global zone boot environment.
   - **Oracle Solaris 10**: Mount point information and boot environment file systems.

4. To display associated zone boot environment details, click the boot environment in the table, then click the **Associated Zone BEs** subtab:

   - **Oracle Solaris 11**: Mount point information, boot environment file systems, zone boot environment information, and non-global zones in the selected boot environment. When a global zone boot environment is selected, the zone boot environment details are listed separately, a zone per row. You can expand all rows to display a complete view of all zone boot environments under the selected global zone boot environment.

- **Oracle Solaris 10**: Mount point information and boot environment file systems.

Monitoring rules and thresholds are defined in the Monitoring tab.

Managing Boot Environments

You can view all available boot environments for a system, and choose to activate an alternate boot environment or delete inactive environments.

Activating and Reboot a Boot Environment

Activating a boot environment makes an inactive or alternate boot environment the active boot environment. You can activate a single boot environment, or you can select an OS group and activate an alternate boot environment for each operating system.

To Activate a Single Boot Environment

Information about boot environments appears in the Boot Environments tab for each operating system.

Perform the following steps to activate a single boot environment:

1. Click an Oracle Solaris 11 global or non-global zone or an Oracle Solaris 10 global zone in the Assets tree.
2. Click the **Boot Environments** tab.

   Existing boot environments for the OS appear in the center pane.

To Activate Boot Environments for All Members of a Group

You can activate an alternate boot environment for all members of a group. If each OS in the group has a single alternate boot environment, all operating systems are booted into the alternate boot environment. When some systems have more than one alternate boot environment, you are prompted to select the alternate boot environment.

Perform the following steps to activate boot environments for all members of an operating system group:

1. Select an Oracle Solaris 11 or an Oracle Solaris 10 Operating system group from the Assets section in the Navigation pane.
2. Click **Activate Boot Environment and Reboot** in the Actions pane.
3. For systems with multiple alternate boot environments, select the one that you want to boot into. A Filter ABE by name option is available to identify similarly named alternate boot environments across multiple targets.

Deleting a Boot Environment

When you delete a boot environment, you delete all associated snapshots and unshared file systems. A snapshot is a non-bootable copy of a boot environment. If there are non-global zone boot environments associated with the global zone boot environment, they are deleted too. Shared file systems are not deleted.

You cannot delete the active boot environment.

1. Click an Oracle Solaris 11 global or non-global zone, or an Oracle Solaris 10 global zone in the Assets tree.
2. Click the **Boot Environments** tab.

3. Select one or more boot environments or snapshots that you want to remove, then click the **Delete** icon. Figure 12–16 shows a snapshot selected for deletion.

Figure 12–16 Delete an Oracle Solaris 11 Snapshot

Overview of Oracle Solaris 11 Boot Environments

Oracle Solaris 11 uses ZFS file systems, where the swap and dump volumes are shared within the pool. For unshared file systems, ZFS is the only supported file system for boot environments on Oracle Solaris 11. With the ZFS-based boot environment, the boot environments are clones of the existing ZFS partitions. This saves disk space and you do not need to reserve disk partitions for additional boot environments. The create and activate boot environment functionality is much faster than in previous versions.

Oracle Solaris 11 boot environments are managed through the `beadm` utility. A new boot environment is created whenever the kernel or system packages are installed or updated. This can result in high disk space utilization levels. By default, Oracle Enterprise Manager Ops Center monitors the disk (zpool) utilization of boot environments. If the utilization levels exceed defined thresholds, you can delete unwanted boot environments.

When a boot environment is created in the global zone, the following occurs:

- A Boot Environment of the source boot environment is created (the boot environment from which it is cloned).
- The currently active boot environment in all of the non-global zones is cloned and it is associated with the global zone boot environment that was just created.

When a global zone boot environment is activated, the active boot environment data set that is associated with that boot environment in each non-global zone is mounted and activated. Only one non-global zone boot environment that is associated with a global zone boot environment can be active.

When a global zone boot environment is deleted, all corresponding zone-specific boot environments are also deleted.

When a non-global zone is deleted, all corresponding boot environments are deleted. The boot environments for other zones are not deleted.

Displaying Oracle Solaris 11 Boot Environment Details

The Oracle Solaris 11 Boot Environments tab provides you with a large amount of information about the boot environments that are associated with the selected physical or virtual Oracle Solaris 11 operating system. This is particularly valuable since Oracle Solaris 11 boot environments are automatically created and they can quickly consume valuable resources. The following information is available:

- Displaying Total ZPools Utilization for Oracle Solaris 11 Boot Environments
- Displaying Oracle Solaris 11 Boot Environments

- Displaying Snapshots for Oracle Solaris 11 Boot Environments
- Displaying File Systems for Oracle Solaris 11 Boot Environments
- Displaying Associated Zone Boot Environments

Displaying Total ZPools Utilization for Oracle Solaris 11 Boot Environments

The first section in the Boot Environments tab is Total ZPools Utilization. It is compressed by default. Use the arrow to expand the table. The amount of resources used by the zpool appears in this section, as shown in Figure 12–17.

Figure 12–17 Oracle Solaris 11 Total ZPools Utilization

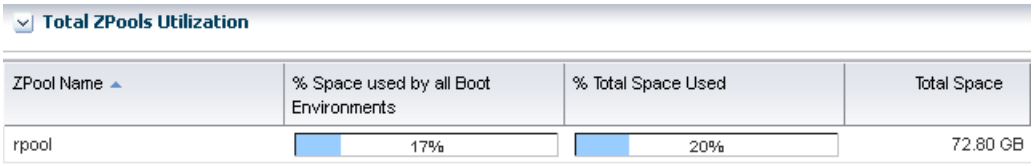

Displaying Oracle Solaris 11 Boot Environments

All boot environments that are associated with the physical or virtual operating system that is selected in the Assets tree appear in the Boot Environments table. As shown in Figure 12–18, the active status, size, and when the boot environment was created appear in this section. A green check mark icon next to the boot environment name identifies the boot environment as active. A green check mark with a green circle and white x indicates that this is the boot environment that is active upon reboot. When two green check marks are visible in the Active column, the boot environment is active and is the active boot environment upon reboot.

Figure 12–18 Oracle Solaris 11 Boot Environments Tab

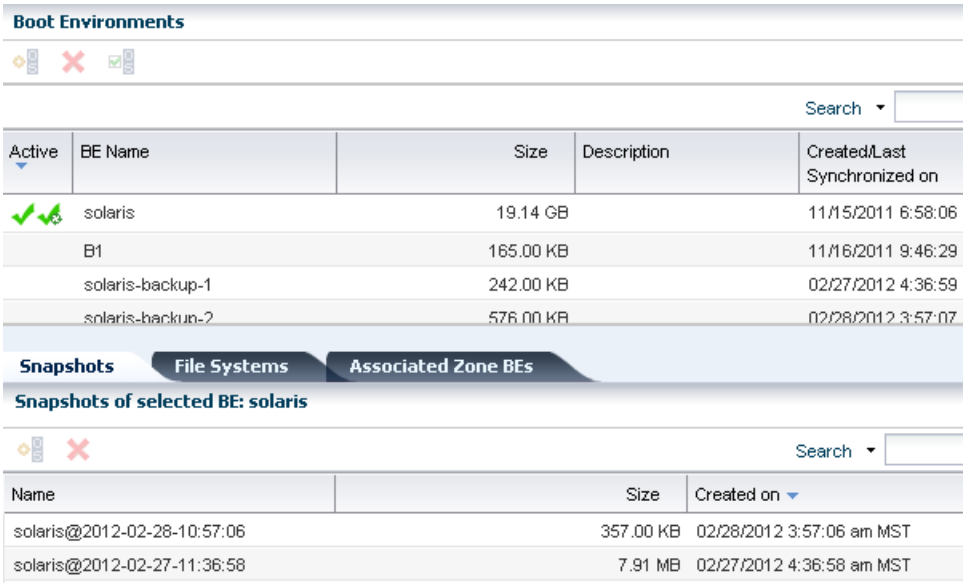

The lower section of the page has three tabs that provide details about the boot environment that you select in the Boot Environments table.

Displaying Snapshots for Oracle Solaris 11 Boot Environments

A snapshot is a point-in-time image of a Volume. It is a non-bootable copy of a boot environment that uses much less disk space than a boot environment. You can create a boot environment from a snapshot.

Non-bootable snapshots are available beginning with Oracle Solaris 11. Select a boot environment in the Boot Environments table to see all associated non-bootable snapshots in the Snapshots tab. You can select a snapshot in this tab and click the icon to create a boot environment from the snapshot.

Figure 12–19 Oracle Solaris 11 Boot Environment Snapshots

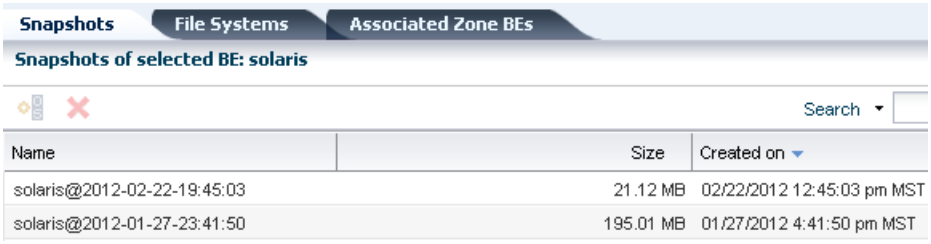

Displaying File Systems for Oracle Solaris 11 Boot Environments

The File Systems tab provides file system details for the boot environment, or alternate boot environment, that you select in the Boot Environments tab. As shown in Figure 12–20, the file system name, type, size and mount location appear in this table. You can also see if the file system is shared or not.

Figure 12–20 Oracle Solaris 11 Boot Environments File Systems

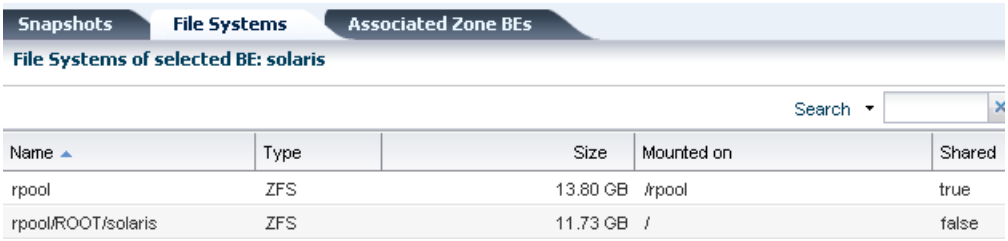

Displaying Associated Zone Boot Environments

The Associated Zone BEs tab is populated with the boot environment details for a selected zone's boot environment. The zone might have multiple boot environments. The table shows when the boot environment selected is active. As shown in Figure 12–21, you can see the boot environment name, the zone name, the size and when the boot environment was created. When you select an operating system that is not a zone, No Data appears in the table.

Figure 12–21 Oracle Solaris 11 Boot Environments Associated Zone BEs

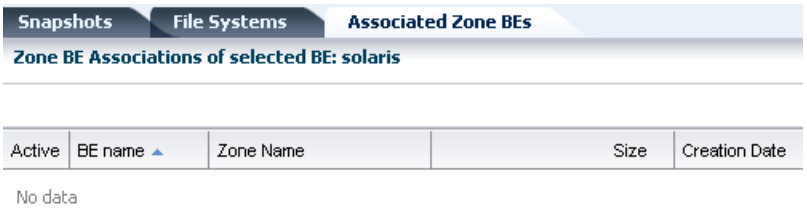

About Boot Environment Profiles and Plans for Oracle Solaris 11

A Boot Environments plan and profile for Oracle Solaris 11 operating systems identifies and defines how the boot environment is updated and activated. You can use an agent-managed or agentless-managed operating system.

Use the profile for Oracle Solaris 11 to perform the following tasks:

- Create a Boot Environments profile that defines the creation policies
- Create an OS update deployment plan, specifying a Boot Environments profile and an OS update profile

You can create profiles with different policy options. The following options are available for Oracle Solaris 11:

- Create a new boot environment only when needed, such as when the OS update operation requires a reboot.
- Always create a new boot environment.
- Never create a new boot environment.
- Activate and reboot.

> **Note:** Oracle Solaris 11 boot environments cannot have spaces in the name. When you name a boot environment, do not use spaces.

To Create a Boot Environments Profile for Oracle Solaris 11

1. Click **Plan Management** in the Navigation pane, then scroll down to the Profiles and Policies directory and click **Boot Environments**.

   The Boot Environments Profiles page appears in the center pane.
2. Click **Create Profile** in the Actions pane.
3. Enter a unique profile name, without spaces, and provide a description to identify the profile. Select the **Oracle Solaris 11** subtype.
4. Choose the boot environment policy for this profile:
   - Create a new boot environment when needed.
   - Always create a new boot environment.
   - Never create a new boot environment.
5. Click the check box to **activate boot environment and reboot** when the job is completed.
6. Choose to have the software automatically create a unique name or enter a specific boot environment name in the field provided, then click **Finish**.

To Copy an Oracle Solaris 11 Boot Environments Profile

Copying a boot environment profile makes a copy of an existing profile.

1. Click **Plan Management** in the Navigation pane, then scroll down to the Profiles and Policies directory and click **Boot Environments**.

   The Boot Environments Profiles page appears in the center pane.
2. Select the profile to copy, then click the **Copy Profile** icon in the center pane.
3. Enter a unique profile name, without spaces, and a description.

4. Change the boot environment parameters, as needed, then click **Finish**.

Creating an Oracle Solaris 11 Boot Environment

A new alternate boot environment is automatically created when you install or update the operating system's kernel or system packages.

You can create a bootable or non-bootable copy an existing boot environment. A non-bootable copy is call a snapshot. You can create a boot environment from a snapshot.

To Create an Oracle Solaris 11 Boot Environment

You can create a bootable or non-bootable (snapshot) copy of an Oracle Solaris 11 boot environment.

1. Click an Oracle Solaris 11 OS, a global or non-global zone, in the Assets tree, then click the **Boot Environments** tab.
2. Select the boot environment to copy, then click the **Create** icon.
3. Select either the **Create a Snapshot** or **Create a BE** option, then select the boot environment to use as the source. Enter a unique name for the boot environment, without spaces, and provide a description that includes the purpose of the boot environment and the name of the source boot environment. Click **Create**.

Figure 12–22 Create Boot Environment for Oracle Solaris 11

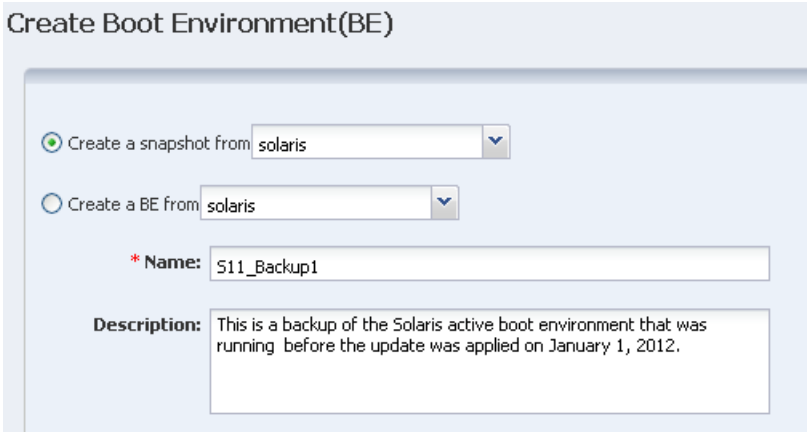

4. When the job finishes, refresh the Boot Environment table to see the new boot environment.

If you have a group that contains only Oracle Solaris 11 operating systems and you select the OS group as the target, a new boot environment is created for every active boot environment in the group. You can assign a common name for all of the boot environments. If you do not assign a common name, default names are assigned.

Overview of Oracle Solaris 10 Boot Environments

Boot environments in Oracle Solaris 10 are managed through the Live Upgrade feature. A boot environment is created by using a script that contains the `lucreate` command and options.

The Live Upgrade feature is available on the following operating systems:

- Oracle Solaris 10 OS for x86 platform versions through Oracle Solaris 10 5/09: Update alternate boot environments for physical systems.
- Oracle Solaris 10 OS for SPARC through Oracle Solaris 10 5/09: Update alternate boot environments for physical and virtual machines, including Oracle Solaris Zones and Oracle VM Server for SPARC (formerly known as Logical Domains).

> **Note:** Do not use Oracle Solaris Live Upgrade on your Enterprise Controller or Proxy Controllers. It does not synchronize all of the files that are required for these Oracle Enterprise Manager Ops Center components.

To create and update alternate boot environments, install the latest Oracle Solaris Live Upgrade packages and patches. This ensures that you have all the latest bug fixes and features. In addition, verify that the disk space and format are sufficient for an alternate boot environment.

Live Upgrade packages and patches are available for each Oracle Solaris software release beginning with Oracle Solaris 10 OS. Use the packages and patches that are appropriate for your software instance.

If you installed Oracle Solaris 10 using any of the following software groups, have the required packages:

- Entire Oracle Solaris Software Group Plus OEM Support
- Entire Oracle Solaris Software Group
- Developer Oracle Solaris Software Group
- End User Oracle Solaris Software Group

When you install one of these Software Groups, verify that you have all of the required packages:

- Core System Support Software Group
- Reduced Network Support Software Group

For detailed requirements to install and use the Live Upgrade feature, see *Oracle Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* at `http://docs.oracle.com/cd/E18752_01/html/821-1910/preconfig-17.html`. For information on other releases, see *Operating Systems Documentation* at `http://docs.oracle.com/en/operating-systems/`. Review and verify that all the packages and patches that are relevant to your system are installed and that the disk is formatted properly before creating a new boot environment.

For Oracle Solaris 10 alternate boot environments, file systems are categorized into the following types:

- **Critical File Systems**: Non-sharable file systems that are required by the Oracle Solaris OS, such as root (`/`), `/usr`, `/var`, and `/opt`. These file systems are separate mount points in the `vfstab` of the active and inactive boot environments and are always copied from the source to the inactive boot environment.
- **Sharable File Systems**: User-defined files, such as `/export`, that contain the same mount point in the `vfstab` in both the active and inactive boot environments. Updating shared files in the active boot environment also updates data in the inactive boot environment. When you create a boot environment, sharable file systems are shared by default. If you specify a destination slice, also known as a partition, the file systems are copied.

- **Swap**: Depends on the type of file system:
  - For UFS file systems, swap is a special sharable volume. Like a sharable file system, all swap slices are shared by default. If you specify a destination directory for swap, the swap slice is copied.
  - For ZFS file systems, swap and dump volumes are shared within the pool.

When creating a new boot environment in Oracle Solaris 10, the entire contents of a slice is copied to the designated new boot environment slice. You might want some large file systems on that slice to be shared between boot environments rather than copied to conserve space and copying time. File systems that are critical to the OS such as root (/) and /var must be copied. File systems such as `/home` are not critical file systems and could be shared between boot environments. Sharable file systems must be user-defined file systems and on separate swap slices on both the active and new boot environments. You can reconfigure the disk several ways, depending your needs.

You can reslice, or partition, the disk before creating the new boot environment and put the sharable file system on its own slice. For example, if the `root (/)` file systems, `/var`, and `/home` are on the same slice, reconfigure the disk and put `/home` on its own slice. When you create any new boot environments, /home is shared with the new boot environment by default. To share a directory, the directory must be split off to its own slice.

Requirements for Oracle Solaris 10 Live Upgrade and Oracle Solaris Zones

The following requirements are needed to use the Live Upgrade feature with zones:

- Agent managed operating system
- At least Oracle Solaris 10 5/09 (update 7) operating system
- Storage library used to house the zones cannot be part of the root pool; you must create a separate pool on a shared file system
- You cannot use the `-p` option to create alternate boot environments

The `-p` option, which copies between two root pools on ZFS configuration, is not supported with the `lucreate` command.

Note: If you plan to use alternate boot environments with zones, you must designate sufficient zone storage space. When you create the zones and configure the zone storage, specify twice the size of the zone file system for the root file system of the zone. For example if your zone root file system was configured as 8 GB, the storage used to back up the zone must be at least 16 GB.

See Updating Zones for how to update Oracle Solaris Zones without using a dual boot environment.

Displaying Boot Environment Details for Oracle Solaris 10

The Oracle Solaris 10 Boot Environments tab provides you with a large amount of information about the boot environments that are associated with the selected Oracle Solaris operating system. A Boot Environments tab is not available for Oracle Solaris 10 non-global zones. The following information is available:

- Displaying Total ZPools Utilization for Oracle Solaris 10 and Earlier Boot Environments

- Displaying Oracle Solaris 11 Boot Environments
- Displaying File Systems for Oracle Solaris 10 and Earlier Boot Environments

Displaying Total ZPools Utilization for Oracle Solaris 10 and Earlier Boot Environments

The first section in the Boot Environments tab is Total ZPools Utilization. It is compressed by default, use the arrow to expand the table. The amount of resources used by the zpool appears in this section, as shown in Figure 12–23.

Figure 12–23 Oracle Solaris 10 and Earlier Total ZPools Utilization

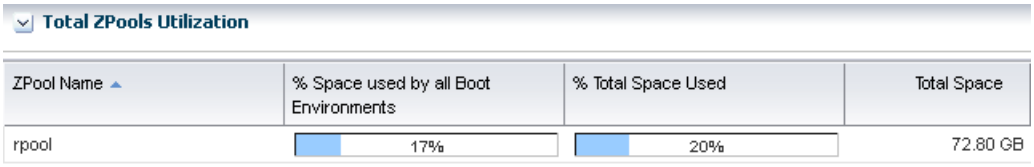

Displaying Boot Environments for Oracle Solaris 10 and Earlier

All boot environments that are associated with the physical or virtual operating system that is selected in the Assets tree appear in the Boot Environments table. As shown in Figure 12–24, the active status, size, description, and when the boot environment was created or synchronized appear in this section. A green check mark icon next to the boot environment name identifies the boot environment as active. A green check mark with a green circle and white x indicates that this is the boot environment that is active upon reboot. When you see two green check marks in the Active column, the boot environment is active and is the active boot environment upon reboot.

Figure 12–24 Boot Environments for Oracle Solaris 10 and Earlier

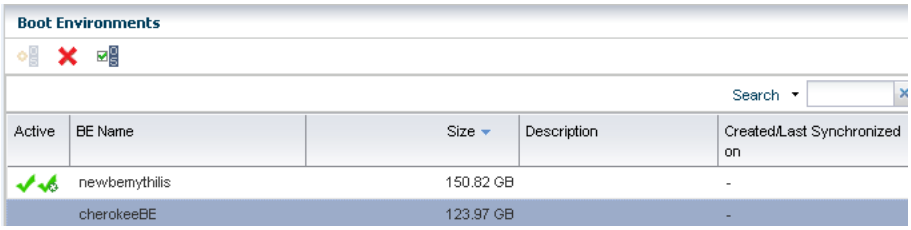

The lower section of the page the File Systems tab that provide details about the boot environment that you select in the Boot Environments table.

Displaying File Systems for Oracle Solaris 10 and Earlier Boot Environments

The File Systems tab provides file system details for the boot environment, or alternate boot environment, that you select in the Boot Environments tab. As shown in Figure 12–25, the file system name, type, size and mount location appear in this table. You can also see if the file system is shared or not.

Figure 12–25 Oracle Solaris 10 Boot Environments File Systems

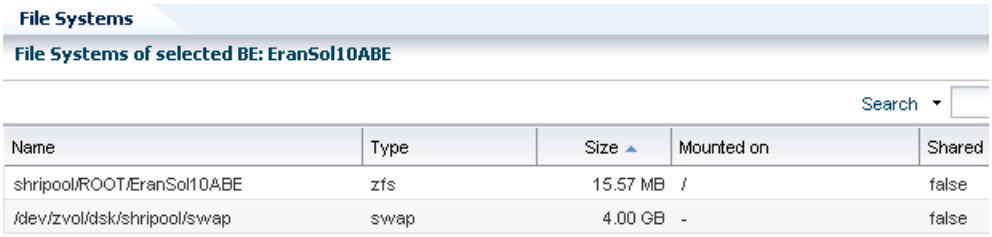

About Boot Environment Profiles and Plans for Oracle Solaris 10 and Earlier

You can use an alternate boot environment that was created outside of Oracle Enterprise Manager Ops Center; however, the preferred method is to create a Boot Environment profile and use the associated plan to create a boot environment. In either case, the operating system must be an agent-managed asset in the Oracle Enterprise Manager Ops Center software.

Use the profile for Oracle Solaris 10 and earlier releases to perform the following tasks:

1. Create a Boot Environments profile to specify how to create boot environments across all your managed system, using `lucreate` scripts.
2. Create your boot environments using the action in the action panel for your assets. This is typically a one time operation, as you'll probably reuse the alternate boot environment.
3. Create a Boot Environments profile to always sync (and possibly) activate. Use this profile for OS update deployment plans.
4. Create an OS update deployment plan specifying the boot environment in step 3 and any OS update profile.

Depending on the Oracle Solaris version that you are using, several methods are available for creating alternate boot environments. The Oracle Enterprise Manager Ops Center software uses Boot Environments profiles to help standardize and simplify the process.

Perform the following tasks before creating a Boot Environments profile:

1. Determining Your Boot Environment Policy
2. Determining Your Live Upgrade Feature Requirements and Options

Determining Your Boot Environment Policy

Each Boot Environments profile can have a different boot environment policy. The following options are available for Oracle Solaris 10:

- Activate Boot Environment and reboot on job completion.
- Synchronize existing Boot Environment before submitting a job.

Determining Your Live Upgrade Feature Requirements and Options

For Oracle Solaris 10 only, determine the file system and swap requirements and the `lucreate` command options that you want to use. If you are using ZFS, create a zpool. See *Oracle Solaris 10 9/10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* available at: `http://docs.oracle.com/cd/E18752_01/html/821-1910` for more information about the requirements and command options for the Live Upgrade feature.

You can create an operational plan that includes your script, or you can enter the `lucreate` options directly in the Boot Environments profile without creating a script.

Overview of Boot Environments Profiles and Plans

The following is a high-level overview of how to create an alternate boot environment for Oracle Solaris 10 and earlier releases.

1. Create a Boot Environments profile.

   The profile identifies and defines how the boot environment is created and activated, including defining the target type, the options used to create the boot environment, and the activation parameters. Review *Oracle Solaris 10 9/10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* at: `http://docs.oracle.com/cd/E18752_01/html/821-1910` for more details on the requirements:

   – Define your file system and swap requirements.

   – Determine the `lucreate` command options. You can enter the options in the Boot Environments profile, or you can create an operational profile that contains a script with the `lucreate` command and options. If you are specifying a variable boot environment name in the script passed to the operational profile, make sure the variable is called *BEName*. When you use *BEName* as the variable name, the operational plan passes on the BE name entered during plan execution to the `lucreate` script.

   – Determine your boot environment policies.

2. Create a deployment plan.

   The Boot Environment deployment plan defines the failure policy and is associated with a single Boot Environments profile. You can add a Boot Environments profile as a task, or step, within a complex deployment plan. See Plan Management for information about complex plans.

3. Execute the deployment plan.

   After the deployment plans are created, a user with the appropriate role and privileges can choose from a list of plans to quickly and consistently create an alternate boot environment.

4. Create and deploy an OS Update plan to update the inactive boot environment. See Updating an Oracle Solaris Boot Environment for the steps to create and deploy the OS Update plan.

For information about system administration tasks such as managing file systems, mounting, booting, and managing swap space, see the *Oracle Solaris System Administration Guide: Devices and File Systems* in the Oracle Solaris 10 Documentation Library.

Creating an Oracle Solaris 10 Boot Environment

Oracle Solaris Live Upgrade scripts are used to create an alternate boot environment. To create a replica of your boot environment, run the script in Oracle Enterprise Manager Ops Center.

The following methods are available to create the alternate boot environment:

- Copy an existing boot environment.
- Execute a Boot Environments deployment plan.

The deployment plan must reference an Boot Environments profile.

- Run an operational plan that contains the `lucreate` script.
- Upload an Oracle Solaris Live Upgrade script as Local Content in Enterprise Manager Ops Center.
  - Run an OS update job and specify a pre-action which runs the script. You can select multiple compatible targets and create an alternate boot environment for each target using the same script at the same time.
  - Create an OS profile, and then run an OS Update job. The profile enables you to define the components and the actions to be performed every time you use the profile to create an alternate boot environment.
- Run an Oracle Solaris Live Upgrade script from the command-line interface. With this method, you must log in to each agent and then run the script to create the boot environment.

When you create the alternate boot environment with an OS Update job, you can choose to run the job immediately, or you can schedule the job to run during your maintenance window. In all methods, the new boot environment is automatically discovered and a new Boot Environments tab appears in the center pane for OS management.

This task describes how to run a New OS Update job to create the alternate boot environment. Although it is a New OS Update job, the sole purpose of the job is to create an alternate boot environment. The job uses the Live Upgrade script that you uploaded to Local Content to create a duplicate of your boot environment.

Defining Deployment Options for Oracle Solaris 10

To create boot environments for Oracle Solaris 10, you can define the options in an Boot Environments profile or you can add a script to an operational profile. The software enables you to store the scripts and parameters and provides a couple of methods of deploying the scripts.

- As a profile
- As part of a plan
- As part of an update job

The scripts use the `lucreate` command to create an alternate boot environment. For a complete list of options, see the Oracle Solaris 10 lucreate(1M) documentation at `http://docs.oracle.com/cd/E19253-01/816-5166/lucreate-1m/index.html`. The following are some commonly used options:

- `-m` option to create a boot environment and split a directory off to its own slice.

  To share a directory, the directory must be split off to its own slice. The directory is then a file system that can be shared with another boot environment. You can use the `lucreate` command with the `-m` option to create a boot environment and split a directory off to its own slice. But, the new file system cannot yet be shared with the original boot environment. You must run the `lucreate` command with the `-m` option again to create another boot environment. The two new boot environments can then share the directory.
- `-s` option to force a synchronization.
- `-p` specifies the ZFS pool in which a new boot environment resides.

The `-p` is not required when the source and target boot environments are within the same pool. The `-p` option does not support the splitting and merging of file systems in a target boot environment, instead use the `-m` option.

Existing boot environments appear in the hierarchy in the Boot Environments tab for a global zone. You can create the boot environment with Oracle Enterprise Manager Ops Center or outside of the software.

Creating an Oracle Solaris 10 Boot Environment From a Deployment Plan

You can only create a boot environment from a global zone in Oracle Solaris 10. The boot environments on Oracle Solaris non-global zones are tightly coupled with the corresponding boot environment on the global zone and cannot be managed independently.

You can use either of the following deployment plans to create an Oracle Solaris 10 boot environment:

- Boot Environment: A simple deployment plan that only creates an alternate boot environment. Use the default alternate boot environment profile and then override the applicable parameters to create a boot environment.
- Software Deployment / Update: A multi-step plan where one of the steps creates a boot environment. See the *Oracle Enterprise Manager Ops Center Updating Your Oracle Solaris 10 Operating System* for an example of how to use this plan.

Perform the following steps to create a boot environment using the Alternate Boot Environment profile and Boot Environment deployment plan:

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.

   Existing boot environments for the OS appear in the center pane.
3. To create a copy, or clone, of the existing active boot environment, click the **Create** icon.
4. Enter a unique name for the boot environment and provide a description. The text describes what the boot environment and from what source it was created. Edit the applicable parameters in the Alternate Boot Environment profile, then click **Confirm**.
5. When the job finishes, refresh the Boot Environment table to see the new boot environment.

Creating an Oracle Solaris 10 Boot Environment from an Operational Plan

An operational profile contains a single script and can define variables. The associated plan launches the profile. An operational plan is associated with a specific version of an operational profile. By default, creating a profile also creates an operational plan.

You can create the boot environments using this script from a few different contexts:

- Simple BE Operational profile
- Boot Environments profile using BE Operational profile from the plan context
- Create Boot Environment from the asset context

You can save a shell script on the Enterprise Controller and download it into the plan, or you can enter the script in a field when you create the plan. Both types of shell scripts are executed by the user. They differ in the location, either on the Enterprise

Controller or on the remote Agent Controller, and the user credentials needed to execute the profile. See Operational Plans and Profiles for more information about operational profiles and plans and the variables you can use.

1. Click **Plan Management** in the Navigation pane, then click **Operational Profiles.** The Operational Profiles appear in the center pane.
2. Click **Create Operational Profile**. The Create Profile - Operation page appears.
3. Enter a name for the new plan and a description of its purpose or role.
4. Select an asset type from the Subtype list.
5. Click **Next**.
6. Select the Operation Type from the drop-down menu, either **EC Shell** or **Remote Shell**.
   - If you select EC Shell, browse to the location of the script in the Script File field, then click **Load Script**.
   - If you select Remote Shell, enter your script in the Script field.
7. Enter a numeric value in the Timeout field, then select Minutes or Seconds. The default timeout is 60 minutes.
8. (Optional) Click **View System Variables** to view the default variables.
9. Click **Next**.
10. (Optional) Specify Additional Variables. You can specify any variable you want. For example, add Alarm_ID$ to add the incident identifier number for easier incident management.

Creating an Update Profile for Oracle Solaris 10 Boot Environment

You can use an Update profile to create and update a boot environment.

1. Expand **Plan Management**, then click **Update Profiles** in the Navigation pane.
2. Click **New Profile** in the Actions pane.
3. Enter a name and description for the profile.
4. Select **Script** from the Type drop-down menu.
5. Select an OS channel from the Distribution drop-down menu.
6. Select **Local** from the Category drop-down menu.
7. Choose the options that you want from the View and the Version drop-down menus.
8. Highlight Local RPMs in the Available Packages / Patches table, select the **Apply to all applicable distributions** check box.
9. Click **Profile Contents** and select an item, then click **Create OS Update Profile**.

Creating an Oracle Solaris 10 Boot Environment From an Update Profile

You can use an Update profile to create and update a boot environment.

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.

   Existing boot environments for the OS appear in the center pane.

3. To create a copy, or clone, of the existing active boot environment, click the **Create** icon.

   The plan for execution that uses the default Boot Environments profile for Oracle Solaris 10 appears.

4. Enter a unique name for the boot environment and provide a description. The text describes what the boot environment and from what source it was created. Edit the applicable parameters in the Alternate Boot Environment profile, then click **Confirm**.

5. When the job finishes, refresh the Boot Environment table to see the new boot environment.

Creating an Oracle Solaris 10 Boot Environment With an OS Update Job

Oracle Solaris Live Upgrade contains a suite of script commands. To create an alternate boot environment with Enterprise Manager Ops Center, use the `lucreate` command to write one or more Oracle Solaris Live Upgrade scripts, add the scripts to the Local Content library in Enterprise Manager Ops Center, then run an OS Update job and select the ABE options.

When you use Enterprise Manager Ops Center to create the alternate boot environment, the scripts must meet the following requirements:

- The script cannot contain parameters.
- The alternate boot environment name must be hard-coded into the script itself or otherwise be provided outside of Enterprise Manager Ops Center.
- The alternate boot environment name defined in the script must match the alternate boot environment name that you use when you run the update job to create the alternate boot environment.
- The script must return 0 on success and non-zero on failure.

For detailed instructions and examples for using the `lucreate` command to create a boot environment, see *Oracle Solaris 10 9/10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* available at: `http://docs.oracle.com/cd/E18752_01/html/821-1910`.

1. Expand **Libraries** in the Navigation pane.
2. Click **Local Content** in the Solaris/Linux OS Updates library.
3. Click **Upload Local Action** in the Actions pane.
4. Enter a name for the file.
5. Enter a brief description of the purpose of the action.
6. In the Action list, click the **Pre-action** type of action to run the script on the managed host before job tasks are carried out.
7. Click the name of the distribution that uses the action in the Distribution list. The Parent field shows the category, based on the type of Action.
8. Click **Browse** to locate and select the file.
9. Click **Upload** to upload the file to the selected distribution.

To Create a Boot Environment With an OS Update Job

1. Click the OS in the Assets section of the Navigation pane.

2. Click **New Update OS Job** in the Update section of the Actions pane. The New Update OS Job Wizard appears.
3. Complete the following Job Information parameters:
   - Enter a job name.
   - Select **Actual Run**, which creates the alternate boot environment when you specify in Step 5.
   - Select the Sequential task execution order.
   - Select the Target Setting: Use the same Targets for all tasks in the job.
   - Select a Task Failure Policy:
     - Complete as much of the job as possible
     - Stop at failure and notify
   - Select the **Boot Environment Type** check box.
   - Select the **Run ABE Pre-action Script** check box.
4. (Optional) To add tasks to the job, click the **Add Task** icon. To edit, click the Profile and Policy fields. Click **Next**.
5. Enter the name of the alternate boot environment, as defined in the script. Select a script, then click **Next**.
6. (Optional) Complete the Boot Environment Workflow, then click **Next**.
   - To synchronize the alternate boot environment with the current boot environment before mounting the alternate boot environment, select the **Sync ABE** check box.
   - To edit the description to describe the state of the Boot Environment, click **Modify Current BE**, and enter text in the Description field.
   - To edit the description to describe the state of the alternate boot environment, click **Modify Alternate BE**, and enter text in the Description field.
   - To switch boot environments after update, select the **Activate and Reboot ABE** check box.
7. Schedule the job, then click **Next**.
   - Run Now: Starts the job immediately after you click Finish in the Job Summary.
   - Start Date: Select a date and time to start the job.
   - On a recurring schedule: Enables you to run the same job on a monthly or daily scheduled time.
8. Review the Job Summary, then click **Finish** to run the job as scheduled in the previous step.

When the job completes, the new alternate boot environment is associated with the operating system. To verify that the alternate boot environment is created, click the operating system in the Assets pane. The Boot Environments tab appears in the center pane. Click the Boot Environments tab to display the new boot environment, as specified in the Live Upgrade script. An OS can have several associated alternate boot environments.

Note: The Boot Environments tab only appears when at least one alternate boot environment associated with the operating system.

Synchronizing Oracle Solaris 10 Boot Environments

You can synchronize, or sync, an active Oracle Solaris 10 global zone boot environment with an inactive boot environment on the same system. Synchronizing boot environments makes the inactive boot environment the same as the currently running boot environment. After you sync the boot environments, you can activate the inactive, or alternate, boot environment.

Activating a Boot Environment

Activating a boot environment switches a new boot environment to become the currently running boot environment when the system reboots.

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.

   Existing boot environments for the operating system appear in the center pane.
3. Click **Activate Boot Environment and Reboot** in the Actions pane.
4. To schedule the activation for a later date or time, select the check box. Click **Next**.
5. If you selected the schedule option, complete the schedule, then click **Next**.
6. Review the Summary, then click **Finish**.

   The new boot environment activates when you reboot the system.

Related Resources for Operating System Management

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` and the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm`.
- See Chapter 14, "Operating System Updates" for information about patching, or updating, your operating systems.
- See Chapter 18, "Oracle Solaris Zones" for information about zones and how you can use Oracle Enterprise Manager Ops Center to efficiently manage all phases of zones lifecycle.
- See Chapter 4, "Monitoring Rules and Policies" for information on how monitoring rules and policies work in the software.

For in-depth information about these products, see the following Oracle documentation:

- For a list of the Oracle Linux documentation available in HTML and PDF formats, visit the Oracle Linux Documentation website at `http://www.oracle.com/us/technologies/linux/index.html`.

- For a list of the Oracle Solaris 11 documentation available in HTML and PDF formats, visit the Oracle Solaris 11 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-11-192991.html.
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html.
- For a list of the Oracle Solaris 8 and 9 documentation, visit the Legacy Solaris Documentation website at http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html.
- For detailed instructions and examples for using the `lucreate` command to create a boot environment, see *Oracle Solaris 10 9/10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* available at: http://docs.oracle.com/cd/E18752_01/html/821-1910 and the Oracle Solaris 10 lucreate(1M) documentation at http://docs.oracle.com/cd/E19253-01/816-5166/lucreate-1m/index.html.
- For more information about JET resources and documentation, see *Solaris 10 10/09 Installation Guide: Custom JumpStart and Advanced Installations* available at http://docs.oracle.com/cd/E18752_01/html/821-1911/index.html.
- For JET documentation and to download additional modules, see http://www.oracle.com/technetwork/systems/jet-toolkit/index.html.

13

Operating System Provisioning

This chapter describes the operating system (OS) provisioning features that are available in Oracle Enterprise Manager Ops Center.

The following information is included:

See Chapter 12, "Operating System Management" for information about monitoring, OS Analytics, and managing boot environments. See Chapter 14, "Operating System Updates" for information about patching and updating your operating systems. See Chapter 11, "Hardware" for information about provisioning firmware.

Introduction to Operating System Provisioning

The provisioning feature provides a method of automatically and consistently installing operating systems on managed systems from the Oracle Enterprise Manager Ops Center UI.

You can provision the following:

- Oracle Solaris operating systems
- Linux operating systems

- Oracle VM Server for SPARC
- Logical Domains
- Oracle Solaris Clusters

This chapter focuses on basic OS provisioning. Many of the concepts apply to other types of provisioning. See Related Resources for Operating System Provisioning for links to information about provisioning Oracle VM Server for SPARC, logical domains, virtual machines, and Oracle Solaris Cluster.

Provisioning an operating system installs a specific operating system release with your defined configuration. Earlier versions of the software used a single OS Provisioning profile to define both the operating system and the configuration. Beginning with Oracle Enterprise Manager Ops Center 12.2.0.0.0, the single profile is replaced with an OS Provisioning profile and an OS Configuration profile.

Note: If you created OS Provisioning profiles in versions of the software earlier than 12*c* Release 2, see Migrating OS Provisioning Profiles to the New Format for how Oracle Enterprise Manager Ops Center updates your profiles and plans to the new format.

The following are needed to define your OS provisioning job:

1. **OS Provisioning profile**: Defines the image, provisioning, and installation requirements, including the basic OS configuration and boot network information.
2. **OS Configuration profile**: Defines the networking configuration. You can use a simple networking interface for any Oracle Solaris or Linux operating system, or advanced networking configurations for Oracle Solaris.

   When you create a configuration profile for Oracle Solaris, you can configure the following advanced networking options:

   – **Link aggregation**: Provides high availability and higher throughput by aggregating multiple interfaces at the MAC layer. Link aggregation enables you to combine the capacity of multiple full-duplex Ethernet links into a single logical link.

   – **IP multipathing (IPMP)**: Provides features such as higher availability at the IP layer. IPMP enables you to configure multiple IP interfaces into a single IPMP group.

   You can implement both Link Aggregation and IPMP methods on the same network because they work at different layers of the network stack.

   After you create the profiles, you create a deployment plan to apply the profiles. As part of applying the plan, you can change some of the options that you defined earlier in the profiles.
3. **Provision OS deployment plan**: Defines the OS Provisioning and OS Configuration profiles to use and the targets to provision. The plan also provides you with an opportunity to provide a specific IP address and to make changes to the network and interface for the target.

Figure 13–1 shows the basic steps that you need to plan for, and complete, a provisioning job:

Figure 13–1 Overview of OS Provisioning

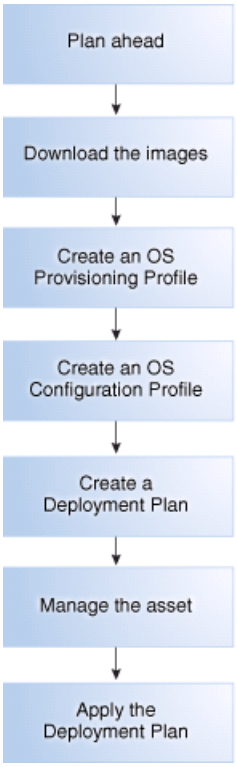

As shown in Figure 13–1, the following are the basic steps that you need to plan for, and complete, a provisioning job:

1. Plan ahead. See Planning for Operating System Provisioning for items to consider before provisioning.
2. Download a file with the OS image into the Software Library. See Images and Images for OS Provisioning for details on downloading images.
3. Create an OS Provisioning profile, edit an existing profile, or reuse an existing profile. See About Oracle Solaris OS Provisioning Profiles for an overview. Creating an Oracle Solaris 11 OS Provisioning Profile or Creating an Oracle Solaris 9 or 10 OS Provisioning Profile.
4. Create an OS Configuration profile, edit an existing profile, or reuse an existing OS Configuration profile. See Creating an Oracle Solaris 11 OS Configuration Profile or Creating an Oracle Solaris 9 or 10 OS Configuration Profile.
5. For provisioning OS on logical domains, you cannot use the OS Provisioning and OS Configuration profiles created for bare-metal provisioning. You must select the Logical Domain subtype when you create OS provisioning profiles for logical domains. See Provisioning an Operating System on Logical Domains for more information about OS provisioning on logical domains.
6. Create or configure a deployment plan that includes the OS Provisioning profile and the OS Configuration profile.

Note: To successfully provision an operating system, the plan must contain an OS Provisioning profile and an OS Configuration profile with the same platform, either SPARC or x86. For Oracle VM Server for SPARC, the control domain (CDom) version must be the same.

---

7. Manage the service processor for one or more systems that you want to provision.
   - To provision an existing system with a new operating system, verify that the service processor is discovered. See Adding Assets Using a Discovery Profile for how to discover a service processor.
   - To provision a bare metal system, manage the service processor. See Declaring Servers for OS Provisioning for details.
8. Apply the deployment plan on one or more targets. When you choose a group as a target, it must be a homogeneous group where all members are the same.

> **Note:** The target must have a discovered and managed service processor for Oracle Enterprise Manager Ops Center to identify the system as a target.

Default Profiles and Plans

Beginning with the 12.2.2.0 release, the software does not automatically generate a default set of Oracle Solaris 11 OS provisioning profiles and plans for each Support Repository Update (SRU) or OS image that you import. Default profiles and plans are still created for earlier versions of Oracle Solaris and for Linux images. This change significantly reduces the number of default profiles and plans that you need to manage. See the *Oracle Enterprise Manager Ops Center Administration Guide* for how to change the configuration to generate a set of default profiles and plans.

In releases prior to 12.2.2.0, when you import or download packages into a software library, the software creates a default OS Provisioning profile, OS Configuration profile, and deployment plan for the specific package. The profiles and plans appear in the Plan Management section of the UI.

All default profiles and plans have a naming convention that begins with *default* and includes the type of profile or plan and ISO image information. For example, the software creates the following for an Oracle Solaris 11.0 SPARC-10.1.0 Oracle Solaris Desktop package:

- OS Provisioning profile: `default-profile-Oracle Solaris 11.0 sparc-10.1.0-OracleSolarisDesktop v1`
- OS Configuration profile: `default-osc-profile-Oracle Solaris 11.0 sparc-10.1.0-OracleSolarisDesktop v1`
- Deployment Plan: `default-profile-Oracle Solaris 11.0 sparc-10.1.0-OracleSolarisDesktop-plan`

The default deployment plan references the associated default profiles for the package. You can edit the default profiles and plans, you can create copies of the default profiles and plans and edit them, or you can create your own profiles and plans. See Plans and Profiles for an overview of deployment and operational plans and version control.

Deployment Plans

A deployment plan defines the OS Provisioning and OS Configuration profiles to apply, the managed targets to provision, the network configuration and IP addresses for the targets, and the tasks to perform. Several different deployment plans let you provision an operating system. The Provision OS plan is a simple plan with the sole task of provisioning the operating system. The Provision OS plan lets you install the OS on a single system, one or more groups of systems, or a combination of systems that are attached to your network. Other plans are multi-step or complex plans where

OS provisioning is one of several tasks performed. See About Deployment Plans That Provision an Operating System for more information.

In some cases, the requirements are determined by the target type that you are provisioning. Each target type has different requirements and options. See Planning for Operating System Provisioning for a list of requirements.

Note: To provision Oracle Solaris 11 or to manage Oracle Solaris 11 boot environments, both the Enterprise Controller and Proxy Controller must be installed on a system that is running Oracle Solaris 11.

To provision Oracle Solaris 9 and 10, you can use a Proxy Controller that is installed on a system that is running either Oracle Solaris or Linux.

To provision Oracle Solaris 10 using JET customization, the Enterprise Controller must be installed on a system that is running an Oracle Solaris operating system.

Roles for Operating System Provisioning

Table 13–1 lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 13–1 OS Provisioning Tasks and Roles

Task	Role
Import or upload image	Storage Admin
Create OS Provisioning profile	Plan/Profile Admin
	Asset Admin
	Update Admin
Create OS Configuration profile	Plan/Profile Admin
	Asset Admin
	Update Admin
Create OS Provisioning Deployment Plan	Plan/Profile Admin
	Asset Admin
	Update Admin
Apply, or deploy a Deployment Plan	Apply Deployment Plans

Actions for Operating System Provisioning

OS Provisioning enables you to deploy an Oracle Solaris or Linux operating system on a managed system or service processor.

The following actions are available for OS provisioning:

- Modify or create an OS Provisioning profile to define the OS platform, image, configuration, file system, naming service, and other installation parameters.

- Modify or create an OS Configuration profile to define the OS platform, OS management, and network interface configuration.
- Modify or create a Deployment Plan to define which OS Provisioning profile and OS Configuration profile to deploy, to identify the target systems, and to begin a provisioning job.

Location of Operating System Provisioning Information in the UI

Operating system information appears in the **Assets** section in the Navigation pane. The operating system appears beneath the system in the asset tree. Click the OS in the **Asset** section to display information.

OS Provisioning profiles, OS Configuration profiles, and Deployment Plans appear in the **Plan Management** section in the Navigation pane. Click a profile or plan to view more details.

Planning for Operating System Provisioning

The following are some of the items to consider before provisioning:

- Do you need the Enterprise Controller and Proxy Controller installed on Oracle Solaris 11?

  See Enterprise and Proxy Controller Requirements for OS Provisioning.
- Do you want to use WAN boot or Dynamic Host Configuration Protocol (DHCP) services to support OS provisioning operations?

  See Using WAN Boot for Oracle Solaris Operating Systems and Using Dynamic Host Configuration Protocol (DHCP).
- Do you want to use advanced networking, either IPMP or Link Aggregation, for Oracle Solaris?

  See Defining IPMP in an OS Configuration Profile and Defining Link Aggregation in an OS Configuration Profile.
- Do you have OS images available in the library?

  See Adding Images to Local Software Libraries.
- Do you have the latest firmware installed on the system?

  See Using the Latest Firmware Version.
- Do you want to create custom scripts to add to the provisioning job?

  You can create a script to perform a task that is not defined in the OS provisioning or OS Configuration profiles and include it in the OS provisioning job. For example, you might want to change permission levels or disable print capabilities. See Creating Custom Scripts.
- Do you want to install an Agent Controller on the new operating system for full management capabilities?

  See Determining Agent Management Mode.
- Do you have networks and IP addresses available for provisioning?

Review the following information to plan for OS provisioning:

- Enterprise and Proxy Controller Requirements for OS Provisioning
- Networking for OS Provisioning

- Using WAN Boot for Oracle Solaris Operating Systems
- Using Dynamic Host Configuration Protocol (DHCP)
- Determining the Network Interface to Use
- Provisioning an OS Using a User-Defined MAC Address
- Defining IPMP in an OS Configuration Profile
- Defining Link Aggregation in an OS Configuration Profile
- Adding Images to Local Software Libraries
- About NVRAC When Provisioning an OS on a SPARC Platform
- Creating Custom Scripts
- Determining Agent Management Mode

Enterprise and Proxy Controller Requirements for OS Provisioning

The operating system that the Enterprise Controller and Proxy Controller are installed on might impact your ability to provision an operating system or Oracle VM Server for SPARC. If you ever plan on provisioning, patching, or managing Oracle Solaris 11, install the Enterprise Controller and Proxy Controller on systems that are running the Oracle Solaris 11 operating system.

Table 13–2 shows the actions that you can perform based upon which operating system that the Enterprise Controller and Proxy Controller are installed.

Table 13–2 Provisioning Actions Determined by the Operating System on which the Enterprise Controller and Proxy Controller are Installed

Action	Enterprise and Proxy Controllers on Oracle Solaris 11	Enterprise and Proxy Controllers on Oracle Solaris 10	Enterprise and Proxy Controllers on Linux
OS Provisioning Oracle Solaris 11	Supported	Not Supported	Not Supported
OS Provisioning Oracle Solaris 11 with JET and DHCP server	JET: Not Supported Oracle DHCP server: Not Supported ISC DHCP server: Supported	JET: Not Supported Oracle DHCP server: Not Supported	JET: Not Supported Oracle DHCP server: Not Supported
OS Provisioning Oracle Solaris 10	JET Supported Oracle DHCP server: Not Supported	JET Supported Oracle DHCP server: Supported	JET: Not Supported Oracle DHCP server: Not Supported
OS Provisioning Oracle Solaris 10 with JET and DHCP Server	JET: Supported Oracle DHCP server: Not Supported	JET: Supported Oracle DHCP server: Supported	JET: Not Supported Oracle DHCP server: Not Supported
OS Provisioning Linux	Supported	Supported	Supported
Provisioning Oracle VM Server for SPARC	Supported	Supported	Enterprise Controller: Supported Proxy Controller: Not Supported
Provisioning Oracle VM Server for x86	Supported	Supported	Supported

Note: To provision Oracle Solaris 10 using JET customization, the Enterprise Controller must be installed on a system that is running an Oracle Solaris operating system.

The specific hardware, OS, and firmware requirements that you need for provisioning Oracle VM Server for SPARC are described in Chapter 19, "Oracle VM Server for SPARC".

Networking for OS Provisioning

The target system boots over the network and gets its network configuration and the location of the install server from a DHCP server or WAN boot.

When provisioning an operating system, the Proxy Controller must be attached to the same subnet as the assets that you want to provision. You can use DHCP or WAN boot (Oracle Solaris only). To improve security and bandwidth, consider establishing a provisioning network for OS deployment and a production network for guest management.

When you install a system from the network, you must provide a method of determining the network configuration (IP address and gateway), which server is going to perform the boot and install, and the installation instructions.

Oracle Enterprise Manager Ops Center uses DHCP services or WAN boot to support OS provisioning operations. DHCP servers enable you to obtain the IP configuration and the rest of the information needed on the NIC. You must configure DHCP services on the Proxy Controller on the same subnet as the target systems to support OS provisioning. Configure the DHCP services in the Oracle Enterprise Manager Ops Center user interface, not from the command line. Oracle Solaris 10 uses the Oracle DHCP server with a Proxy Controller that is running Oracle Solaris 10. Oracle Solaris 11 uses an ISC DHCP server.

WAN boot enables you to provision Oracle Solaris 10 or 11 on a SPARC platform across the network. With WAN boot, the software explicitly configures the information in the Open Boot PROM (OBP) and uses WAN boot for installation.

WANBoot has a number of benefits over broadcast-based installation:

- Not restricted to a single subnet
- Does not require special DHCP configuration or DHCP helpers
- Uses standard HTTP and HTTPS protocols, which cross firewalls much more easily than NFS-based package installations.

For more information, see Using WAN Boot for Oracle Solaris Operating Systems and Using Dynamic Host Configuration Protocol (DHCP).

Using WAN Boot for Oracle Solaris Operating Systems

The following information is in this section:

Overview of WAN Boot

The WAN boot installation method enables you to boot and install software over a wide area network (WAN) by using HTTP. By using WAN boot, you can install the Solaris OS on SPARC based systems over a large public network where the network infrastructure might be untrustworthy. You can use WAN boot with security features to protect data confidentiality and installation image integrity.

WAN boot is the default connection for Oracle Solaris 11 provisioning. Oracle Solaris 11 provisioning does not use a Flash Archive (FLAR) image.

Oracle Solaris 10 provisioning can use a WAN boot or DHCP connection. For Oracle Solaris 10, you need a FLAR to use WAN boot. When you do not use the FLAR, you must enable DHCP before you can provision.

With a WAN boot connection, Oracle Solaris 10 provisioning enables you to provision a FLAR image on a SPARC system using an HTTP web server. WAN boot installation is useful when DHCP does not meet your organization's security policies or you have SPARC-based systems that are located in geographically remote areas and you need to install servers or clients that are accessible only over a public network. Because WAN boot uses an HTTP server, it works across your corporate firewall and does not require DHCP or a JumpStart boot server to be on the same network as the client systems.

> **Note:** When provisioning Oracle Solaris with WAN boot, the boot interface only uses untagged networks. Since the OpenBoot PROM (OBP) is not aware of VLAN tagging and cannot handle VLAN-tagged network packets, do not use a tagged network as the boot network.

The WAN boot installation method uses port 5555 and HTTP to boot and install software on SPARC-based ILOM, ALOM, or M-series systems over a wide area network (WAN). For Oracle Solaris 11, you can edit the SMF service to change the default port from 5555 to another port. See the *Oracle Enterprise Manager Ops Center Administration Guide* for how to reconfigure the default WAN boot port.

The WAN boot security features protect data confidentiality and installation image integrity over a large public network where the network infrastructure might be untrustworthy. You can use private keys to authenticate and encrypt data. You can also transmit your installation data and files over a secure HTTP connection by configuring your systems to use digital certificates. For more information about secure WAN boot installation configuration, see the *Security Configurations Supported by WAN Boot* section of the *Oracle Solaris 10 10/09 Installation Guide: Network-Based Installations* document at `http://docs.oracle.com/cd/E19253-01/821-0439/wanboottasks2-30/index.html`.

> **Note:** WANBoot is not available on older hardware.

Requirements for a WAN Boot Connection

The following are required to use WAN boot with Oracle Enterprise Manager Ops Center:

- Oracle Solaris 11
  - The target is a SPARC ALOM-CMT, ILOM-SPARC or M-Series platform that has a supported OBP or XCP. See Checking OBP Support for WAN Boot on the Client for how to check your client OBP. For M-Series, the XSCF Control Package (XCP) file should be at least version 1082. The XCP file contains the

hardware's control programs and includes the XSCF firmware and the OpenBoot PROM firmware.

- The Enterprise Controller is installed on an Oracle Solaris operating system. You can use a SPARC or x86 platform for the Enterprise Controller.
- WAN boot is enabled for Oracle Solaris 11 in Administration. See Disabling and Enabling WAN Boot.

Note: When the target does not have the required OBP firmware version, the Oracle Solaris 11 provisioning profiles do not appear in the target list for the server.

- Oracle Solaris 10
  - The target is a SPARC (ALOM-CMT, ILOM-SPARC or M-Series) and has the minimum OBP or XCP version
  - Use a FLAR image. WAN boot is only supported in Oracle Solaris 10 if you use flash archives. ISO images require DHCP.
  - The Enterprise Controller is installed on an Oracle Solaris operating system. You can use a SPARC or x86 platform for the Enterprise Controller.
  - WAN boot is enabled for Oracle Solaris 10 in Administration. See Disabling and Enabling WAN Boot.
  - The target has the required OBP firmware installed. To determine if your client system has a WAN boot-enabled PROM, see Checking OBP Support for WAN Boot on the Client.

Note: When the target does not have the required OBP firmware version, Oracle Solaris 10 provisioning reverts to a DHCP connection, or OBP/XCP versions + if an ISO is used.

- Verify that the `/opt/SUNWjet/etc/server*interface*` file on the Proxy Controller is updated to use the Proxy IP to target the network.

Checking OBP Support for WAN Boot on the Client

To determine if your client system has a WAN boot-enabled PROM, perform the following to check the client Open Boot PROM (OBP) for WAN boot support.

1. Log in to a terminal window as root.
2. Enter the following to check the OBP configuration variables for WAN boot support:

   ```
   # eeprom | grep network-boot-arguments
   ```

3. The OBP supports WAN boot installations when the variable *network-boot-arguments* appears, or when the command returns the output *network-boot-arguments: data not available*. For example:

   ```
   # eeprom | grep network-boot-arguments
   network-boot-arguments: data not available
   ```

4. If the command in Step 2 does not return any output, the OBP does not support WAN boot installations. Use Firmware Provisioning to update the OBP to the required level.

Setting Up a WAN Boot Connection

When Oracle Enterprise Manager Ops Center is installed on an Oracle Solaris operating system, the Enterprise Controller is automatically configured to be a WAN boot server.

Oracle Solaris 11 uses WAN boot. For earlier versions of Oracle Solaris, WAN boot is the default connection for provisioning when the requirements are met and you choose to use a FLAR image. When you launch an OS provisioning on an eligible SPARC-based system and you choose a FLAR image, the software automatically uses WAN boot. If you have a group of systems to provision, the software determines whether to use WAN boot or DHCP for each system.

Disabling and Enabling WAN Boot

WAN boot is automatically installed and enabled when the Enterprise Controller is running on an Oracle Solaris operating system. You can disable or enable WAN boot in the Enterprise Controller configuration file.

To Disable and Enable WAN Boot

1. Expand the **Administration** section in the Navigation pane, then click **Enterprise Controller**.
2. Click the **Configuration** tab.
3. Select **OS Provisioning** from the Subsystem menu.

Figure 13–2 Enterprise Controller's Configuration Tab

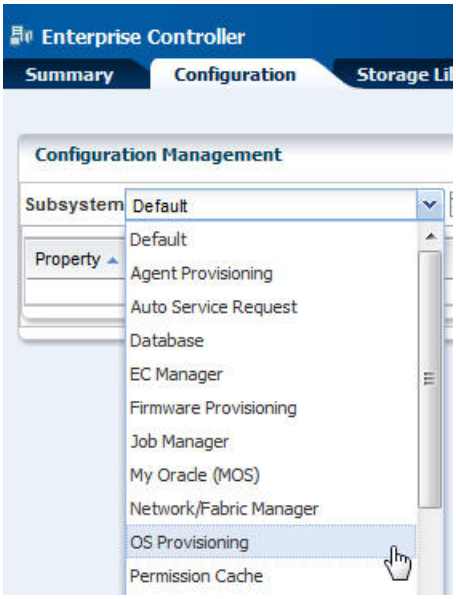

4. Scroll down to the WAN boot property:
   - For Oracle Solaris 11, see the following property: `usesS11WANBoot`.
   - For Oracle Solaris 10, see the following property: `usesS10WANBoot`.

When *true* appears in the value column, WAN boot is enabled, as shown in Figure 13–3.

Figure 13–3 WAN Boot Configuration

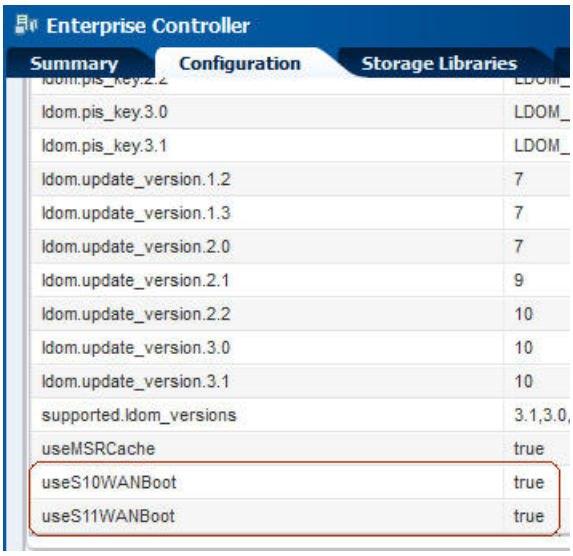

5. (Optional) To disable WAN boot, change the value for the property to **false**.
6. (Optional) To enable WAN boot, change the value for the property to **true**.

Using Dynamic Host Configuration Protocol (DHCP)

DHCP dynamically assigns IP addresses to devices on a network. A typical OS provisioning job requires an installation server and a DHCP server on the same subnet as that of the client systems. A JumpStart boot server must be on the same subnet as that of the client systems.

Before you can provision, you must configure DHCP services on the Proxy Controllers. You can use basic DHCP services, with or without defined subnets, or an external DHCP server. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about how to configure DHCP and subnets for OS provisioning.

Note: Oracle Solaris 10 supports an Oracle Solaris DHCP server. The external DHCP-related files are copied only if the Proxy Controller is running on an Oracle Solaris 10 operating system.

Oracle Solaris 11 only supports an ISC DHCP server.

Verify that the Dynamic Host Configuration Protocol (DHCP) services are enabled on Proxy Controllers. You cannot create a profile or assign any network if the DHCP services are not enabled. The Install Server option to provision an OS on a server is not enabled if the DHCP is not enabled on any of the interfaces.

Determining the Network Interface to Use

The OS Configuration profile lets you define all network interfaces you want to use on the operating system. As part of the OS Configuration profile for Oracle Solaris, you have the option to establish Link Aggregation or IPMP network interfaces that the target system will use after the operating system is configured.

The OS Configuration profile lets you define all network interfaces you want to use on the operating system. When you use an on board interface for the provisioning network, you can pair the network with option card interfaces for Link Aggregation. Before you provision, you must know your network architecture. For example, the PCIe slot and Net*n* connection. Figure 13–4 is an example of the PCIe slots on a T4-4 server.

Figure 13–4 PCIe Slots on a T4-4 Server

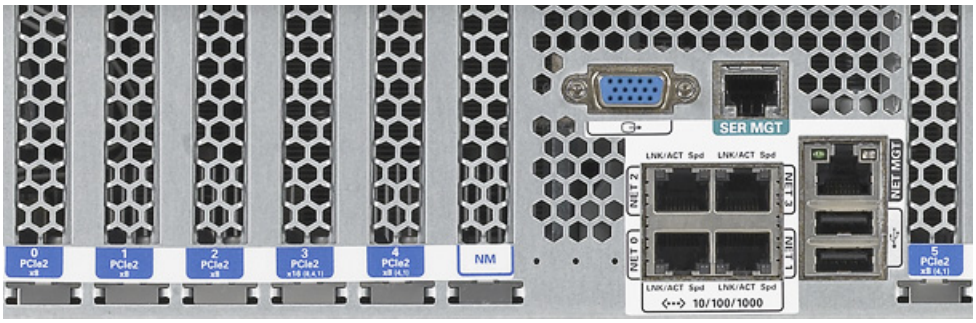

You can use a built-in network interface or a network that is connected to a specific port on a network interface card (NIC). This information is not available from the Oracle Enterprise Manager Ops Center UI. You must contact your network administrator for the details.

The OS Configuration profile lets you define one or more interfaces. When you specify the network interfaces, you select an interface for the controller from a list of 32 interfaces. The 32 interfaces that appear in the wizard are all possible network interfaces, not available interfaces. Your network administrator can give you the list of available networks and interfaces. As shown in Figure 13–5, the primary interface is `net_0` and is the boot interface in the OS Configuration profile. Always define the first interface that appears in the table as the boot interface. You can change the primary interface to a different network when you apply the plan.

Figure 13–5 Specify Network Interfaces in the OS Configuration Profile

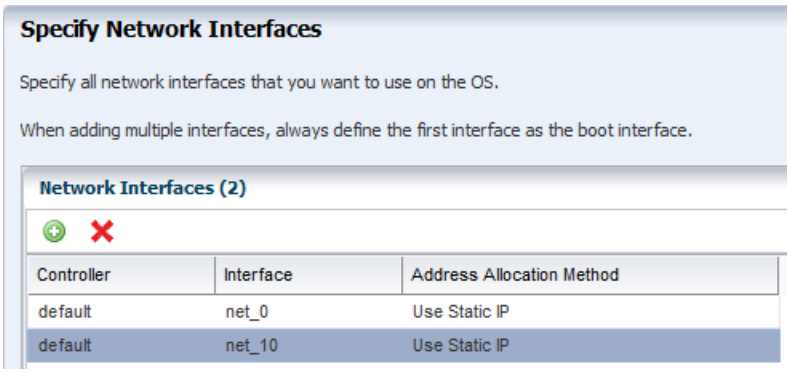

The Boot Interface Resources Assignments page in the OS Configuration profile lets you provide the network resources and host name for each target. By default, the first network listed is the IP address and host name for the primary boot interface. If you do not enter a host name, your DNS server provides the name.

As shown in Figure 13–6, you specify the network resources for the boot interface of each target when you apply the plan, including assigning the network and the IP address for each target. Instead of using the network interface (NIC) to perform an OS

provisioning job, you can provide a MAC address for the service processor. When you provide the MAC address, the DNS server provides the host name.

Figure 13–6 Assign Boot Interface Resources in the Deployment Plan

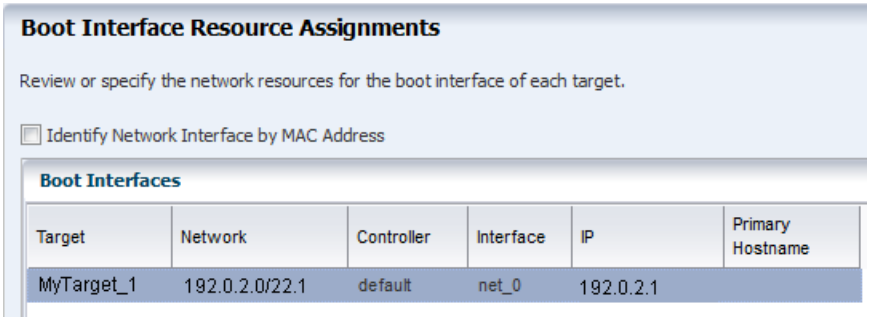

Note: When you apply the plan, the OS Configuration provides a list of available tagged and untagged networks for the boot interface. However, OS provisioning cannot boot from a tagged network and the networks will only be configured in untagged mode.

You can change the networking when you apply a plan that includes OS provisioning, including changing the primary interface to a different interface for a specific target. The flexibility in defining networking is useful when you want to perform OS provisioning and boot on a backup or provisioning network, but you need the host name to match the primary interface. The first listed interface in the OS Configuration profile is the primary interface and is the interface to use when setting the system host name. You can set a second network interface to be the boot interface.

As shown in Figure 13–7, the deployment plan gives you the opportunity to define which network is the primary network when you have multiple network interfaces.

Figure 13–7 Network Resource Assignments in the Deployment Plan

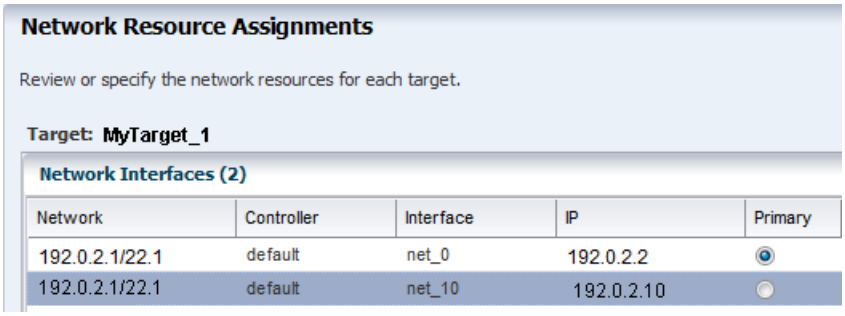

Refreshing the Oracle Solaris 11 Service

Oracle Enterprise Manager Ops Center creates an Oracle Solaris 11 `installadm` Automated Installer service when you first configure the Proxy Controller. If the service is not created during configuration, the software creates the service when you run the first Oracle Solaris 11 OS provisioning job. The service creates and updates the Oracle Solaris 11 Image Packaging System (IPS), which contains the packages that you need to install, provision, and update your Oracle Solaris 11 operating system.

The Oracle Solaris 11 `installadm` service creates and adds the existing network interfaces in the `/var/ai/ai-webserver/listen-addresses.conf`. When you add a

new network interface, you must refresh the `installadm` service to enable Oracle Solaris 11 AI service access on that interface.

Note: When you add a new network interface, run the `svcadm refresh system/install/server` command to refresh the service to enable Oracle Solaris 11 AI service access on that interface. Use the `installadm list` and the other options for `installadm` to check the status. See the *Oracle Enterprise Manager Ops Center Command Line Interface Guide* for more details.

You cannot use new interfaces to provision or update Oracle Solaris 11 until you refresh the service.

Provisioning an OS Using a User-Defined MAC Address

Instead of using the IP address and NIC to perform an OS provisioning job, you can provide a MAC address for the service processor.

To view the MAC address, expand **Assets** in the Navigation pane, then select the service processor. The MAC address appears on the right side of the Summary section of the Dashboard tab, as shown in Figure 13–8.

Figure 13–8 Dashboard Page Showing MAC Address

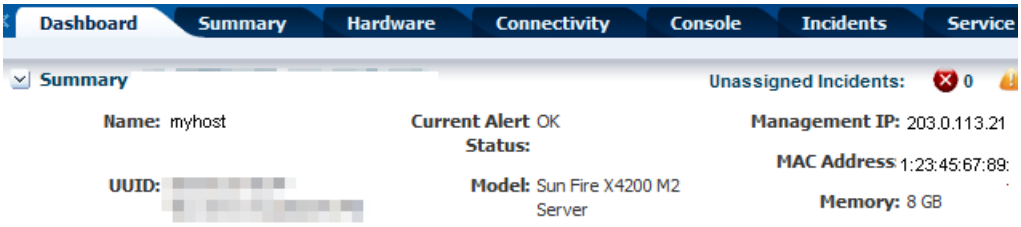

Special OS Provisioning and OS Configuration profiles are not required to use a MAC address for provisioning an operating system When you apply a plan that includes the OS Provisioning and OS Configuration profiles, you step through the plan to verify the configuration and provide final information before starting the job.

The Boot Interface Resources Assignments page lets you provide the network resources and host name for each target. By default, the first network listed is the IP address and host name for the primary boot interface. Alternatively, you can choose to provide the MAC address. Click **Identify Network Interface by MAC Address** to display the MAC Address field, as shown in Figure 13–9. Enter the MAC Address and the IP Address. When you provide the MAC address, the DNS server provides the host name.

Figure 13–9 Boot Interface Resource Assignments Using MAC Address

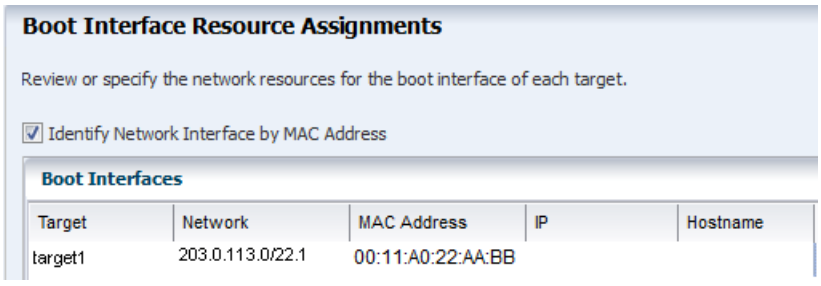

Defining IPMP in an OS Configuration Profile

IP multipathing (IPMP) groups provide network failover for your Oracle Solaris operating system, Oracle VM Server for SPARC system, and guests. Use IPMP to improve overall network performance by automatically spreading out outbound network traffic across the set of interfaces in the IPMP group.

You can configure one or more physical interfaces into an IPMP group. After configuring the IPMP group, the system monitors the interfaces in the IPMP group for failure. If an interface in the group fails or is removed for maintenance, IPMP migrates, or fails over, the failed interface's IP addresses. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections. The network access changes from the failed interface to the standby interface in the IPMP group and the data address of the failed interface migrates to the standby interface. See IP Multipathing Groups and Creating IPMP Groups for more information about IPMP groups.

Defining Link Aggregation in an OS Configuration Profile

Link aggregation, as defined in the IEEE802.3ad standard, is an Oracle Solaris feature that enables you to pool several datalink resources into a single logical link to improve network performance and availability. Figure 13–10 is an example of a link aggregation configured on a system. The aggregation, aggr1, has three underlying datalinks, net0, net1, and net2. The datalinks are dedicated to serving the traffic that traverses the system through the aggregation.

Figure 13–10 Link Aggregation

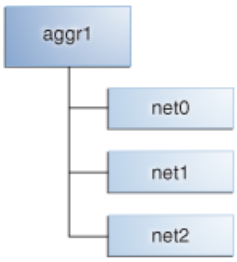

In an aggregated link, two or more NICs form a group and all members of the link aggregation provide network access at the same time. In addition to the high availability and load balancing that an IPMP group provides, an aggregated link can provide increased throughput when the network ports are also aggregated.

Link aggregation has the following features:

- Increased bandwidth: The capacity of multiple links is combined into one logical link.
- Automatic failover and failback: By supporting link-based failure detection, traffic from a failed link is failed over to other working links in the aggregation.
- Improved administration: All underlying links are administered as a single unit.
- Less drain on the network address pool: The entire aggregation can be assigned one IP address.
- Link protection: You can configure the datalink property that enables link protection for packets flowing through the aggregation.
- Resource management: Datalink properties for network resources as well as flow definitions enable you to regulate applications' use of network resources.

When you create an OS Configuration profile, link aggregation is a networking option that is available for Oracle Solaris and Oracle VM Server for SPARC. To define link aggregation networking, you must define a load balancing policy and a MAC address policy.

Aggregated interfaces are treated as a single network interface. Oracle Enterprise Manager Ops Center includes any link aggregations in the list of available NICs, as if the link aggregation were an individual interface. To assign a network with a link aggregation to an Oracle VM Server or global zone, select the link aggregation from the NIC list. You can view the link aggregation details on the Oracle VM Server's or global zone's Network tab. See Link Aggregation in Chapter 17, "Networks for Virtualization", Managing Global Zone Networks in Chapter 18, "Oracle Solaris Zones", or in the tasks in Chapter 19, "Oracle VM Server for SPARC" for more information.

Load Balancing Policy

- L2: Determines the outgoing link by hashing the MAC (L2) header of each packet
- L3: Determines the outgoing link by hashing the IP (L3) header of each packet
- L4: Determines the outgoing link by hashing the TCP, UDP, or other ULP (L4) header of each packet

Link Aggregation Control Domain (LACP)

If the aggregation topology involves a connection through a switch, determine whether the switch supports LACP. When the switch supports LACP, you must configure LACP for the switch and the aggregation.

- LACP Mode: Select **No** when the switch does not support LACP. When the aggregation topology involves a connection through a switch that supports LACP, configure LACP for the switch and the aggregation and define whether LACP runs in Active or Passive mode.
- LACP Timer: Define the timer as either Short or Long.

MAC Address Policy

- Auto: Use MAC address of any network interfaces in the Link Aggregation
- Fixed: Use MAC address of a specific network interface. Select the network interface to use in the next step.

Link aggregations perform similar functions as IPMP to improve network performance and availability.

When interfaces are aggregated, they are treated as a single network interface. Oracle Enterprise Manager Ops Center displays the link aggregation in the list of available NICs as if it were an individual interface. You can assign a network with a link aggregation to a non-global zone, and select the link aggregation from the NIC list.

Adding Images to Local Software Libraries

The images and supporting metadata that you use to provision and update operating systems are stored in software libraries, as shown in Figure 13–11.

Figure 13–11 Software Libraries

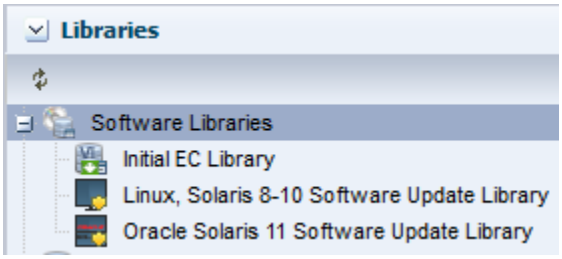

The software libraries shown in Figure 13–11 are created when you install Oracle Enterprise Manager Ops Center:

- **Oracle Solaris 11 Software Library**: Acts as a local copy of the Oracle Solaris 11 Image Packaging System (IPS) repository. This library contains the packages to install, provision, and update Oracle Solaris 11 operating systems.
- **Linux, Oracle Solaris 8-10 Software Library**: Contains Knowledge Base metadata, operating system package and patch content for Linux and Oracle Solaris and operating systems.
- **Initial EC Library**: Stores the operating system (and firmware) images that you download.

You can use the following methods to add images:

- **Upload ISO Image**: Copies the ISO image from a system's web browser to the library.
- **Import Image**: Copies the ISO or FLAR from a file system location on the Enterprise Controller system to the library.
- **Download OS Image**: Downloads the OS image from My Oracle Support to the library.

See Chapter 5, "Software Libraries" for more information about software libraries and adding images to the library.

Using the Latest Firmware Version

Verify that you have the latest firmware version before provisioning an operating system, including Oracle VM Server for SPARC and bare metal provisioning.

About NVRAC When Provisioning an OS on a SPARC Platform

When you run an OS provisioning job on a SPARC machine, Oracle Enterprise Manager Ops Center resets the configuration to the factory default configuration and removes the user-defined commands that are executed during start-up and that are stored in the NVRAMRC file in the non-volatile RAM (NVRAM). The Control Domain OS Provisioning profile does give you the option to preserve the information in the NVRAMRC file. See Chapter 19, "Oracle VM Server for SPARC" for more information about provisioning a Control Domain and Oracle VM Server for SPARC.

Creating Custom Scripts

You can create a script and reference the script in the OS Provisioning profile. When the script is saved in a directory that the Enterprise Controller can access, Oracle Enterprise Manager Ops Center deploys the script as part of the provisioning job. You

can save scripts in a local directory of the Enterprise Controller, or in a directory that the Enterprise Controller mounts using NFS.

You cannot use custom scripts when provisioning Oracle Solaris 11.

Determining Agent Management Mode

You can manage an operating system in one of two modes: agent managed or agentless managed. The management mode determines the features that are enabled for your operating system.

When you choose the agent managed mode, you can perform software updates and create operating system reports. When you choose agentlessly managed, SSH credentials are required to monitor the operating system. You can change the management mode after the OS is provisioned.

Agent managed is the more robust management mode because the Agent Controller enables a greater level of communication with the Proxy Controller and Enterprise Controller than the agentless managed operating systems. You can use the features and perform the actions described in this chapter with an agentless managed operating system, but OS update functionality requires an agent managed operating system. You can manage your operating systems by installing an Agent Controller on the OS or by using SSH to perform tasks. See Using Agent Management for Operating Systems for more details on managing operating systems with or without an Agent Controller. See Virtualization Agent Controllers for more information about Agent Controllers specifically used for virtualization.

About OS Provisioning Profiles

To complete provisioning, you must have an OS Provisioning profile and an OS Configuration profile. OS Provisioning profiles define the provisioning and installation details.

The following information is covered in this section:

- About Oracle Solaris OS Provisioning Profiles
- About Linux Provisioning Profiles

See Default Profiles and Plans for information about default OS Provisioning profiles. When you download an OS image, the job saves the image in the Initial EC Software Library. You can create an OS Provisioning Profile from that image:

- To create an OS Provisioning profile for Linux and Oracle Solaris versions earlier than Oracle Solaris 11, add the images into the Initial EC Library.
- To create an OS Provisioning profile for Oracle Solaris 11, add the images to the Oracle Solaris 11 Software Library.

OS Provisioning profiles are in the Plan Management section of the UI. To view the profiles, expand Plan Management, scroll down to Profiles and Policies, then open the OS Provisioning folder. You can create new profiles or copy and edit existing profiles in Plan Management.

You can create one or more new libraries to organize and save the images to a location other than the Initial EC Software Library. You can save the images in a local software library on the Enterprise Controller or in a Network Attached Storage (NAS) software library that you create on an NFS server that the Enterprise Controller can access. See Chapter 5, "Software Libraries" for more information about software libraries and importing images.

After you download the OS images, you can copy or create new OS Provisioning profiles. You can reuse the profiles in a variety of plans that have OS provisioning as a step.

MPxIO is highly desirable on SPARC, whether for a standalone Global Zone or an Oracle VM Server for SPARC Control Domain. MPxIO is enabled on the default OS Provisioning profiles.

See Defining Link Aggregation in an OS Configuration Profile and Defining IPMP in an OS Configuration Profile for information about Link Aggregation and IPMP networking options.

About Oracle Solaris OS Provisioning Profiles

For each OS Provisioning profile, you specify the OS image, OS setup parameters, user account details, ISCSI disk usage, file system parameters, and the naming service for the operating system.

When using an OS Provisioning profile to install Oracle Solaris 11, the profile always installs the latest version of the Oracle Solaris 11 operating system.

When you specify the naming service in the OS Provisioning profile, you must ensure that you enter the correct information in each of the fields. IP addresses 0.0.0.0 or 255.255.255.255 are allowed. Enter each IP address in a new row in the Name Server field. In the Domain Name Search List field, enter each domain name, such as `1domain.com` and `2domain.com`, on a new line.

See Default Profiles and Plans for information about default OS Provisioning profiles and plans. When you have the software configured to automatically create a default profiles, each profile is defined by the OS image that is in the Software Library. When you import an OS image into the software library, the software creates a default OS Provisioning profile. The profile name is the OS description prepended with the term *default-profile*. For example, `default-profile-Oracle Solaris 11.0 sparc 10.1.0-OracleSolarisDesktop`. You can edit the default profile, or you can copy the default profile to create a new profile that you can edit to define your parameters.

JumpStart Enterprise Toolkit for Oracle Solaris 9 and 10

For Oracle Solaris 9 and 10 only, you can optionally use JumpStart Enterprise Toolkit (JET) modules to specify additional Installation Parameters. Oracle Solaris 11 uses the Automated Installer (AI) instead of JET.

Within the Oracle Enterprise Manager Ops Center UI, there are 2 methods of influencing the JET template variables:

- Import a JET Template
- Add JET variables to the OS provisioning profile

You cannot manipulate the JET template in the UI. When you want to make changes to a template, make the changes and then import the template.

To learn more about JET and how to use JET in a profile, see About JumpStart Enterprise Toolkit (JET) for Oracle Solaris.

About Linux Provisioning Profiles

Each profile is defined by the OS image that is in the Software Library.

When you have the software configured to automatically create a default profiles, each profile is defined by the OS image that is in the Software Library. When you import an

OS image into the software library, the software creates a default OS Provisioning profile. The profile name is the OS description prepended with the term *default-profile*. You can edit the default profile, or you can copy the default profile to create a new profile that you can edit to define your parameters. See Default Profiles and Plans for information about default OS Provisioning profiles and plans.

When you specify the naming service in the OS Provisioning profile, ensure that you enter the correct information in each of the fields. IP addresses 0.0.0.0 or 255.255.255.255 are allowed. Enter each IP address in a new row in the Name Server field. In the Domain Name Search List field, enter each domain name, such as `1domain.com` and `2domain.com`, on a new line.

About OS Configuration Profiles

OS Configuration profiles define the operating system, network configuration details, host name, and server pool configuration.

The OS Configuration profile enables you to specify and assign the following network resources:

- Controller
- Interface
- Address Allocation Method
- Network
- IP address

A server pool is a group of one or more virtualization hosts with the same processor architecture that have access to the same virtual and physical networks, and storage resources. Server pools provide load balancing, high availability capabilities, and sharing of some resources for all members of the pool. Once created, you can edit the server pool settings.

You can create server pools for Oracle VM Server (for SPARC and x86) and for Oracle Solaris Zones. When you want the server to be added to a server pool, you can configure the OS Configuration profile to assign the newly provisioned server to a compatible server pool or you can create a new server pool based on the attributes of the newly provisioned server and assign default server pool settings. See Chapter 21, "Server Pools".

Two advanced network interface options are available for Oracle Solaris and Oracle VM Server for SPARC systems:

- **Link Aggregation**: Provides high availability and higher throughput by aggregating multiple interfaces at the MAC layer.
- **IP Multipathing (IPMP)**: Provides features such as higher availability at the IP layer.

You can implement both methods on the same network because they work at different layers of the network stack.

Defining Link Aggregation in an OS Configuration Profile

Link aggregation groups two or more NICs. All members of the link aggregation provide network access at the same time and are treated as a single network interface. The link aggregation appears in the list of available NICs in the UI as an individual interface.

Link aggregation is a networking option when you create an Oracle Solaris or Oracle VM Server for SPARC OS Configuration profile. You define the link aggregation load balancing policy and a MAC address policy.

1. Expand **Plan Management** in the Navigation pane.
2. Select **OS Configuration** in the **Profiles and Policies** tree. A list of existing OS Configuration profiles appears in the center pane.
3. Click **Create Profile** in the Actions or center pane.
4. Name the profile and enter a profile description. Select **Solaris** as the Subtype and **OSP SPARC** or **OSP x86** as the Target Type. Click **Next**.
5. The default setting is to automatically manage the OS with Oracle Enterprise Manager Ops Center and Deploy the Agent Controller. This option provides the most robust management capabilities. If you do not want to enable SAN storage connectivity, deselect Enable Multiplexed I/O (MPxIO). Click **Next**.
6. Select **Use Link Aggregation** for the Networking Services, then click **Next**.
7. Select a link aggregation name, define the Load Balancing Policy, LACP Mode, MAC Address Policy and the number of link aggregations you want. Click **Next**.
   - Load Balancing Policy determines the outgoing link:
     - L2: Hashes the MAC (L2) header of each packet
     - L3: Hashes the IP (L3) header of each packet
     - L4: Hashes the TCP, UDP, or other ULP (L4) header of each packet
   - LACP

     Configure the LACP for the switch and aggregation when your link aggregation topology has a switch connection that supports LACP.
   - MAC Address Policy
     - Auto: Use MAC address of any network interfaces in the Link Aggregation.
     - Fixed: Use MAC address of a specific network interface. Select the network interface to use in the next step.

Figure 13–12 Specify Link Aggregations

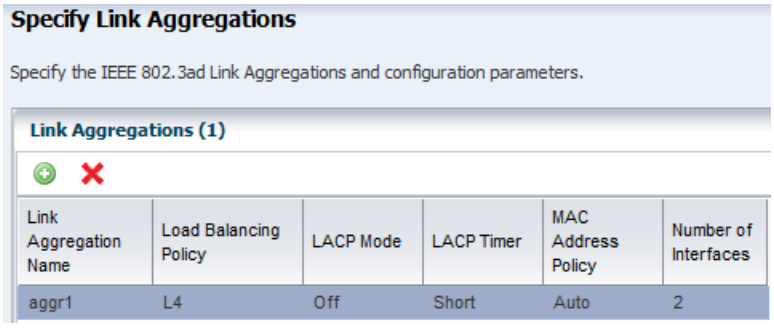

8. Specify the interface for each link aggregation. The number of interfaces is determined by the number that you defined in step 7. A list of all possible interfaces appears in the wizard. Work with your network administrator to know the interfaces that are available and which interfaces to configure.

Figure 13–13 Specify Link Aggregation Interfaces

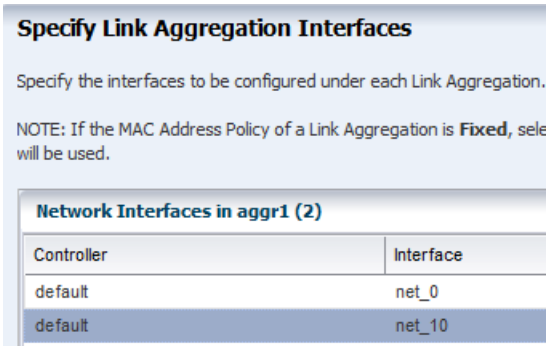

9. Review the summary of the parameters selected, then click **Finish** to create the OS Configuration profile for link aggregation.

Defining IPMP in an OS Configuration Profile

IP multipathing (IPMP) groups provide network failover for your Oracle Solaris operating system, Oracle VM Server for SPARC system, and guests.

You can configure one or more physical interfaces into an IPMP group. After configuring the IPMP group, the system monitors the interfaces in the IPMP group for failure. If an interface in the group fails or is removed for maintenance, IPMP migrates, or fails over, the failed interface's IP addresses. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections. The network access changes from the failed interface to the standby interface in the IPMP group and the data address of the failed interface migrates to the standby interface. See IP Multipathing Groups and Creating IPMP Groups for more information about IPMP groups.

1. Expand **Plan Management** in the Navigation pane.
2. Select **OS Configuration** in the **Profiles and Policies** tree. A list of existing OS Configuration profiles appears in the center pane.
3. Click **Create Profile** in the Actions or center pane.
4. Name the profile and enter a profile description. Select **Solaris** as the Subtype and **OSP SPARC** or **OSP x86** as the Target Type. Click **Next**.
5. The default setting is to automatically manage the OS with Oracle Enterprise Manager Ops Center and Deploy the Agent Controller. This option provides the most robust management capabilities. If you do not want to enable SAN storage connectivity, deselect Enable Multiplexed I/O (MPxIO). Click **Next**.
6. Select **Use IPMP** for the Networking Services, then click **Next**.
7. Use the default IPMP group name, or click the field and enter a name. Select the Failure Detection Policy, and enter the number of interfaces you want. Click **Next**.

   Probe based failure detection probes the target systems to determine the condition of the interface. Each target system must be attached to the same IP link as the members of the IPMP group.

Figure 13–14 IPMP Groups

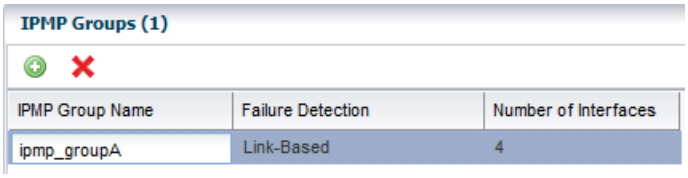

8. Specify the network interfaces. Select an interface from the list. For each interface, select either the Failover or Standby Interface check box. If you use Link and Probe based failure detection, you do not need to provide test IP addresses. The number of interfaces is determined by the number that you defined in step 7.

> **Note:** A list of all possible interfaces appears in the wizard. Work with your network administrator to know the interfaces that are available and which interfaces to configure.

Figure 13–15 Specify IPMP Interfaces

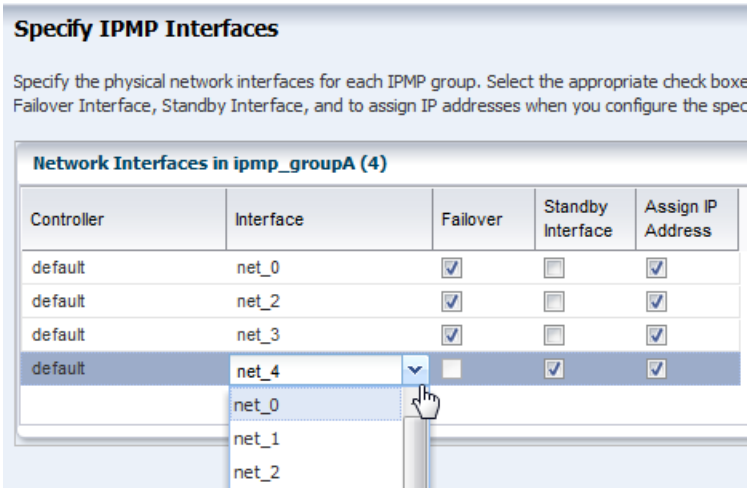

9. Review the summary of the parameters selected, then click **Finish** to create the OS Configuration profile for IPMP groups.

Migrating OS Provisioning Profiles to the New Format

If you have OS Provisioning profiles in a version of Oracle Enterprise Manager Ops Center earlier than 12c Release 2, use the Upgrade feature to upgrade to Oracle Enterprise Manager Ops Center 12.2 and automatically upgrade the profiles to the new format. The software upgrade feature automatically migrates the OS Provisioning profiles to an OS Provisioning profile and an OS Configuration profile. The original names of the migrated profiles are appended with *osp* or *osc*.

In addition to updating the profiles, existing deployment plans that use an OS Provisioning profile are also updated to use the new OS Provisioning and OS Configuration profiles. Migrating to the new format does not change the version number of the profile or plan.

For every OS provisioning profile created in Oracle Enterprise Manager Ops Center 12.1, a new OS Provisioning profile and a new OS Configuration profile are created in the 12.2 release. The new profiles have the following naming convention: *<12.1_profile_*

name>-osp and *<12.1_profile_name>_osc* respectively. Each profile is version one. For example, an OS Provisioning profile is named *S11_SPARC_LargeServer* in the 12.1 release. In the 12.2 release, the profile is converted into the following two profiles: *S11_SPARC_LargeServer_osp* and *S11_SPARC_LargeServer_osc*. The original profile *S11_SPARC_LargeServer* is deleted.

In the 12.1 release, both bare metal provisioning and logical domain guest provisioning were done using the same profiles. Beginning with 12.2, you can only provision a guest with the Logical Domain profile sub-type. When you upgrade from release 12.1 to release 12.2, a new Logical Domain profile is created for each Oracle Solaris 11 bare metal provisioning profile that was created in 12.1. The release 12.1 logical domain OS Provisioning profile is converted in the 12.2 release to a new logical domain OS Provisioning profile and a new OS Configuration profile and the original profile is deleted. The new profiles use the following naming convention: *<12.1_profile_name>-logicaldomain-osp* and *<12.1_profile_name>-logicaldomain-osc* respectively. For example, an OS Provisioning profile is named *S11_SPARC_LargeServer* in the 12.1 release. In the 12.2 release, the profile is converted into the following two profiles: *S11_SPARC_LargeServer_osp* and *S11_SPARC_LargeServer_osc*.

The deployment plans conversion is similar to the profile conversion. New deployment plans are created from the old plans, the plan version is updated, and the plan name is prefixed with <12.1_plan_name>. For example, the name for the new logical domain plan is *<12.1_plan_name>-logicaldomain-osp-plan*. The new plans use the new OS Provisioning and OS Configuration profiles. For Provision OS plans, the just the version of the 12.1 plan is updated and the OS Provisioning and OS Configuration profiles are added to the plan.

About Deployment Plans That Provision an Operating System

Deployment plans execute the OS Provisioning profile and OS Configuration profile on the targets you select, enabling you to provision in a consistent and repeatable way.

> **Note:** When adding an OS Provisioning profile and an OS Configuration profile to a plan, use the profiles that reference the same platform subtype.

Deployment plans are all based on defined templates that provide a sequence of steps to perform a task. After you create the profiles to define your tasks, you select a deployment plan template and create a plan that uses specific profiles. You can reuse the profiles in different plans to create consistency.

The following simple, multi-step, and complex plans include steps for OS provisioning:

- Provision OS: Use this plan to provision an operating system.
- Install Server: Use this plan to provision an operating system on the server and update the OS. The Update Software step enables you to update the OS or install additional OS packages. This step is run as part of the OS provisioning job.
- Configure M-Series Hardware, Create and Install Domain: Use this plan to configure an M-Series server, create dynamic system domains, provision OS on the domains, and update the domains.
- Configure and Install Dynamic System Domain: Use this plan to create dynamic system domains, provision and update OS on the domains.

- Configure Server Hardware and Install OS: Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- Configure and Install Logical Domains: Use this plan to create logical domains and provision OS on the logical domains.

When you select a plan to apply, a list of eligible targets appears in the target selector list. Targets are eligible when they meet the criteria of the profiles, such as type of platform, and for which you have the correct permissions to perform the provisioning tasks. Before deploying the plan on the selected targets, you have the ability to review, add, and override the configuration settings for the plan. For example, you can change the IP address and the boot interface.

Provisioning Oracle Solaris 11

The following information is in this section:

About Oracle Solaris 11 and Provisioning

Oracle Solaris 11 uses a new OS provisioning technology, called the Automated Installer. This feature replaces the older JumpStart Enterprise Toolkit (JET) technology that Oracle Enterprise Manager Ops Center uses to provision earlier versions of Oracle Solaris.

Oracle Enterprise Manager Ops Center reduces the complexity by using a local copy of the Oracle Solaris 11 Software Library on the Enterprise Controller and creating an Automated Installer server on the Proxy Controller for your use.

> **Note:** To provision an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must both be running on an Oracle Solaris 11 operating system. The repository resides on the Enterprise Controller and the Automated Installer server resides on the Proxy Controller. When the Enterprise Controller and Proxy Controller are not running on Oracle Solaris 11, the Oracle Solaris 11 library and OS provisioning actions are not available.

Oracle Enterprise Manager Ops Center can create and run multiple `installadm` services on the same Proxy Controller, one for each solaris-auto-install mini root. For example, Oracle Solaris 11, 11.1, 11.1.6.4.0, or whenever the solaris-auto-install version increases in an SRU (Support Repository Update). An SRU is a package of bug fixes and updates that releases on a regular basis. SRUs eliminate the ad hoc patching from Oracle Solaris 10 and earlier versions of the operating system. Each Oracle Solaris SRU builds upon, and only contains the changes from, the preceding Oracle Solaris 11 update. Oracle Solaris 11 uses a 5-digit taxonomy to define the SRU. The digits represent `Release.Update.SRU.Build.Respin`. For example, Oracle Solaris 11.1.6.4.0.

The Oracle Solaris 11 network architecture is significantly different from previous releases of Oracle Solaris. The implementation, the names of the network interfaces, the commands, and the methods for administering and configuring them is different from previous versions of Oracle Solaris. These changes were introduced to bring a more consistent and integrated experience to network administration, particularly as administrators add more-complex configurations including link aggregation, bridging, load balancing, or virtual networks. In addition to the traditional fixed networking configuration, Oracle Solaris 11 introduced automatic network configuration through network profiles.

The OS Provisioning and OS Configuration profiles that you use for provisioning Oracle Solaris 11 contain all of the information needed, such as type of target, OS image, time zone and language setup disk partitions, naming services and network details. With Oracle Solaris 11, you can define link aggregations or IPMP groups for advanced networking.

You can provision Oracle Solaris 11 zones as part of the OS installation. After a system is bootstrapped with a minimized operating system, the operating system is installed from the Oracle Solaris 11 Software Update Library in Oracle Enterprise Manager Ops Center. The zones are provisioned during the initial system reboot after the base operating system is installed.

Steps for Oracle Solaris 11 Provisioning Plan

A provisioning plan includes an Oracle Solaris 11 OS Provisioning profile and an OS Configuration profile. You can create a new plan or you can copy an existing plan and edit it to create a new plan.

> **Note:** When you create the plan, both the OS Provisioning profile and the OS Configuration profile must have the same platform subtype, either Solaris SPARC or Solaris x86.

Prerequisites

Perform the following before you provision the operating system:

- Verify that the Oracle Solaris 11 Software Library is configured on the Enterprise Controller and the package you want is available in the library.

  If the image is not in the Oracle Solaris 11 Software Library, import the OS image.

  > **Note:** Uploading the packages from Oracle to the library can take several hours.

- (Optional) Edit an existing OS Provisioning profile or create a new profile.
- Discover the service processors of the target systems.
- Verify that any scripts the profile uses are in a directory that the Enterprise Controller can access. You can save scripts in a local directory of the Enterprise Controller, or in a directory that the Enterprise Controller mounts using NFS.
- When you are provisioning a dynamic system domain of an M-Series server, the domain must have an IP address.

Note: It is a good practice to place the systems that you are going to provision in Maintenance Mode so that you can take the system offline without generating alerts and incidents.

Steps to Provision Oracle Solaris 11

Complete the following steps to provision an operating system:

1. Create an Oracle Solaris 11 Provisioning profile. See Creating an Oracle Solaris 11 OS Provisioning Profile.
2. Create an Oracle Solaris Configuration profile.
3. Create a deployment plan that uses Oracle Solaris 11 OS Provisioning and OS Configuration profiles, or verify that a plan and profiles are available and configured with the parameters you want to use.

   See Specifying Common Oracle Solaris 11 Parameters for more information.
4. Discover the service processor of the target system.
5. Place the asset in Maintenance Mode to prevent events related to a system going offline.
6. Select the deployment plan and define the targets for the plan. The target must be an x86 server or WAN boot-capable SPARC server. Review the configuration parameters and make any last minute changes in the plan before applying the plan to the target system.

Specifying Common Oracle Solaris 11 Parameters

Oracle Solaris 11 uses the following components for provisioning an operating system:

- Oracle Solaris 11 Software Library: A local version of the software package repository. You can update the Oracle Solaris 11 Software Library, as needed, and then provision multiple systems without using a network connection to Oracle for each provisioning job. See Software Libraries for more information about adding content to libraries.
- Installation manifest: Defines the system configuration, including what software to install and details on the virtualized environments to provision. A default manifest is included with each Image Packaging System (IPS) software repository.
- DCHP or WAN boot connection
  - x86 client: Requires a DHCP connection.
  - SPARC client: Requires a DHCP or WAN boot connection. Oracle Enterprise Manager Ops Center automatically sets up WAN boot connection.

Note: Oracle Solaris 11 only supports a ISC DHCP Server. Oracle Solaris DHCP Server is not supported.

Creating an Oracle Solaris 11 OS Provisioning Profile

You can create multiple profiles to respond to subtle variations in hardware attributes, software profiles, or your organization's requirements.

When you create an Oracle Solaris 11 OS Provisioning profile, you select the architecture, either SPARC or x86, to display the boot image and distribution for your

architecture. You must provide non-root user credentials and root user credentials. You can only use non-root user credentials to login or SSH to the client after install.

Oracle Solaris 11 System Software Groups

Oracle Solaris 11 provides three system software group packages that install different sets of packages appropriate for a larger server, a smaller server or non-global zone, or a graphical desktop environment. The boot image is associated with the software group.

As shown in Figure 13–16, you must select one of the following software groups:

- **large-server**: Provides common network services for an enterprise server. This group package also contains hardware drivers that are required for servers, such as InfiniBand drivers.
- **small-server**: Provides a smaller set of packages to be installed on a small server or non-global zone.
- **desktop**: Provides the GNOME desktop environment and other GUI tools such as web browsers and mail. It also includes drivers for graphics and audio devices.

Figure 13–16 System Software Groups

See the Oracle Solaris 11 Package Lists documentation for a detailed comparison of the three packages and list of contents.

Oracle Solaris 11 Feature Software Groups

The boot image might have additional software groups that you can select, such as trusted-desktop, storage-server and storage-nas. Each group has a tool tip that describes the group, as shown in Figure 13–17. You can select one or more of these groups.

Figure 13–17 Feature Software Groups

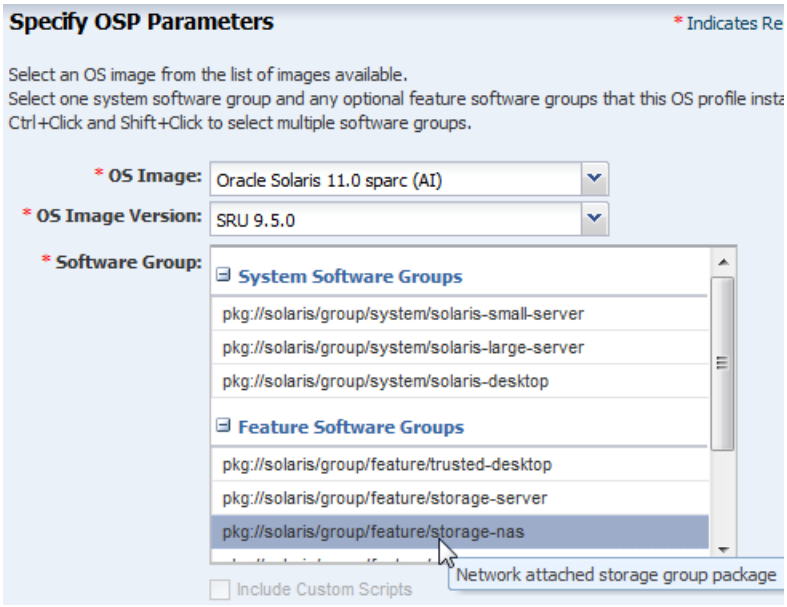

To install the operating system on an iSCSI disk, select the **Use iSCSI Disk** option when you create the profile, then specify the iSCSI disk settings. When you use this option, you must provide the following parameters when you deploy the OS Provisioning plan:

- Storage server IP
- SCSI disk LUN

Note: When you specify the naming service in the OS Provisioning profile, each IP address in the Name Server field must be entered in a new row. IP addresses 0.0.0.0 or 255.255.255.255 are allowed. In the Domain Name Search List field, enter each domain name, such as `1domain.com` and `2domain. com`, on a new line.

To Create an Oracle Solaris 11 OS Update Profile for Provisioning

You might create an update profile to install a specific Oracle Solaris 11 version or package, uninstall a package, or install a script. First, create a user-defined Oracle Solaris 11 Update Profile to define the action. When you create or edit an OS Provisioning profile, you can add your OS Update profile.

This example shows how to create an OS Update profile that installs the Firefox Extensions package.

1. Expand **Plan Management** in the Navigation pane, then select **Update Profiles** under **Profiles and Policies**.
2. Click **New Profile**.
3. Enter a name and description for the profile, then select **Solaris 11** as the Type.
4. Select Distribution, Category, View, and Version from the Search Criteria. Select Show Only **Support Repository Updates** to further filter the list. The available search criteria are determined by your selection. In this example, select **Applications/Plug-ins and Run-times** from the Category menu and **Latest Versions**. A list of available packages appears on the right side of the page.

5. Highlight the package. Click the **View Details** icon, which is the first icon, to get more information about your selection. To add the package components to the Profile Contents, select the components, then click the **Install** icon.

Figure 13–18 Create OS Update Profile

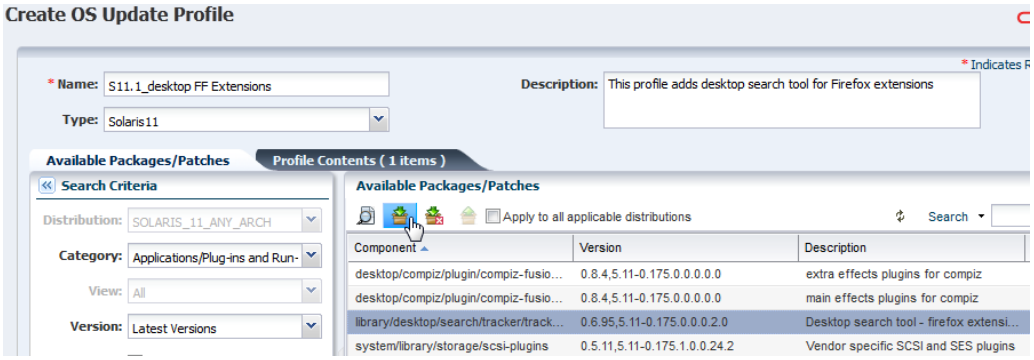

6. Click **Create OS Update Profile**. The profile appears in the list of Update profiles.

To Create an Oracle Solaris 11 Provisioning Profile

You can use the default profiles, copy a default profile to create a new profile, or create a new profile.

See the *Oracle Enterprise Manager Ops Center Provisioning Oracle Solaris 11 Operating System* document for how to create an OS provisioning profile, an OS Configuration profile, and a Provision OS deployment plan to complete a provisioning job.

This example shows how to create a new profile to provision a system with Oracle Solaris SPARC.

1. Expand **Plan Management** in the Navigation pane, then select **OS Provisioning** in the **Profiles and Policies** tree.
2. Click **Create Profile** in the Actions pane.
3. Name the profile and complete the profile description. A detailed description will help when determining which profile to use when several are available. Select Solaris x86 or Solaris SPARC from the Subtype and Target Type options, then click **Next**.

Figure 13–19 Identify Oracle Solaris 11 OS Provisioning Profile

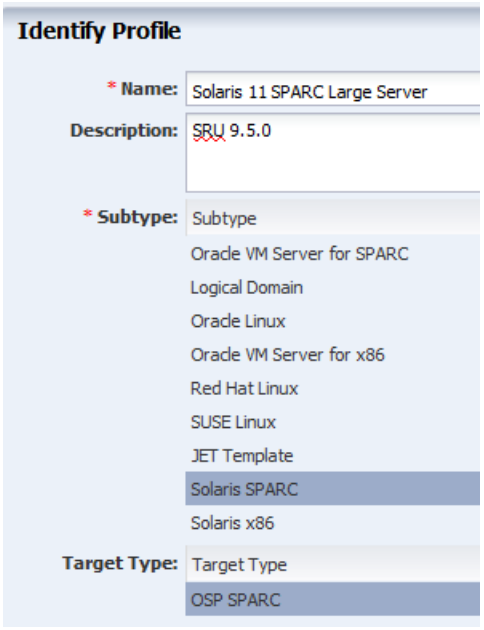

4. Select the OS image, OS Image Version, and Software group. Optionally, you can select a user-defined Solaris 11 Update profile. Click **Next**.

Figure 13–20 Specify OSP Parameters

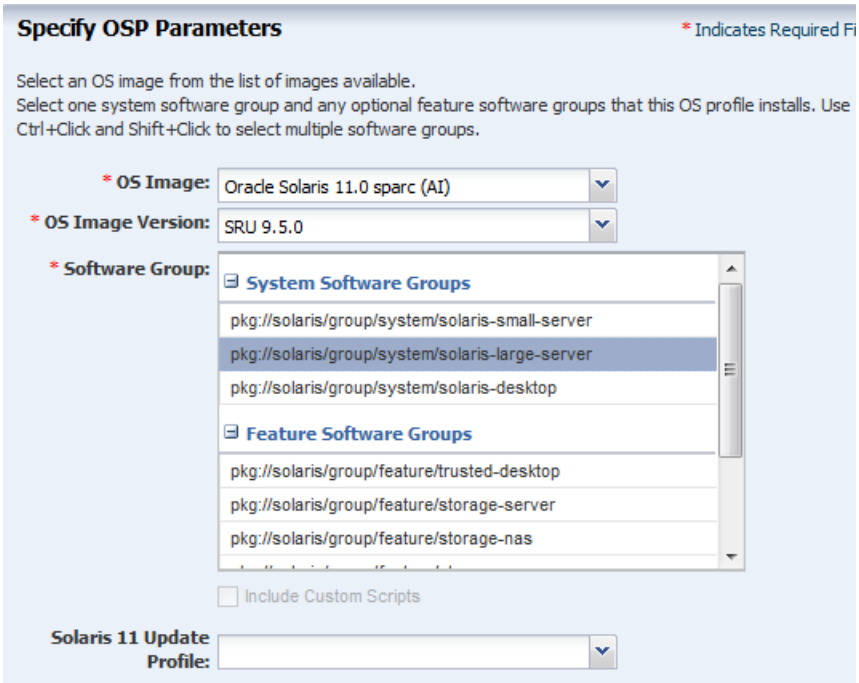

5. Edit the OS Setup parameters for language, time zone, terminal type, console serial port and baud rate, and NFS4 Domain, as needed. Enter a password for root in the Root Password and Confirm Password fields. Click **Next**.
   - Language: Select a language from the list.
   - Time Zone: Specify the time zone for the OS.

- Terminal Type: Select a terminal type from the list.
- Console Serial Port: To monitor the installation using a serial connection, select the correct console serial port device.
- Console Baud Rate: To monitor the installation using a serial connection, select the correct serial port device baud rate.
- NFS4 Domain: Enter the NFS4 domain name that the target system will use. The dynamic value for NFSv4 domain name enables the NFSv4 domain to be derived dynamically, at run time, based on the naming service configuration. You can also provide valid domain name to hard code the value for NFSv4 domain.
- Password: Enter the root password for the root user on systems provisioned using this profile. Re-enter the password for confirmation. The default password is *admin*.
- Manual Net Boot: Select this option when you want to manually control booting from the network. When you select this option, you are prompted to manually boot the system before the provisioning job completes.
  - For DHCP servers: Use the `boot net - install` command to manually boot the system over the network.
  - For WAN boot servers: Set the WAN boot parameters in the Open Boot PROM (OBP) before running the `boot net - install` command.

Note: The client-ID value of the WAN boot parameters must use the following format: 01<*macaddress*>. For example, client-id=0100123FF4E56E.

Figure 13–21 Specify OS Setup

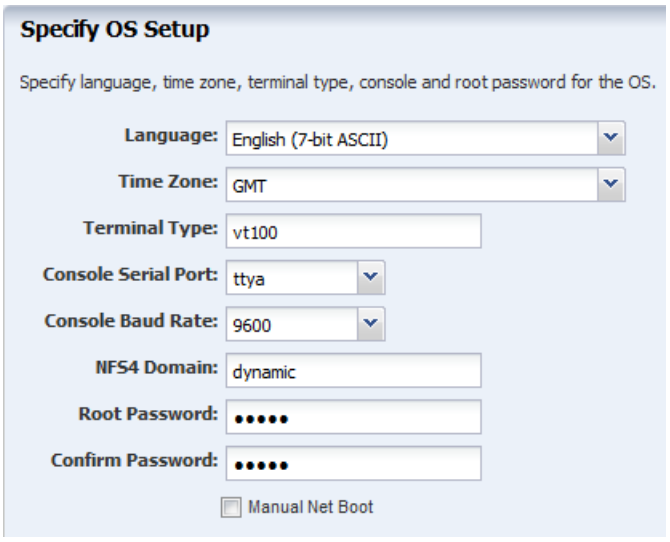

6. Specify the User Account details by entering a user name and password, then click **Next**.
7. On the Specify iSCSI Disk Image page, select the **Use iSCSI Disk** check box if you want to use an iSCSI disk for OS provisioning. When you select **Use iSCSI Disk**, another check box appears. Select **Manually Specify iSCSI Disk** to manually define the iSCSI disk resource assignments later. Click **Next**.

Figure 13–22 Specify iSCSI Disk Usage

Note: To assign a volume group and automatically create a new iSCSI disk, the DSL library must be attached to the server.

8. Review and edit the default file system layout, then click **Next**.

   This example uses the default file system layout. To specify changes to the default File System space, click the size field for the file system, and redefine.

Figure 13–23 Specify File System Layout

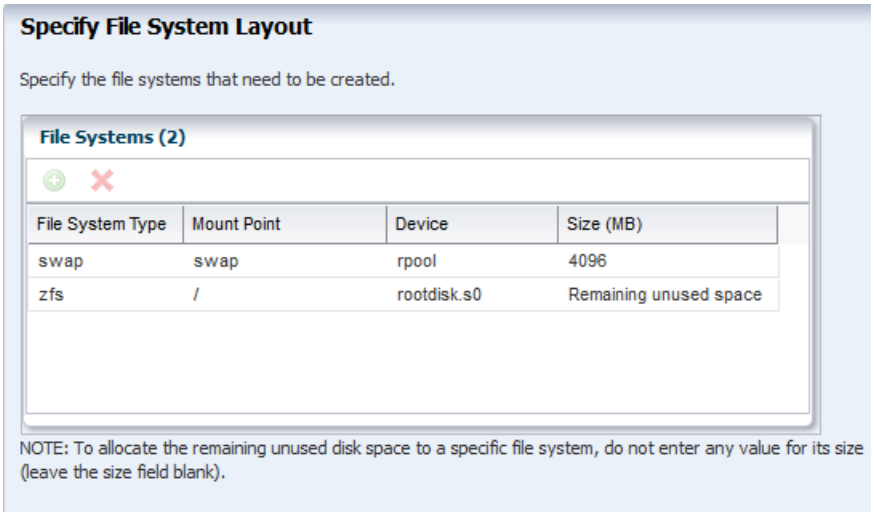

9. Select the Naming Service as **None**, **DNS**, **NIS**, or **LDAP**, then click **Next**.

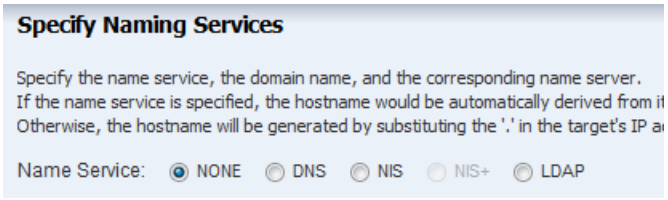

   When using a naming service, select the service and complete the required fields.

10. Review the parameters and click **Finish** to create the OS Provisioning profile for provisioning Oracle Solaris 11 operating system.

Figure 13–24 Oracle Solaris 11 OS Provisioning Profile Summary

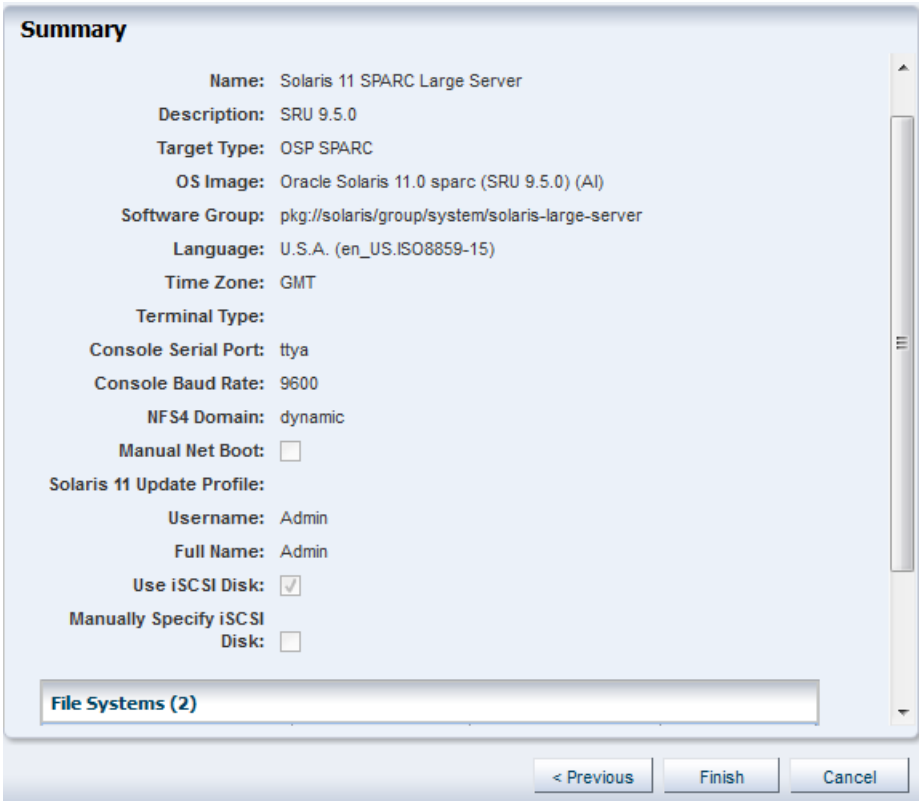

The profile appears in the center pane and in the Profiles and Policies section of **Plan Management**.

Creating an Oracle Solaris 11 OS Configuration Profile

The OS Configuration profile defines the networking configuration. You can use advanced networking configurations for Oracle Solaris.

When you create a configuration profile for Oracle Solaris, you can configure the following advanced networking options:

- **Link aggregation**: Provides high availability and higher throughput by aggregating multiple interfaces at the MAC layer. Link aggregation enables you to combine the capacity of multiple full-duplex Ethernet links into a single logical link.
- **IP Multipathing (IPMP)**: Provides features such as higher availability at the IP layer. IPMP enables you to configure multiple IP interfaces into a single IPMP group.

You can implement both Link Aggregation and IPMP methods on the same network because they work at different layers of the network stack.

After you create the profiles, you create a deployment plan to apply the profiles. As part of applying the plan, you can change the options that you defined earlier in the profiles.

Note: When you specify the network interfaces, select an interface for the controller from a list of 32 interfaces. The 32 interfaces that appear in the wizard are all possible network interfaces, not available interfaces. Your network administrator can give you the list of available networks. By default, the first interface listed is the boot interface.

To Create an OS Configuration Profile

1. Expand **Plan Management** in the Navigation pane, then select **OS Configuration** in the **Profiles and Policies** tree.
2. Click **Create Profile** in the Actions pane.
3. Name the profile and complete the profile description. A detailed description will help when determining which profile to use when several are available. Select Solaris x86 or Solaris SPARC from the Subtype and Target Type options, then click **Next**.

Figure 13–25 Identify Oracle Solaris 11 OS Configuration Profile

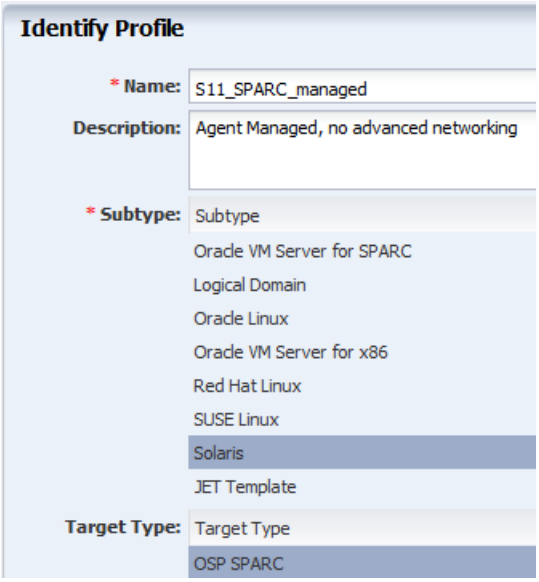

4. Click **Next** to accept the default selection to **Automatically manage with Oracle Enterprise Manager Ops Center** and **Deploy the Agent Controller**. This option provides the most robust management capabilities.

Figure 13–26 OS Management

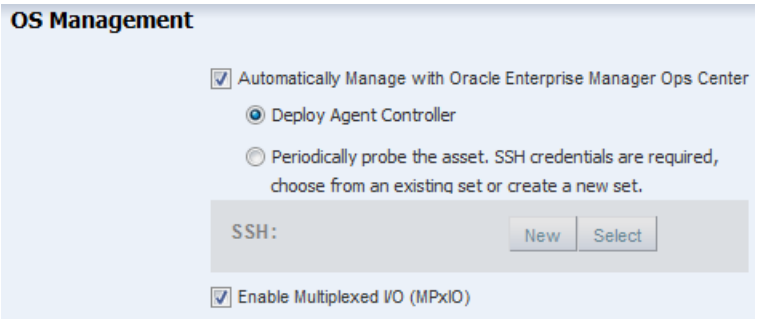

5. Specify the networking you want to establish for the target system.
   - **Use Link Aggregation**, go to Step 6.
   - **Use IPMP**, go to Step 7.
   - **None**, go to Step 8.

Figure 13–27 Specify Networking

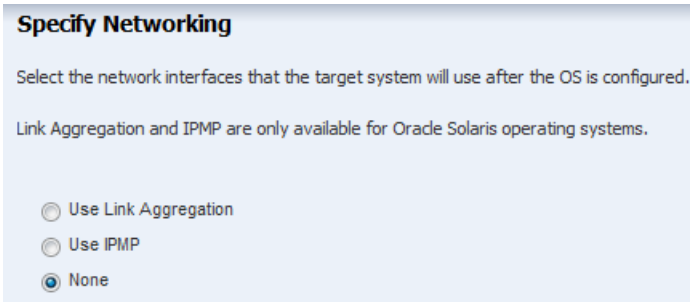

6. If you selected Link Aggregation in Step 5, complete the following steps for each link aggregation:
   a. Specify the Link Aggregation parameters. See Defining Link Aggregation in an OS Configuration Profile for details on the parameters.

Figure 13–28 Specify Link Aggregations

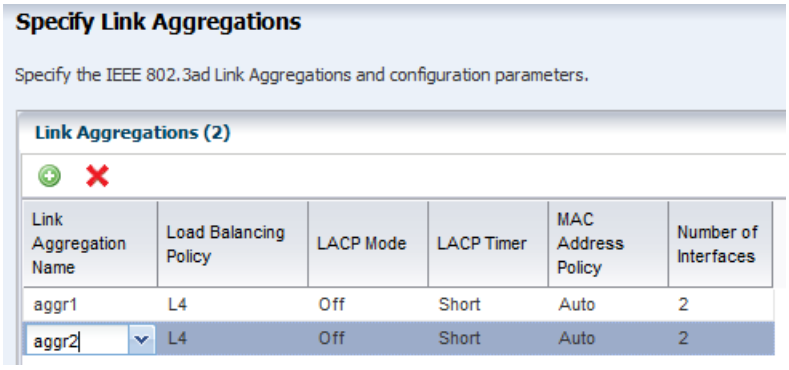

   b. Specify the Link Aggregation interfaces for each Link Aggregation.

      The 32 interfaces that appear in the list (net_0 - net_31) are possible network interfaces, not available interfaces. Contact your network administrator for a list of available interfaces.

Figure 13–29 Specify Link Aggregation Interfaces

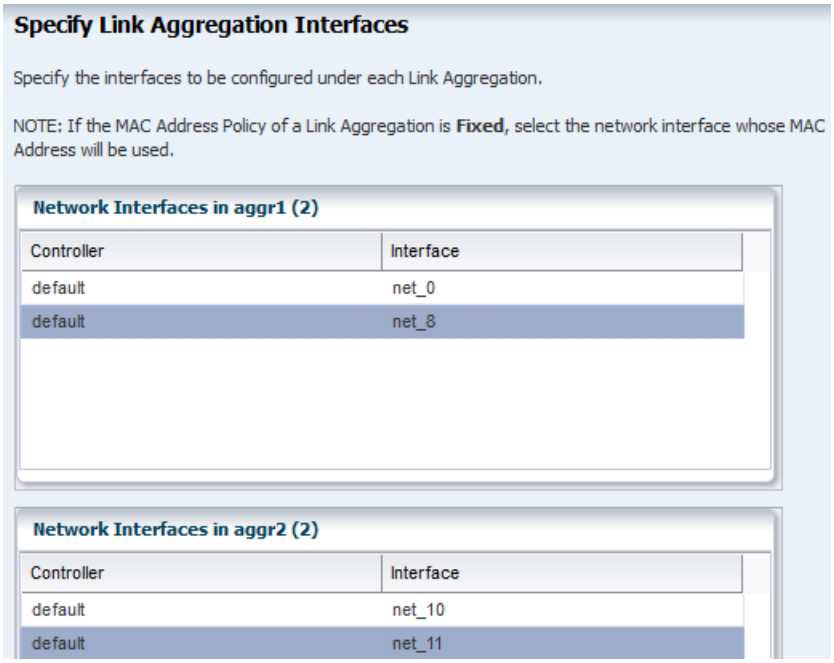

c. Go to Step 9.

7. If you selected Use IPMP in Step 5, complete the IPMP parameters. See Defining IPMP in an OS Configuration Profile for details.

a. Complete the Failure Detection method, either Link-Based or Link Based + Probe Based, and the number of interfaces.

Figure 13–30 Specify IPMP Groups

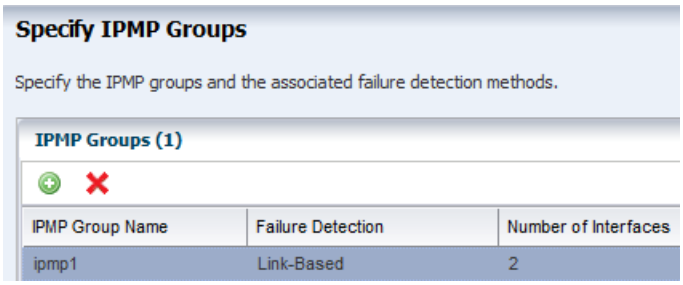

b. Specify the interface. Select the check boxes to define the IPMP configuration for **Failover** and **Standby Interface**. Select **Assign IP Address** to assign the IP address.

Figure 13–31 Specify IPMP Interfaces

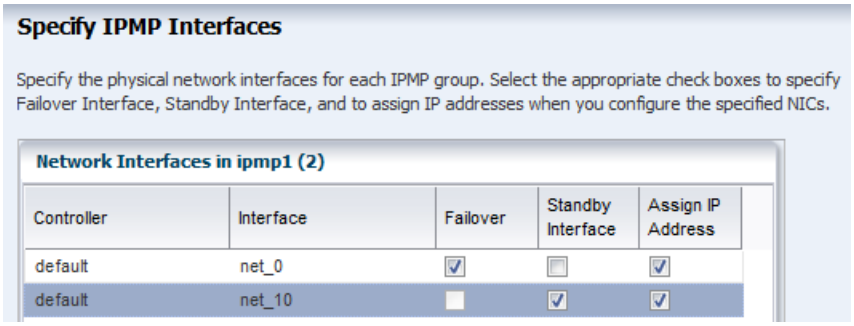

c. Go to Step 9.

8. If you selected None for Networking in Step 5, define the network interfaces that you want to use on the operating system. The boot interface is typically net_0, as shown in Figure 13–32. You define the IP address when you apply the plan to a target system. Click the **Add** icon to add other interfaces.

Figure 13–32 Specify Network Interfaces

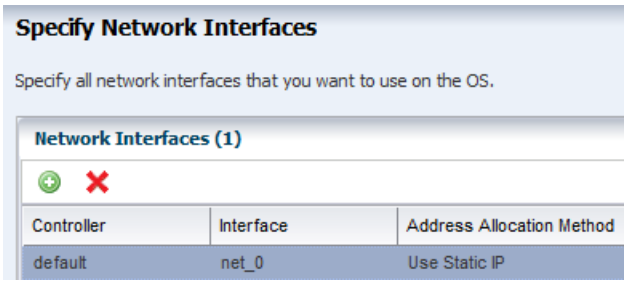

9. Review the Summary page, then click **Finish**. The new configuration appears in the table in the center pane. Click the profile to view the details.

Provisioning Oracle Solaris 9 and 10

You can create OS profiles for provisioning Oracle Solaris 9 or 10 on x86 or SPARC platforms. See Planning for Operating System Provisioning for the requirements needed for OS provisioning, including network requirements.

The OS Provisioning and OS Configuration profiles collect all the information such as type of target, OS image, time zone and language setup, required JET modules, disk partitions, naming services and network details.

Before You Begin

Perform the following before you provision the operating system:

- Import the OS image. Uploading the packages from Oracle to the library can take several hours.
- (Optional) Edit an existing OS Provisioning profile or create a new profile.
- Discover the service processors of the target systems.
- Verify that the Dynamic Host Configuration Protocol (DHCP) services are enabled on Proxy Controllers. You cannot create a profile or assign any network if the DHCP services are not enabled. The Install Server option to provision OS on a server is not enabled if the DHCP is not enabled on any of the interfaces.

> **Note:** Oracle Solaris 10 supports an Oracle Solaris DHCP Server. The external DHCP-related files are copied only if the Proxy Controller is running on an Oracle Solaris 10 operating system.

- Verify that any scripts the profile uses are in a directory that the Enterprise Controller can access. You can save scripts in a local directory of the Enterprise Controller, or in a directory that the Enterprise Controller mounts using NFS.
- When you are provisioning a dynamic system domain of an M-Series server, the domain must have an IP address.

> **Note:** It is a good practice to place the systems that you are going to provision in Maintenance Mode so that you can take the system offline without generating alerts and incidents.

Complete the following steps to provision an operating system:

1. Verify that the OS Provisioning profile and OS Configuration profile are available and configured with the parameters you want to use. See Specifying Common Oracle Solaris 9 and 10 Parameters for more information.
2. Create a deployment plan that enables OS provisioning. The plan must contain an OS Provisioning profile and an OS Configuration profile that use the same subtype, either Oracle Solaris SPARC or x86.
3. Discover the service processors of the target systems.
4. Place the asset in Maintenance Mode to prevent events related to a system going offline.
5. Select the deployment plan and define the targets for the plan. Make any last minute changes in the plan, then submit the job.

See the *Oracle Enterprise Manager Ops Center Provisioning Oracle Solaris 10 Operating System Guide* for an end-to-end example.

Specifying Common Oracle Solaris 9 and 10 Parameters

To create a profile that installs the Oracle Solaris 9 or 10 OS, specify the following parameters:

- Manual reboot: By default, the profile reboots the OS. You can choose the Manual Net Boot option to enable manual control of network boot operations for the target system.

> **Note:** The Enterprise Controller cannot remotely control the network boot process on systems that do not have a service processor. When your target system does not have a service processor, you must select the Manual Net Boot option.

- Custom Scripts: When you specify the OS Parameters, you have the option to include custom scripts. This feature is disabled for Oracle Solaris 11 profiles.

About JumpStart Enterprise Toolkit (JET) for Oracle Solaris

For Oracle Solaris 9 and 10 only, you can optionally use JumpStart Enterprise Toolkit (JET) modules to specify additional installation parameters. Oracle Solaris 11 uses the Automated Installer (AI) instead of JET.

JET provides a framework to simplify and extend the JumpStart functionality provided within the Oracle Solaris 9 and 10 operating system. The `SUNWjet` and `JetFLASH` packages are installed on the Proxy Controller during installation when the Proxy Controller is installed on an Oracle Solaris 10 operating system.

Using JET provides more options for defining the Jumpstart parameters. When you install JET on a JumpStart server, you have the following advantages:

- Install multiple versions of Oracle Solaris
- Deploy flash archives
- Utilize multiple boot methods
- Install recommended patches
- Configure all your network interfaces

Note: You cannot define IPMP groups or link aggregation for a JET template profile.

Table 13–3 describes the JET modules that are installed on the Proxy Controller.

Table 13–3 JET Modules and Associated Packages

JET Module Name	JET Package	Description
base_config	SUNWjet	Provides the standard installation configuration for the client, including the information required to set up the JumpStart server to allow the client to boot and build.
custom	SUNWjet	Adds functionality to the JumpStart framework to handle packages, patches, scripts, and files.
flash	JetFLASH	Adds the ability for the JumpStart server to deliver Solaris images in Solaris Flash format.

Within the Oracle Enterprise Manager Ops Center UI, there are 2 methods of influencing the JET template variables:

- Import a JET Template
- Add JET variables to the OS provisioning profile

You cannot manipulate the JET template in the UI. When you want to make changes to a template, make the changes and then import the template.

See Appendix C, "JumpStart Enterprise Toolkit" for a list of SUNWjet parameters, to learn more about JET and how to use JET in a profile.

To Create a JET Template

To create a profile that uses the JumpStart Enterprise Toolkit (JET), select a JET template that defines all the parameters for OS provisioning.

Place the JET template on a directory that the Enterprise Controller can access. You can also create a JET template on the Enterprise Controller in the directory `/opt/SUNWjet/Templates`, using the following command:

```
./make_template template_name
```

A sample template is provided. You can make a copy and change the values in the JET template as required. During provisioning, the OS provisioning parameters are read from the template. After you create the JET templates you want and save them on the Enterprise Controller, you can use them in your Oracle Solaris provisioning profiles.

Creating an Oracle Solaris 9 or 10 OS Provisioning Profile

The OS Provisioning profile defines the OS provisioning parameters, including the platform-specific OS image and software package, file system layout, user accounts, naming services, and other installation requirements.

Complete the following steps to create an OS Provisioning profile:

1. Expand **Plan Management** in the Navigation pane.
2. Select **OS Provisioning** under the **Profiles and Policies** section.
3. Click **Create Profile** in the Actions pane.
4. Define the following profile parameters in the **Create Profile-OS Provisioning** wizard, then click **Next**.
   - Name: The name of the profile.
   - Description: A description of the profile.
   - Subtype: Select Solaris SPARC or Solaris x86.
   - Target Type: Select the target type, either SPARC or x86.
5. Select an **OS Image**, **OS Image Version**, and **Software Group**. This example does not include custom scripts. Click **Next**.
6. Specify the following OS setup parameters:
   - Language: Select a language for the OS.
   - TimeZone: Specify the time zone for the OS.
   - Terminal Type: Enter a terminal type, if other than the default type listed.
   - Console Serial Port: A default port appears in the wizard. If incorrect, select the correct console serial port device for your environment. This port enables you to monitor the installation using a serial connection.
   - Console Baud Rate: A default serial port device baud rate is provided. If incorrect, select the correct baud rate for your device.
   - NFS4 Domain: Enter the NFS4 domain name that the target system will use. The dynamic value for NFSv4 domain name enables the NFSv4 domain to be derived dynamically, at run time, based on the naming service configuration. You can also provide valid domain name to hard code the value for NFSv4 domain.
   - Password: Enter the root password for the root user on systems provisioned using this profile. Re-enter the password for confirmation.
7. Click **Next** to skip specifying the Installation Parameters.
8. Review and edit the default file system layout, then click **Next**.

To specify changes, click the size field for the file system and redefine the size. To add another file system, click the **Add** icon and complete the fields.

9. Select the naming service you want, or select **None**, then click **Next**.
10. Review the parameters selected for Oracle Solaris 10 operating system provisioning, then click **Finish** to save the profile.

Creating an Oracle Solaris 9 or 10 OS Configuration Profile

The OS Configuration profile defines the networking configuration.

When you create a configuration profile for Oracle Solaris, you can configure the following advanced networking options:

- **Link aggregation**: Provides high availability and higher throughput by aggregating multiple interfaces at the MAC layer. Link aggregation enables you to combine the capacity of multiple full-duplex Ethernet links into a single logical link.
- **IP Multipathing (IPMP)**: Provides features such as higher availability at the IP layer. IPMP enables you to configure multiple IP interfaces into a single IPMP group.

You can implement both Link Aggregation and IPMP methods on the same network because they work at different layers of the network stack.

After you create the profiles, you create a deployment plan to apply the profiles. As part of applying the plan, you can change some of the options that you defined earlier in the profiles.

Provisioning an Operating System on Logical Domains

To provision an operating system on logical domains, you must select the Logical Domain subtype when you create OS Provisioning and OS Configuration profiles. You cannot use the OS profiles that are created for bare-metal provisioning. Refer to Provisioning OS on Logical Domains in Chapter 19, "Oracle VM Server for SPARC" for more information about creating OS provisioning and configuration profiles for provisioning OS on logical domains.

Provisioning an Operating System on an Oracle Solaris Cluster

To provision an operating system on an Oracle Solaris Cluster, you provision the same operating system on all nodes of a cluster. The Cluster OS profile handles the pre-action and post-action operations. See the Oracle Solaris Cluster documentation and Appendix A, "Oracle Solaris Cluster" for how to install a new Oracle Solaris Cluster and to maintain existing Oracle Solaris Clusters.

Provisioning Linux

You can create OS Provisioning and OS Configuration profiles for provisioning Linux OS on x86 systems. The profiles collect all the information such as type of target, OS image, time zone and language setup disk partitions, naming services and network details.

Provisioning Oracle Linux and other supported versions of Linux is very similar to provisioning Oracle Solaris 10 x86.You add the Linux image, create the OS

Provisioning and OS Configuration profiles, create a Provision OS deployment plan, then apply the plan to provision the operating system.

The OS Provisioning plan will use Kickstart as the install mechanism to perform the installation. You do not need to do anything to enable Kickstart.

Provisioning requires a DHCP-enabled network interface for the boot interface. You can add multiple networks, as long as the networks are available and defined in the Enterprise Controller. You can select a NIC from the list of available logical interfaces for each network or you can use the Address Allocation Method for the selected networks. You cannot use the Address Allocation method for the boot interface. When you use a static IP address, you must provide the IP address when you apply a deployment plan that uses the profile. The IP address is assigned to the target system after provisioning.

Note: When you specify the naming service in the OS Provisioning profile, each IP address in the Name Server field must be entered in a new row. IP addresses 0.0.0.0 or 255.255.255.255 are allowed. In the Domain Name Search List field, enter each domain name, such as `1domain.com` and `2domain.com`, on a new line.

Before You Begin

Perform the following before you provision the operating system:

- Import the OS image. Uploading the packages from Oracle to the library can take several hours.
- (Optional) Edit an existing OS Provisioning profile or create a new profile.
- Discover the service processors of the target systems.
- Verify that the Dynamic Host Configuration Protocol (DHCP) services are enabled on Proxy Controllers. You cannot create an OS Configuration profile or assign any network if the DHCP services are not enabled. The Install Server option to provision OS on a server is not enabled if the DHCP is not enabled on any of the interfaces.
- Verify that any scripts the profile uses are in a directory that the Enterprise Controller can access. You can save scripts in a local directory of the Enterprise Controller, or in a directory that the Enterprise Controller mounts using NFS.

Note: It is a good practice to place the systems that you are going to provision in Maintenance Mode so that you can take the system offline without generating alerts and incidents.

To Provision Linux

Complete the following steps to provision an operating system:

1. Verify that the Linux image you want to use is available in the library, or import the Linux image. See Chapter 5, "Software Libraries" and the `Deploy How To library` for more information.
2. Create an OS Provisioning profile and an OS Configuration profile using Linux as the Target Type. Configure the profiles with the parameters you want to use. See Specifying Common Linux Parameters and Specifying SuSE Parameters for more information.

3. Create a Provision OS deployment plan or other deployment plan that enables OS provisioning.
4. Discover the service processors of the target systems.
5. Place the asset in Maintenance Mode to prevent events related to a system going offline.
6. Select the deployment plan and click **Apply Plan** to define the targets for the plan. Make any last minute changes in the plan, then submit the job.
7. When the job completes and the new operating system is provisioned, take the asset out of Maintenance Mode.

Specifying Common Linux Parameters

Specify the following parameters:

- Installation number: The number that enables you to install all of the Linux software that is included in your subscription.
- Partition action: Use this parameter when you want to change the disk partition of the system.
  - You can opt to remove all the existing Linux partitions and retain the non-Linux partitions. You can provide specification for the new partitions.
  - You can opt to preserve all the existing partitions. You must define new partitions, outside of the partitions that exist, in which to install the OS.
  - You can opt to remove all the existing partitions. Define specification for the new partitions.
- Install protocol: Specify HTTP or NFS as the install protocol.
- Kernel parameters: Enter kernel parameters for the GRUB menu of the target system, when needed.
- MD5 Checksum: Select this option to use MD5 encryption for user passwords.
- Reboot action: Select whether you want to reboot the target system after OS installation.
- Disk label initialization: Select this option to initialize labels on new disks. This option creates labels that are appropriate for the target system architecture.
- Shadow passwords: Select this option to use an `/etc/shadow` file to store passwords on the target system.
- Clear master boot record: Select this option to clear all invalid partition tables.
- Linux packages: You can specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude any package, enter the package name preceded by a dash (-).

Specifying SuSE Parameters

To create a profile that installs the SuSE Linux OS, specify the following parameters:

- FTP proxy server: Enter the name of the FTP proxy server to support FTP services.
- HTTP proxy server: Enter the name of the HTTP proxy server to support HTTP services.
- Install protocol: Specify HTTP or NFS as the install protocol.

- Enable proxy servers: Select this option to enable the FTP and HTTP proxy servers that you specified in the FTP Proxy Server and HTTP Proxy Server fields.
- Kernel parameters: Enter kernel parameters for the GRUB menu of the target system, when necessary.
- Reboot action: Select whether you want to reboot the target system after OS installation.
- Linux packages: You can specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude any package, enter the package name preceded by a dash (-).

Related Resources for Operating System Provisioning

See the following workflows and how to documentation in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm for end-to-end examples and workflows.

- *Oracle Enterprise Manager Ops Center Deploy Operating Systems Workflow*
- *Oracle Enterprise Manager Ops Center Provisioning Oracle Solaris 11 Operating Systems*
- *Oracle Enterprise Manager Ops Center Provisioning Oracle Solaris 10 Operating Systems*

For information related to this feature, go to one of the following resources in the *Oracle Enterprise Manager Ops Center Feature Reference Guide*:

- See Chapter 4, "Monitoring Rules and Policies" for information on how monitoring rules and policies work in the software.
- See Chapter 12, "Operating System Management" and Chapter 14, "Operating System Updates" for information about managing, patching, or updating, your operating systems.
- See Chapter 18, "Oracle Solaris Zones" for information about zones and how you can use Oracle Enterprise Manager Ops Center to efficiently manage all phases of zones lifecycle.
- Chapter 19, "Oracle VM Server for SPARC" for the requirements and information needed to provision Oracle VM Server for SPARC and domains
- See Chapter 20, "Oracle VM Server for x86" for the requirements and information needed to provision Oracle VM Server for SPARC and domains.
- See Appendix A, "Oracle Solaris Cluster" for how to use Oracle Enterprise Manager Ops Center to install and upgrade Oracle Solaris Clusters.

See Chapter 11, "Hardware" for information about provisioning firmware.

For more information about how to set up and manage the Oracle Enterprise Manager Ops Center infrastructure, including DHCP and WAN boot, see the *Oracle Enterprise Manager Ops Center Administration Guide*.

For in-depth information about Oracle Linux, Oracle Solaris, DHCP, WAN boot, and related features, see the following Oracle documentation:

- For a list of the Oracle Linux documentation available in HTML and PDF formats, visit the Oracle Linux Documentation at http://www.oracle.com/us/technologies/linux/index.html.

- For a list of the Oracle Solaris 11 and 11.1 documentation available in HTML and PDF formats, visit the Oracle Solaris 11 Documentation at http://www.oracle.com/technetwork/documentation/solaris-11-192991.html.
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html.
- For a list of the Oracle Solaris 8 and 9 documentation, visit the Legacy Solaris Documentation at http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html.
- For more information about JET resources and documentation, see *Solaris 10 10/09 Installation Guide: Custom JumpStart and Advanced Installations* available at http://docs.oracle.com/cd/E18752_01/html/821-1911/index.html.
- For JET documentation and to download additional modules, see http://www.oracle.com/technetwork/systems/jet-toolkit/index.html.

14

Operating System Updates

This chapter describes the operating system update features that are available in the software.

The following information is included:

Introduction to Operating System Updates

Oracle Enterprise Manager Ops Center reduces the complexity of updating a large number of systems, standardizes the update installation process, minimizes downtime, and enables you to choose the level of automation.

You can maintain your Oracle Solaris, Linux, and Microsoft Windows operating systems to the recommended and latest updates and perform complex update tasks in a consistent manner. For most platforms, the update features help you to perform the following tasks:

- Manage different operating system update conditions that exist for installing an update
- Identify dependencies
- Download update packages or updates from the appropriate vendor sites
- Run an update simulation to test the update in your environment

- Rollback your systems to a previous state if an update is not stable in your environment
- Maintain consistent component configuration of your systems to the latest security updates

The list of supported operating system releases and functionality is available in the *Oracle Enterprise Manager Ops Center Certified Systems Matrix Guide*.

The update features include:

- Catalogs
- Reports
- Update Profiles
- Update Policies
- Deployment Plans

Oracle Enterprise Manager Ops Center provides one stop solution for all the requirements for updating your operating systems. You can use Update Profiles and plans to define which components must be installed and the level of automation during the installation.

You create update jobs with operating system update profiles and policies to update an operating system. OS update profiles and policies define which updates to install, and how the update job proceeds after determining the update dependencies and user interaction. A set of system-defined update profiles is available in the UI. You can copy the profiles and edit them to create your own customized profiles.

When you run an OS Update job, Oracle Enterprise Manager Ops Center performs the following actions:

1. Locks the asset.
2. Creates a snapshot for Oracle Solaris 8-10 and Linux operating systems.
3. Queues the job on the Enterprise Controller for the associated Proxy Controller. The Agent Controller retrieves the job and performs the tasks on the asset.
4. Saves the job log on the Enterprise Controller.
5. Unlocks the asset.

A variety of operating system reports are available to give you insight into the operating system compliance status, the state of your operating system update levels, and provide information about the recommended updates and packages.

You can view your managed operating systems with the All Assets view or the Operating Systems view, as shown in Figure 14–1 and Figure 14–2.

Figure 14–1 All Assets View

Figure 14–2 Operating Systems View

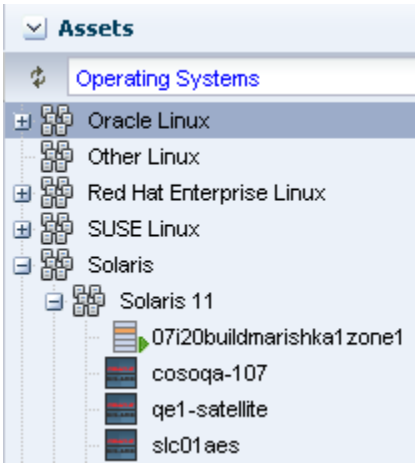

Requirements for Updating Operating Systems

The Enterprise Controller obtains information about latest updates from the Knowledge Base, Oracle Solaris 11 parent repositories on Oracle.com, and vendor sites. You must have an Internet connection to obtain the updates and packages from the various locations. In the absence of an Internet connection to the Enterprise Controller, you can get the latest updates by using a special script, called the harvester, to create local versions of the Knowledge Base and Oracle Solaris 11 Software Update Library. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about connection modes and how to use the harvester script.

The operating system update feature has the following requirements:

- Agent managed: Linux and Oracle Solaris 8, 9, and 10 operating system must be agent-managed to perform most operating system update tasks. You can update Oracle Solaris 11 and Microsoft Windows with an agentlessly managed operating system. See Using Agent Management for Operating Systems for details and how to switch management modes.
- Access to updates: The software is designed to use a secure Internet connection to obtain operating system updates as they become available. If you cannot use an Internet connection from within your datacenter, you can run the software in disconnected mode and download the latest updates to an external storage device and use that to update the software libraries.

- Established Software Update Library: For Linux or Oracle Solaris, you must create an update library to store the update information. See Chapter 5, "Software Libraries" for information on establishing and using software update libraries.
- Compatible configuration for Oracle Solaris 11: To update an Oracle Solaris 11 operating system, the Enterprise Controller and the Proxy Controller must be running on an Oracle Solaris 11 operating system.
- Valid update credentials: You must have update credentials for each platform vendor. You must provide a valid My Oracle Support (MOS) account for Oracle Linux and Oracle Solaris. Valid vendor credentials are required for SuSE Linux and Microsoft Windows. Typically, you supply the credentials when you install Oracle Enterprise Manager Ops Center. See `Authentications` in the *Oracle Enterprise Manager Ops Center Administration Guide* for how to add or edit your update credentials.

Manage the operating systems with an agent for the most robust feature set, including reports, monitoring, and analytics. See Discovering and Managing Assets for more about the discovery and management process. See Using Agent Management for Operating Systems for the differences between the two modes and how to change the agent management mode.

Methods of Running an Update Job

You can use the following methods to run an update job:

- Deployment plans that use the update profiles and policies. A deployment plan defines the profiles used, in some cases the plans used, and the sequence of operations (or steps). The following deployment plans include update: Install Server, Software Deployment/Update, and Update Solaris 11 OS.
- System catalogs for Oracle Solaris 8 - 10 and Linux.
- Reports.
- Update profiles.

Options Available When Running an Update Job

You have the following options available when running an update job:

- Select update profiles and policies.
- Select different targets for each task in the job.
- Select job simulation mode. Simulating a patching job helps to estimate the time required to run the job, to know the patch dependencies and the expected job result. In the simulation mode, you can select to download the required updates.
- Failure policy to determine the action when a task fails.

Roles for Operating System Updates

Table 14–1 lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator when you do not have the necessary role or privilege to complete a task.

See Asset Management for details on using the discovery feature to add assets and for more information on creating user-defined groups. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise*

Manager Ops Center Administration Guide for information about the different roles and the permissions they grant.

Table 14–1 Operating System Management Roles and Permissions

Task	Role
New Update OS Job	Update Admin
Deploy or Update Software	Update Admin
Simulate an OS Update	Update Admin or Update Sim Admin
Compare System Catalog	Update Admin
Create Catalog Snapshot	Update Admin
View and Modify Catalog	Update Admin
Update Management Credentials	Security Admin
Any Actions related to changing credentials	Security Admin
Import image	Storage Admin
Upload image	Storage Admin
Upload image	Storage Admin
Unconfigure, SCCM Configuration	Ops Center Admin
Reboot, upgrade Agent Controller	Asset Admin
Edit Tags	Asset Admin
Edit Attributes	Asset Admin

Actions for Operating System Updates

Two management modes are available, agent managed and agentlessly managed. To perform updates to your Linux or Oracle Solaris 8, 9, or 10 operating system, the operating system must be a managed asset. Oracle Solaris 11 and Microsoft Windows do not require an agent managed operating system.

You can manage your operating systems by installing an Agent Controller on the OS or by using SSH to perform tasks. You must have an agent managed operating system to use the OS Update features, including the system catalogs, OS update jobs, and reports. See Using Agent Management for Operating Systems for more details, including when an agent is required, how to determine the management mode, and how to change the management mode.

After you manage your assets, you can perform the following actions:

- Monitor your physical and virtual operating systems
- View OS utilization for Oracle Solaris and Linux operating systems
- Provision Oracle Solaris and Linux operating systems
- Manage Oracle Solaris boot environments
- Create a System Catalog
- Create an OS Report
- Update Oracle Solaris and Linux operating systems

- Update Microsoft Windows operating systems

Location of Operating System Updates in the User Interface

To see operating system information, expand Assets in the Navigation pane, then select the operating system. You can use the Operating Systems filter to just display the operating systems, as shown in Figure 14–3.

Figure 14–3 Operating System Filter

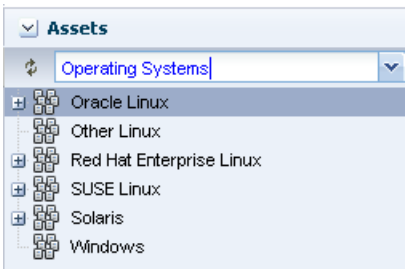

Using System Catalogs

A system catalog is a software inventory of installed instances and versions of Oracle Solaris 8 - 10 OS updates, Linux RPMs, and local software. Oracle Enterprise Manager Ops Center automatically takes a snapshot of the operating system after executing any job on the OS, including when you discover and manage the operating system, start the Agent Controller, and when you update the operating system. A snapshot is stored on the Enterprise Controller as a catalog with the time stamp and job details after every update job that you run on a system.

Note: Oracle Solaris 11 operating systems do not have or use snapshots. or system catalogs. Instead, the Oracle Solaris 11 operating system manages the OS packages.

You can create a new catalog at any time and use it to record the state of a system. Catalogs enable you to rollback your system to any previous configuration or to create a profile that you can use to apply a consistent configuration throughout your data center.

The Catalog List contains all of the snapshots. When you create a historical catalog, the current state of the selected system is identified and stored as the previous catalog of the system. The saved previous catalog is the most recent system catalog.

Note: You can create a historical catalog only for the current state of the system.

The catalog list always provides the listing of the most recent catalog. The software updates the catalog list whenever you update a system or create an historical catalog. You can identify the current catalog by the time stamp. You can use an historical catalog to create a profile and apply it to configure other systems.

Viewing and Modifying a Catalog

When you are using dual boot environments for Oracle Solaris Live Upgrade (Oracle Solaris 10), the catalog displays the inventory of the active boot environment of the operating system. To view the catalog of an alternate boot environment (ABE), you must first activate the ABE from the UI, and then wait for the job to finish. The software updates the current catalog and contains the ABE catalog information and OS software components. This automatically updates the catalog of any zones.

When you have an alternate boot environment, you cannot create and compare catalogs until you activate the ABE. By default, only the catalogs of the active boot environment are compared.

Comparing System Catalogs

You can compare two managed systems or two system catalogs for differences in the installed update components. You can also compare the current system catalog and saved snapshots of the same managed system to examine the differences in the components that are installed and uninstalled after executing a job.

Use the Compare Catalogs option to change the software components of a particular operating system to that of the source system.

The following options are available when you compare catalogs:

- Differences Between Systems: Displays the difference between the source and the target systems update components. The difference appears in the Compare Catalog window.
- Tasks to Make Target Like Source: Creates the list of components that must be installed on the target system. Select Include for the components to install on the target system.

About Operating System Update Reports

To ensure that your managed systems are up-to-date, you must determine which updates (packages and patches) and actions to apply to your system. The OS update reports help you to determine the updates that are applicable to your systems and how many of the applicable updates are compliant or not compliant for the selected systems.

The OS Update reports enable you to check for new updates and update your systems. You can get a general report, or test a system for available fixes.

When you create a report, you select the criteria that are relevant to you, such as a list of hosts that have a specific patch or a list of hosts that do not have a specific patch. See Creating an Operating System Report in Chapter 10 for detailed information about OS Update reports and Table 10–3 for a list of available reports for operating system.

Creating Update Policies

An Update policy defines the level of interaction required during an OS update job for Oracle Solaris and Linux operating systems. Policies define how to answer any questions that are raised during installation, uninstallation, or upgrade of Oracle Solaris and Linux updates.

The following system-defined policies are available in the software:

- Ask for All: The software stops the update job for each action and consults you on all action to take.

- No to All: The software denies all actions.
- Yes to All: The software confirms all actions.

> **Note:** By default, all operating system update plans use the **yes to all** policy.

You can create customized policies that answer differently depending on the specific OS update component. For example, when you want to review the questions and manually supply answers for patch 121288, which is part of the Oracle Solaris 10 Recommended updates from February 23, 2012, you can choose to set a specific policy for that patch, as shown in Figure 14–4.

Figure 14–4 OS Update Policies and User Defined Policy Details

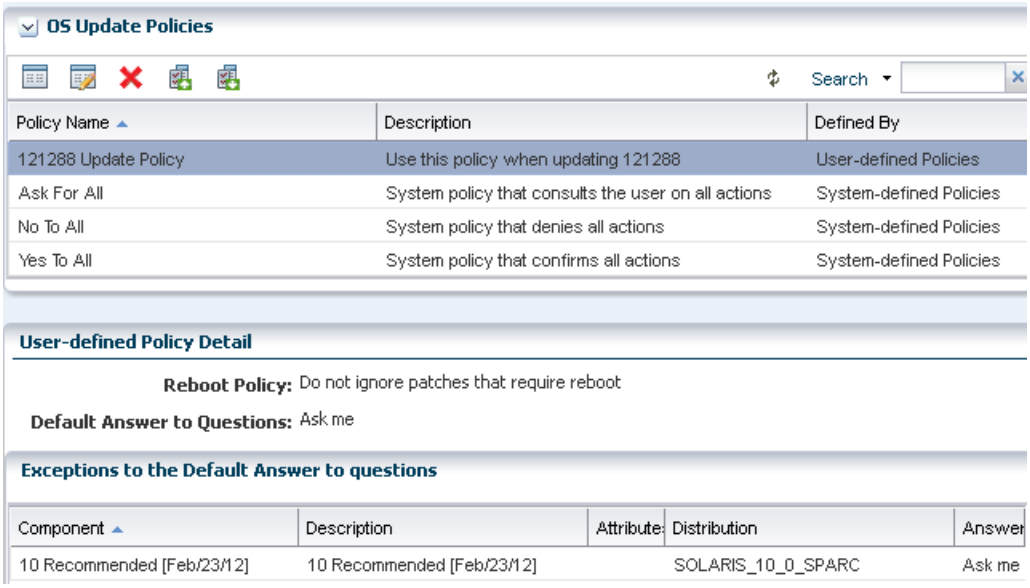

Policy settings are hierarchical. When there is not a policy setting for a component, the policy for that component's parent applies. For example, it is possible to create a policy that allows the system to install a given component but prohibits installation of certain specific versions of that component.

> **Note:** The policy only applies to actions that are implicitly generated by the dependency resolver. If a conflict occurs between a profile and policy, the profile overrides the policy.
>
> The update policy is not applicable when updating Microsoft Windows operating systems.

To Create a User-Defined Policy

1. Expand **Plan Management**, then click **Update Policies** in the Navigation pane.
2. Click **Create Policy** in the Actions pane.
3. Identify the policy by providing a name and description.
4. Click the **Available Packages / Patches** tab to view all available packages and updates. Use the search criteria to refine the list:

- Distribution: The operating system distribution, such as Oracle Solaris 9 for SPARC, Oracle Solaris 10 for SPARC, or Oracle Solaris 10 for x86
- Category: Package category, such as Cluster, Recommended Software Configuration, or Hardware
- View: View all, available, withdrawn, modified packages/updates, or reboot
- Version: View All Versions or filter for the Latest Versions

Figure 14–5 Create OS Update Policy

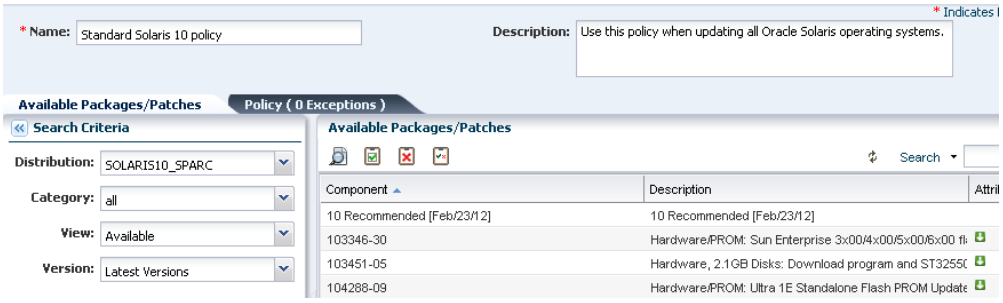

5. Select the component, then click an icon for the type of policy you want for the component, either **Answer Questions Yes**, **Answer Questions No**, or **Answer Questions Ask User**.
6. (Optional) To view details of the component, click the **View** icon.
7. Click **Create OS Update Policy**.

Creating Update Profiles

An update profile defines the component configuration of the systems that you want to manage. Update profiles specify which components are to be installed and which are prohibited, and any additional actions to be performed on an Oracle Solaris or Linux OS.

Use profiles to accomplish the following:

- Manage multiple systems in a consistent manner
- Automate repetitive administration jobs
- Record the requirements of your enterprise
- Automatically configure servers and workstations
- Manage dependencies and ensure consistency

The profile settings Required, Not Allowed, and Upgrade affect a managed host only during the actual deployment of that profile. At any time, you can run a job that contradicts the settings of a previously used profile; therefore you must understand your system settings and requirements thoroughly.

Predefined profiles are provided to perform common system-wide checks and to automate the operating system updates. These profiles cannot be edited or deleted.

You identify the profile type when you create the profile. The profile type is a tag that filters the required profiles when you create a deployment plan.

The following are profile types:

- Install: Indicates that new components are added, or installed. Use the Install profile type for Oracle Solaris operating system updates, baselines, and patchclusters.
- Upgrade: Indicates that existing components are upgraded.
- Script: Indicates that action scripts are executed.

> **Note:** You can create profiles that perform all of the actions for the profile type. The profile tag filters the required profiles in deployment plans. See About Complex Deployment Plans for more information.

To Create a New Profile

1. Expand **Plan Management**, then click **Update Profiles** in the Navigation pane.
2. Click **Create New Profile** in the Actions pane.
3. Enter a profile name and brief description of the profile.
   a. Enter a name and description for the profile.
   b. Select a Profile Type tag, either **Upgrade**, **Install**, or **Script**, to categorize and filter the profiles later.
   c. Select a distribution from the drop-down list. For example, SOLARIS10_SPARC.
   d. (Optional) You can further define the criteria by choosing a category, view, and version from the drop-down lists.

Figure 14–6 OS Update Profile Search Criteria

Figure 14–7 Create an OS Update Profile

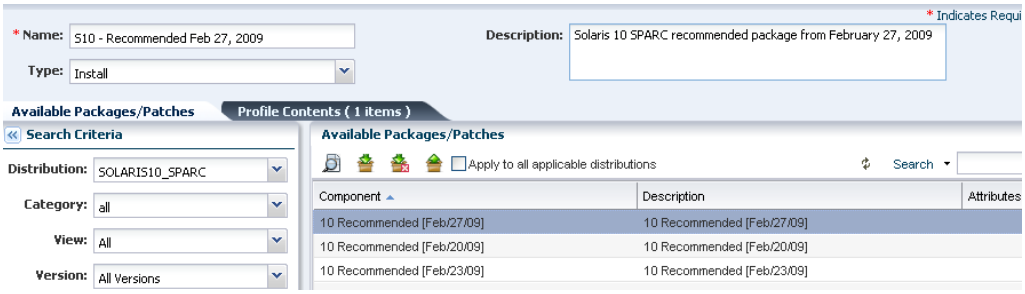

4. Locate and select a Component from the Component tree.
5. If required, select the check box to specify that the component is added to all applicable distributions.

> **Note:** This only applies to distributions that are active at the time the profile is created. As new distributions are activated you must edit the profile to explicitly add any components for those distributions.

6. Specify whether the action is **Required**, **Upgrade**, or **Uninstall**.

> **Note:** Some actions might not apply. For example, a component cannot be Required if the system does not have the information about how to obtain the component.

7. (Optional) You can repeat the preceding actions to select multiple components for the same or different operating systems.
8. Click **Save as Named Profile**. When an existing profile has the same name, you are asked to confirm that you want to replace the profile.

> **Note:** You cannot replace system-defined profiles.

Updating Oracle Solaris 11 Operating Systems

Oracle Solaris 11 uses a different update mechanism than earlier versions of the operating system. Oracle Solaris 11 uses packages to update the operating system and any non-global zones. The packages are part of an Image Packaging System (IPS) that is integrated with the ZFS file system.

When you install Oracle Enterprise Manager Ops Center on an Oracle Solaris 11 operating system, you create a local software package repository. The repository, called the Oracle Solaris 11 Software Update Library, is either on a file system on the same system as the Enterprise Controller, or on an NFS server share that the Enterprise Controller can access. You will populate the local library with packages from the parent repository instead of the Knowledge Base.

> **Note:** To update an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must be running on an Oracle Solaris 11 operating system. When Oracle Enterprise Manager Ops Center is not running on Oracle Solaris 11, the Oracle Solaris 11 update actions are not available.

The ZFS integration automatically creates an alternate boot environment every time an operating system is installed or updated. You can quickly and easily create an alternate boot environment when needed, and manage existing boot environments. Using an alternate boot environment provides a safe method of testing an update before deploying it to your live environment.

Unlike earlier versions of Oracle Solaris, you cannot run an ad-hoc operating system update job or browse package contents for Oracle Solaris 11. The best method of updating an Oracle Solaris 11 operating system is with a deployment plan.

To upgrade the operating system to the latest version of the Oracle Solaris 11 package, choose **Upgrade all components**.

Before you run an update job, verify that the Oracle Solaris 11 Software Update Library contains the package, or latest version of the package, that you want to use.

See Libraries for Oracle Solaris 11 for more information.

> **Note:** An Oracle Solaris 11 OS update job installs the latest available version of Oracle Solaris 11. An Oracle Solaris 11 OS provisioning job will install a specific Oracle Solaris 11 SRU version.

You will complete the following actions to update your Oracle Solaris 11 operating system:

1. Create an Oracle Solaris 11 update profile.
2. Optionally, create operational profiles that contain scripts to perform preinstall and postinstall actions.
3. Create an Oracle Solaris 11 update plan.
4. Select the target asset, then select the Deploy/Update Software action.

All update plans use the **yes to all** policy by default. If you do not want to automate the update by answering yes to all questions, you can create one or more customized update policies.

Create an update profile to define the images to use. After you have created the profile, you can create a deployment plan and choose the update and operational profiles to perform the tasks that you want for that plan. You can copy and edit plans to create customized plans for different purposes or targets.

An Oracle Solaris 11 update deployment plan provides a framework of steps and actions that you can perform to update your Oracle Solaris 11 operating systems.

The Oracle Solaris 11 Update deployment plan is a multi-step plan. The Update OS step is required, all other steps are optional. You cannot add steps to this type of plan.

Figure 14–8 Oracle Solaris 11 Deployment Template

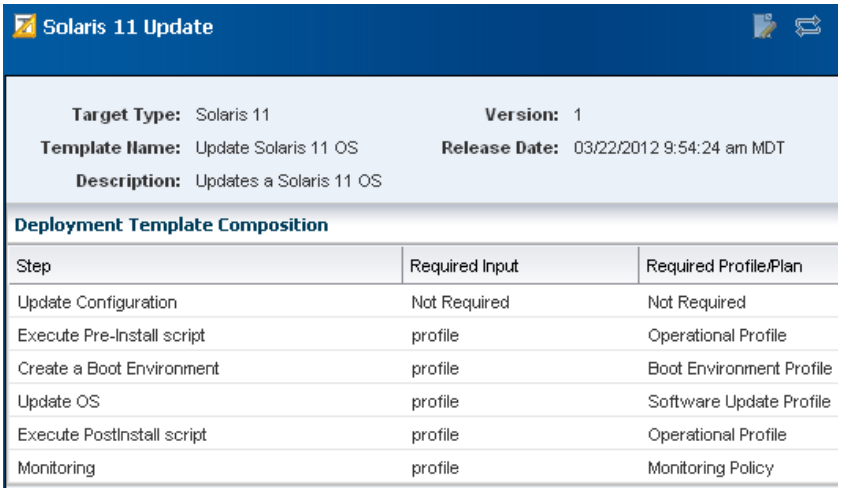

The Oracle Solaris 11 Update deployment plan includes the following steps:

1. Update Configuration: A profile is not attached to this step. The configuration is not required. At runtime, you can choose the type of update job, either an update simulation or an actual job.
2. Preinstall script: Optionally, you can associate an operational profile that contains a script that performs an action before you apply the update.

3. Create an alternate boot environment: You can create an alternate boot environment before the update job. This is optional.
4. Update the operating system: The profile that applies the update packages.
5. Post install script: Optionally, you can associate an operational profile that contains a script that performs an action after you apply the update.
6. Monitoring: Optionally, you can enable monitoring.

You can choose a failure policy for the plan, either to stop the update when a failure occurs, or to complete as many of the steps as possible.

Updating Oracle Solaris 8, 9, and 10 and Linux Operating Systems

You will use a similar methodology and set of procedures to update your Linux and Oracle Solaris 8, 9, and 10 operating systems.

You can use the following update methods and options to update these operating systems:

- Use predefined or custom profiles and associated deployment plans to update a system or group of systems.
- Use a system to create a simple update job without creating a profile. Use this method to apply a single patch quickly.
- Use the compliance reports output to update your OS. Use this method to make your systems compliant with newly released updates.
- Use compare catalogs to roll a system back to its previous state.

Table 14–2 Methods of Creating an Operating System Update Job

Update Method	Oracle Solaris 8, 9, and 10	Linux
Create and deploy an Update Plan	Yes	Yes
Create a new Update OS Job	Yes	Yes
Create Update Profile and Policy	Yes	Yes
Modify or compare a System Catalog	Yes	Yes
Create an Oracle Solaris Update Compliance report	Yes	No
Create a Compliance Report	Yes	No

In addition to the methods described previously, you can use Live Upgrade and alternate boot environments to update your Oracle Solaris OS with a minimum of downtime.

About Operating System Update Jobs

Creating a new update job enables you to use custom or predefined profiles. Use this method for complex update scenarios or to apply updates consistently across many systems. You can run the OS update plan in simulate or deploy mode. In simulate mode, you can choose whether to download the updates.

The New Update OS Job option enables you to create customized update jobs. When creating a job, you define how the software performs the job, set the automation level of the job, and select a policy from the list of available policies. You can run the job in simulation mode or run the actual job. Simulation mode determines the actions and

results of a job, and estimates the amount of time required to complete the job. You can use a job simulation to determine if your job can succeed based on your policy and profile responses. You can run a simulation with or without downloading updates.

Note: To use an alternate boot environment (ABE) and run ABE pre-action scripts for Solaris OS, see the procedures in Oracle Solaris Boot Environments.

Updating an Operating System From a Deployment Plan

Deployment plans enable you to control how the update is performed and apply updates consistently across many systems. With this method, you can use custom or predefined profiles to perform complex update scenarios or to apply updates consistently across many systems.

You can use the following deployment plan templates to update a supported operating system:

- Software Deployment/Update: Use this plan to apply script based update profiles.
- Configure Server Hardware and Install OS: Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- Configure and Install Dynamic System Domain: Use this plan to create dynamic system domains, provision and update OS on the domains.
- Install Server: Use this plan to provision and update the OS.

You will complete the following actions to use plans to update your Oracle Solaris or Linux operating system:

1. Create an OS update profile.
2. Optionally, create operational profiles that contain scripts to perform preinstall and postinstall actions.
3. Create a deployment plan.
4. Complete the plan, select the target asset, then select the Deploy/Update Software action.

The settings and values in the profiles bound to each step are defaults. You can modify the settings and values when you apply the plan. The profile settings and values are constrained by the target systems to which the plan is applied. All update plans use "yes to all" policy.

Updating an Operating System by Modifying a System Catalog

A system catalog contains a list of operating system software components that are installed on a managed system. Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems.

Updating an operating system by modifying a system catalog provides the following advantages:

- Enables you to create a quick ad hoc job
- Provides an easy method of applying a single patch, baseline, or package
- Enables you to update an operating system without creating a profile for a one-time job

Updating an Operating System From an Operating System Report Result

You can generate compliance reports for an operating system from which you can create an operating system update job.

The Oracle Solaris Update Compliance report is similar to a Recommended Software Configuration report. The report uses the Oracle Solaris update patch bundles as the recommended software configurations. You can use this report to check how compliant a system is with a particular Oracle Solaris update and bring the operating system into compliance.

See Chapter 10, "Reports" for information about generating these reports. You can generate these reports for non-compliant components. The report result appears with the option to install the updates, packages, updates, and incidents.

The report results are stored in the database associated with the Enterprise Controller. The software maintains a history of the reports for analysis purposes.

From the report result, you can initiate a job to install the non-compliant component updates. The New Update OS Job Wizard starts, enabling you to enter job information and to schedule the job. The required data for profiles, policies, and targets are automatically pre-populated in the New Update OS Job Wizard.

Note: Updating a version of Oracle Solaris is not the same as upgrading to a new version.

For example, you can run the Oracle Solaris Update Compliance reports for Oracle Solaris update releases and bring systems into compliance with those bundles. Oracle Solaris update release bundles contain the equivalent set of updates to the corresponding update. You can use them to bring pre-existing packages up to the same software level as the corresponding update. However, this feature does not perform a full Oracle Solaris upgrade from one release to another. The update release bundles do not contain additional packages that are in the update releases and they do not change the first line of `/etc/release` to specify an upgrade has taken place, although they do append a line to `/etc/release` to specify that the update bundle is applied.

Using a System Catalog

A system catalog is a list of operating system software components that are installed on a particular managed system. An initial catalog is created after the system is discovered and managed.

After an operating system is available and selected, you can view and modify the catalogs and create historical catalogs (snapshots of the system).

Modifying a catalog is an alternate way to run an operating system update job to install, uninstall, or upgrade a component. Modifying a catalog does not require an update profile to run the update job and is a quick way of changing the component configuration of a system.

You can compare the system catalogs of two managed systems, view the summary of the comparison, and you can choose to make the target system the same as the source system.

Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems. Alternatively, a catalog can be saved as a profile, and then an operating system update job can be run using this profile.

You can run an operating system update job, or you can use the simulate feature to run an update simulation before you apply updates.

Create an Update Profile From a System Catalog

You can save the catalog of a system as a profile. Using this profile, you can create the systems with the required configuration in your data center.

Updating Linux Operating Systems

Use profiles and associated deployment plans to update Linux on a single system or on a group of systems, use the New Update OS Job option to create a customized update job, or use a report to update your Linux operating system. You can use the Host Compliance Report or System Catalog Report to update your Linux operating system.

See Chapter 10, "Reports" for information about the types of reports available for Linux. See the *Oracle Enterprise Manager Ops Center Certified Systems Matrix* for a list of supported Linux operating systems.

Custom or predefined profiles enable you to perform complex update scenarios and to apply updates consistently across many systems.

The following deployment plan templates are available for updating a Linux operating system:

- Software Deployment/Update: Use this plan to apply script-based update profiles.
- Configure Server Hardware and Install OS: Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- Configure and Install Dynamic System Domain: Use this plan to create dynamic system domains, provision and update OS on the domains.
- Install Server: Use this plan to provision and update the OS.

Complete the following actions to use plans to update your Linux operating system:

1. Create an OS update profile.
2. Optionally, create operational profiles that contain scripts to perform preinstall and postinstall actions.
3. Create a deployment plan.
4. Complete the plan, select the target asset, then select the Deploy/Update Software action.

The profile settings and values for each step are the default settings and are constrained by the target systems to which the plan is applied. All update plans use a **yes to all** policy. You can modify the settings and values when you apply the plan.

Creating a new update job enables you to use custom or predefined profiles. Use this method for complex update scenarios or to apply updates consistently across many systems. You can run the OS update plan in simulate or deploy mode. In simulate mode, you can choose whether to download the updates.

To create a customized update job, use the New Update OS Job option. When creating a job, you define how the software performs the job, set the automation level of the job, and select a policy from the list of available policies. You can run the job in simulation mode or run the actual job. Simulation mode determines the actions and results of a job, and estimates the amount of time required to complete the job. You can

use a job simulation to determine if your job can succeed based on your policy and profile responses. You can run a simulation with or without downloading updates.

Determining if the Latest RPM Package Manager is Installed

Use the Service Pack Compliance report to help you to determine if a system is compliant with a service pack. The report provides information on updates created by the publication and release of a service pack, enabling you to determine whether the target system has the latest service package installed. For example, you can run the Service Pack Compliance report and identify required security updates based on packages that have a security flag.

The report results identify whether the operating system is compliant with a specified package. If the operating system is not compliant, the results suggest a course of action. When you have a package that is older than the selected pack, the report will recommend that you upgrade to that package. When the operating system has a package that is newer than the selected package, the report will recommend that you downgrade to the package you selected from the Services menu. For example, if you installed a newer kernel version, the report results recommend downgrading to the version that is in the package you selected.

To Determine if the Latest RPM is Installed

1. Expand **Reports**, then click **Additional Reports** in the Navigation pane.
2. Click **Service Pack Compliance Report** in the Actions pane.
3. Type a report name and, optionally, a description.
4. Select **Not Compliant**.
5. Select an RPM package from the Services list, then click **Next**.

Figure 14–9 Service Pack Compliance Report

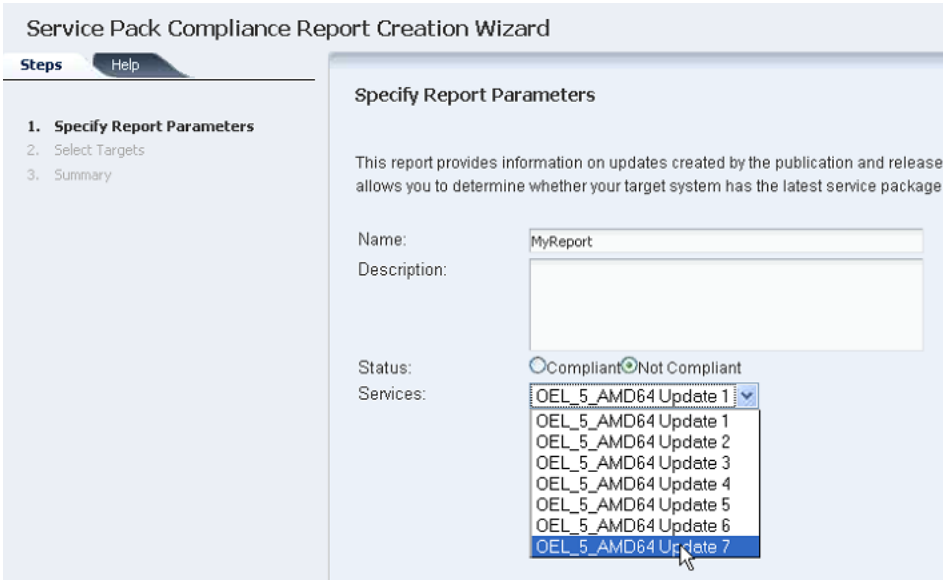

6. Expand Servers in the Select Targets pane and select a Linux operating system. Click **Add to Target List**, then click **Next**.
7. Click **Run Report**.

Uploading RPMs

Oracle Enterprise Manager Ops Center automatically uploads RPMs when a New OS Update job is launched, or when you choose to run an OS Update job simulation and select the option to download the patches. Outside of an update job, you can use the Bulk Upload Packages and Patches option to upload all RPMs from a DVD or from the YUM repository. See Uploading Software in Bulk for more information. To upload a single RPM, use the Upload Local Software option, as described in Uploading a Local Software Package.

Updating an Oracle Solaris Boot Environment

You always need an update profile to update an Oracle Solaris Boot environment. You can use the update profile in an update deployment plan or the Update Job Wizard.

You can update an alternate boot environment as part of a Software Deployment / Update deployment plan by selecting the alternate boot environment as the target. See the *Oracle Enterprise Manager Ops Center Updating Your Oracle Solaris 10 Operating System* for an example of how to use this plan.

You can create a customized update job, including the option to use an alternate boot environment (ABE) to perform a live upgrade of your Oracle Solaris 10 operating system. With Live Upgrade, you create an inactive ABE, update and patch the ABE, synchronize the ABE and BE, and then switch boot environments. When you switch boot environments, the patched and tested ABE becomes the active boot environment.

Note: Do not use Live Upgrade on your Enterprise Controller or Proxy Controllers. Live Upgrade does not synchronize all of the files that are required for these components.

You must run a separate update job for systems that use an ABE from those that do not use an ABE. When creating a job, you must define the following job parameters:

- Name and description of the update job.
- Alternate Boot Environment: Whether to use an alternate boot environment.
- Profile: Defines what updates are to be installed, uninstalled, or updated on an operating system. Select a profile from the list of predefined and customized profiles.
- Policy: Defines how a job is performed and sets the automation level of the job. Select a policy from the list of available policies. You can also create your own policies.
- Target Settings: Defines whether the target is different or similar for each task in the job.
- Actual Run: Defines whether this job is in simulation mode. You can choose to deploy the job, or to run a job simulation. A job simulation determines the actions and results of a job, and estimates how much time is required to complete the job. A job simulation also indicates whether your policy and profile responses will enable the job to succeed.
- Task Execution Order: Specifies whether the tasks is run in parallel or sequentially.
- Task Failure Policy: Specifies the action to take if a task fails.
- Targets: Select one or more target hosts for this job.

To create an ABE as part of this job, you must write at least one script that uses the `lucreate` command and then upload the script to the Local Content.

> **Note:** The ABE name defined in the script must match the ABE name that you use when you run the update job to create the ABE.

To Update a Boot Environment

Perform the following steps to update a boot environment:

1. Click **Assets** in the Navigation pane.
2. Expand All Assets, or use the All Assets filter to locate the Oracle Solaris 10 operating system instance.
3. Click **New Update OS Job** from the Actions pane. The New Update OS Job Wizard is displayed. The Job Information window is displayed first.
4. Complete the following Job parameters:
   - Type a job name.
   - Select the Run Type:
     - Simulation. To download the required updates as part of the simulation, select the Download check box.
     - Actual Run. Updates the operating system.
   - Select the task execution order:
     - Sequential
     - Parallel
   - Choose the Target Setting:
     - Use the same Targets for all tasks in the job
     - Use different Targets for each task in the job
   - Choose the Task Failure Policy:
     - Complete as much of the job as possible
     - Stop at failure and notify
   - Select the ABE check box.
   - (Optional) To create an alternate boot environment during this job by running an ABE Pre-Action Script, click the Enable check box.

   > **Note:** You must create the script and upload it to the library before you can use this option.

5. Define the profile, policy and target for each task, or edit the profile and policy.
6. (Optional) To edit the profile or policy of the default task, click the Profile or Policy cell for the task to display a drop-down menu. Select the profile or policy from the menu.
7. (Optional) To add a new task, click the **Add** (+) icon.
   - A second row appears. Click the Profile cell for that row to display a drop-down menu. Select the new profile that you want to add.

- To change the policy for the new profile, click the Policy cell and select a new policy from the drop-down menu.
- When you chose the parameter to use a different target for each task, click the Targets cell to display the Select Targets page. Select one or more target from the list of Available Items, then click Select to include the asset in the Target List. Click **Add to Target List** to close the page.
- Click **Next**.

8. If you selected the option to create an ABE as part of the job, the Create ABE page appears.
9. When you have only one ABE, the Boot Environment Workflow page appears, go to step 10. When you have multiple alternate boot environments, the ABE Selection page appears.
   - One or more of the targets has more than one possible associated ABE. Select the ABE from the drop-down menu for each of the Targets. You can use the **Select ABE** field to filter for the ABE name.
   - Click **Next**. The Boot Environment Workflow page is displayed.
10. If you selected Simulation in the job parameters, the boot environment workflow cannot be edited. Skip to step 12.
11. If you selected Actual Run in the job parameters, you can edit the pre-actions and post-actions in the workflow.
    - Pre-actions by default will unmount and then mount the ABE. To synchronize the ABE with the BE before mounting, click the **Sync ABE** check box.
    - Post-Actions by default will unmount the ABE.
      – Click **Modify Current BE** to edit the description of the current boot environment. You might use this to describe the state of the current BE. For example, Boot environment running Oracle Solaris 10 5/08 operating system before applying the Oracle Solaris 10 operating system September baseline.
      – Click **Modify Alternate BE** to edit the description of the ABE. You might use this to describe the state of the ABE. For example, boot environment running Oracle Solaris 10 5/08 operating system after applying the Oracle Solaris 10 operating system September baseline.
    - Click **Activate and Reboot ABE** to switch boot environments after update.
12. Schedule the job, then click **Next**.
    - Run Now starts the job immediately after you click Finish in the Job Summary.
    - Start Date enables you to select a date and time to start the job.
    - On a recurring schedule enables you to run the same job on a monthly or daily scheduled time.
13. Review the Job Summary, then click **Finish** to run the job as scheduled in the previous step.

Updating Microsoft Windows Operating Systems

You can update your managed Microsoft Windows operating systems by using the Microsoft System Center Configuration Manager (SCCM) and Windows Management Instrumentation (WMI) software.

Oracle Enterprise Manager Ops Center uses the Microsoft SCCM 2007 and WMI software to update your managed Windows operating systems. The Windows Update function depends on the SCCM's agent installed on the managed systems.You can configure SCCM to install agents on your managed Windows systems either automatically or through a manual process.

You must have access to SCCM software that is configured for software updates. The Enterprise Controller connects to the SCCM software to get the latest updates and packages. The SCCM software connects to the Microsoft website through the Internet and downloads the metadata that is used for compliance analysis. You can connect Oracle Enterprise Manager Ops Center to the Microsoft website to download updates that are then handled by SCCM for installation.

You do not need any authentication to access the Microsoft website. However, you must provide authentication information to access the SCCM server.

About Windows OS Update Jobs

Oracle Enterprise Manager Ops Center contains the following options in an update job to maintain control and consistency across your data center:

- Groups: Help you to organize your assets in the user interface and act as targets for many types of jobs.
- Roles: Determine the tasks that you can perform on a specific piece of an asset or a group of assets.
- Reports: Enable you to run compliance reports and create update jobs from the compliance reports.

You can define the following job parameters while creating a windows update job:

- Name and Description: Identify the name of the report against which you want to create a Windows operating system update job. Provide a detailed description that clearly identifies the job in the historical record.
- Reboot behavior: Enables you to select the reboot behavior when a reboot is required after running the new update job. You can choose to reboot the system immediately following the update operation or to reboot the system at the default setting of the SCCM server.
- License Terms: Enables you to review the license terms and either accept or decline them. The License Terms window appears only when the updates in the report require you to review the License Terms.
- Schedule: Enables you to schedule when the update job runs.

Modify the Registry

Due to some changes that Microsoft introduced for registry key ownership, you must manually modify the registry and change the ownership permissions for the Administrators group.

Note: You must modify the registry on a Windows Server 2008 R2. Other Windows servers, such as 2008 Server SP2, do not require you to modify the registry.

Configure Oracle Enterprise Manager Ops Center for Updating the Windows Operating System

Oracle Enterprise Manager Ops Center uses the DCOM wire protocol (MSRPC) to access the Windows Management Instrumentation (WMI) and get Windows update information. It uses the software update capability of the Microsoft System Center Configuration Manager (SCCM) to update any managed Windows operating systems.

Before you can use the software to update your Windows systems, configure it to interact with the identified Microsoft System Center Configuration Manager (SCCM). In addition, you might need to modify the WMI registry.

To configure Oracle Enterprise Manager Ops Center to interact with the identified SCCM, you must have the following credentials:

- SCCM Server
  - Server Name
  - Domain Name
  - Site Name
  - User Name
  - Password
- SCCM Share
  - URL
  - Domain Name
  - User Name
  - Password

The configuration information appears in the Configuration tab of the Windows Update window.

Note: Oracle Enterprise Manager Ops Center uses the same SCCM credentials to access the SCCM server and enable the SCCM share. Use the <*domain*> format for the Domain Name field. Do not use the <*domain*>\<*username*> format. Entering an incorrect format for the Domain Name field returns a configuration task. In this case, unconfigure the SCCM and configure the SCCM again with the correct format for the credentials.

Creating an Update Job for the Windows Operating System

You can use the output from compliance reports to update your Windows operating system to comply with the newly released updates.

From the results of the Windows Host Compliance Report and the Windows Incident Compliance Report, you can make your systems compliant by initiating an update job for the Windows operating system.

Figure 14–10 Process for Updating Windows Operating System

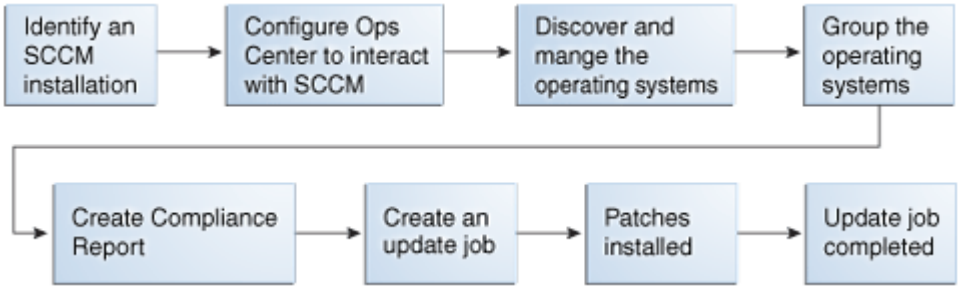

The Create New Windows Update Job Wizard enables you to create an update job. When creating a new update job, you must define the following job parameters:

- Name and Description for the new Windows software update job.
- Reboot behavior: Lets you select whether you want the system to reboot immediately following the update operation or at the default setting of the SCCM server.
- License Terms: Lets you review the license terms and either accept or decline them. The License Terms window appears only when the updates in the report require license terms that must be reviewed.
- Schedule: Lets you decide how you want to schedule the execution of the new update job.

Related Resources for Operating System Updates

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- See Chapter 5, "Software Libraries" for how to add and update operating system packages and images.
- See Chapter 12, "Operating System Management" for details on managing operating systems, including Using Agent Management for Operating Systems, Overview of Oracle Solaris Boot Environments, Overview of Oracle Solaris 11 Boot Environments, and Overview of Oracle Solaris 10 Boot Environments.
- See Chapter 10, "Reports" for details on the operating system Update reports.
- See Chapter 18, "Oracle Solaris Zones" for information about zones and how you can use Oracle Enterprise Manager Ops Center to efficiently manage all phases of zones lifecycle.

For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` and the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm`.

For in-depth information about these products, see the following Oracle documentation:

- For a list of the Oracle Linux documentation available in HTML and PDF formats, visit the Oracle Linux Documentation website at `http://www.oracle.com/us/technologies/linux/index.html`.
- *Transitioning From Oracle Solaris 10 to Oracle Solaris 11 Guide* at `http://docs.oracle.com/cd/E23824_01/html/E24456/docinfo.html`

- Oracle Solaris 11 Information Library at http://docs.oracle.com/cd/E23824_01/index.html
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html.
- For a list of the Oracle Solaris 8 and 9 documentation, visit the Legacy Solaris Documentation website at http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html.

Part IV

Virtualize and Cloud Management

Part IV contains the following chapters:

15

Getting Started with Virtualization

The following chapters in this section describe the virtualization and cloud management features in Oracle Enterprise Manager Ops Center:

- Chapter 15, "Getting Started with Virtualization"
- Chapter 16, "Storage Libraries for Virtualization"
- Chapter 17, "Networks for Virtualization"
- Chapter 18, "Oracle Solaris Zones"
- Chapter 19, "Oracle VM Server for SPARC"
- Chapter 20, "Oracle VM Server for x86"
- Chapter 21, "Server Pools"
- Chapter 22, "Virtual Datacenters"

Introduction to Virtualization

Virtualization technologies are designed to extend your hardware, operating systems, storage and network resources. Oracle Enterprise Manager Ops Center provides centralized management and optimization features for the following Oracle virtualization technologies:

- **Oracle Solaris Zones (SPARC and x86)**: Creates isolated, secure virtual operating systems, called zones, within an Oracle Solaris 10 or 11 operating system on either a SPARC or x86 platform. The default operating system controls the CPU, memory, and network resource allocation for the zones.
- **Oracle VM Server for SPARC**: Creates virtual machines, called logical domains or guests, within a single SPARC machine. Each logical domain has its own resources, such as a boot environment, CPU threads, memory, I/O devices, and its own operating system. Logical domains can run different operating systems. You can create zones within logical domains.
- **Oracle VM Server for x86**: Creates virtual machines within a single x86 platform. Virtual machines can run Linux, Oracle Solaris, and Windows guests.

You can use the virtualization management capabilities to create a comprehensive cloud management solution.

Review the following information to get started with virtualization and cloud management:

- Concepts for Virtualization
- Preparing for Virtualization

Concepts for Virtualization

As a virtualization administrator, you can create virtual operating systems, virtual systems, or a virtual data center in the cloud. All of these options require access to storage and networks. Information about core network and storage functionality is in Part II of the document, see Chapter 16, "Storage Libraries for Virtualization" and Chapter 17, "Networks for Virtualization" for more targeted information for virtualization technologies.

You should be aware of the following Oracle Enterprise Manager Ops Center concepts:

- Connection Modes
- Agent Controllers and VC Agent Controllers
- User Permissions and Roles for Virtualization
- Groups
- Server Pools

To take full advantage of the features in Oracle Enterprise Manager Ops Center, run the application in connected mode use agent-managed assets, groups, and server pools.

Connection Modes

The connection mode determines whether your instance of Oracle Enterprise Manager Ops Center is connected to the Internet. Plans and profiles that rely on images and packages have access to the versions that are stored in the library in Oracle Enterprise Manager Ops Center. In Connected mode, it is quick and easy to use the UI to add the latest images and packages from Oracle and vendor sites. If your site uses the product software in Disconnected mode, a manual process and scripts are available for you to update the images and packages in your local knowledge base. Alternatively, you can operate in Disconnected mode and then change to connected mode to update the images, then change back to Disconnected mode. As an administrator, you can easily and quickly change the connection mode. See the *Oracle Enterprise Manager Ops Center Administration Guide* for the procedure for changing the connection mode. See *Oracle Enterprise Manager Ops Center Using Disconnected Mode* for how to obtain updates while in disconnected mode.

Agent Controllers and VC Agent Controllers

Several types of agent controllers are available for you to manage your physical and virtual assets. The actions that you can perform are determined by the type of agent installed. See Virtualization Agent Controllers for more information.

User Permissions and Roles for Virtualization

The required roles for using a feature are listed in each chapter:

- Roles for Oracle Solaris Zones

- Roles for Oracle VM Server for SPARC
- Roles for Oracle VM Server for x86
- Roles for Server Pools
- Roles for Managing Virtual Datacenter

See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Contact your administrator to add roles to your user account. You must log out and then log in to see the actions for your revised role.

Groups

Oracle Enterprise Manager Ops Center uses groups to help organize and define a set of assets. An asset can belong to multiple groups. One of the main advantages of using groups is that you can apply profiles and plans to a group of like assets instead of individually.

When you discover an asset, the software automatically adds the asset in a system group based on the type of asset. For example, operating systems are automatically added to an Operating Systems group and to a group for that specific release, such as Oracle Linux, Solaris 11 or Solaris 10.

You can create your own user-defined groups and add assets to them. You can manually add assets to groups, or you can create rules that automatically add assets at discovery. See Using Groups for more information about groups and creating groups.

Server Pools

You can create a server pool for your virtualization hosts to share resources among all members of the pool. A server pool enables you to balance the load on the resources and provides high availability capabilities. To be a member of a server pool, one or more virtualization hosts must have the same processor architecture and have access to the same virtual and physical networks and storage resources. See Chapter 21, "Server Pools" for details on server pools.

Preparing for Virtualization

In preparing for virtualization, you need to know what types of virtualization you want to use and which features you want to use. You must create storage libraries to store the metadata and networks to assign to the zones or logical domains.

- Storage for Zones and Logical Domains
- Networks for Zones and Logical Domains

> **Note:** If you want to update Oracle Solaris 11 operating systems or create zones, the Enterprise Controller and Proxy Controllers must be running on a system that is running the Oracle Solaris 11 operating system.

Storage for Zones and Logical Domains

For zones, you must provide storage for the zone and zone metadata.

Zone data is the data that results from its operations. You can store zone data in a local library or a SAN storage library. For zone migration, store the zone data in a SAN storage library.

Zone metadata is the configuration of the zone's operating system, CPU, memory, and network. You can store metadata in a local library or in a NAS library. For zone migration, store the metadata in a NAS storage library.

When you create a zone, you assign it to one of the storage libraries associated with its virtual host. See Chapter 16, "Storage Libraries for Virtualization" for more details.

For Oracle VM Server for SPARC, you must provide storage for the logical domains. When you associate storage resources with Oracle VM Server for SPARC, the storage becomes available for the logical domains. A virtual disk server (vds), *primary-vds0*, is added to the Control Domain by default. The vds provides virtual disk service to the logical domains to access the storage disks that are not directly assigned to them.

Assigning a PCIe bus or PCIe HBA to a logical domain results in exclusive storage resource for the logical domain. A virtual disk server is also created in the logical domain so that it can provide virtual disk service to other domains.

Networks for Zones and Logical Domains

You can establish and manage different types of network infrastructure for Oracle Solaris Zones, Oracle VM Server for SPARC and Oracle VM Server for x86 virtualization technologies.

See Chapter 17, "Networks for Virtualization" for the requirements and how networks are connected or assigned to the virtualization hosts, virtual host, and the operating system. For core networking information, see Chapter 7, "Networks".

Virtualization Agent Controllers

A customized discovery profile identifies the assets and credentials needed to add assets to the user interface. Hardware assets are managed using a set of credentials. Operating systems and virtualization software are managed using an Agent Controller installed on the system or using only a set of credentials.

For an asset to be managed, it must be accessible on a network that is associated with a Proxy Controller. When you discover assets, you are given the option to manage them with or without an agent. Agent managed assets enables you to use Oracle Enterprise Manager Ops Center to monitor, update, and virtualize the assets. When an operating system is managed using only a set of credentials, it is considered agentlessly managed. Some features, such as analytics of your virtual environment, are not available when the operating system is managed agentlessly.

Agent Controllers run in two modes, basic and virtualization management. The basic Agent Controller provides support for monitoring and updating of the host operating system. The Virtualization Controller Agent, or VC Agent, supports basic and virtualization control for Oracle VM Server systems and Oracle Zones.

Two types of Virtualization Agent Controllers exist:

- **Zones Virtualization Controller Agent**: Manages the zones running on a global zone and manages zones that are running on the logical domains. Using this agent enables full zone monitoring and management actions. This is also known as the Zones VC agent.
- **Oracle VM Server Virtualization Controller Agent**: Manages the logical domains that are running on the Control Domain. The Oracle VM Server, Control Domain and operating system are reflected in the UI. Using this agent enables full monitoring and management actions for the Oracle VM Server system. This is also known as the Oracle VM Server for SPARC VC agent or LDom VC agent.

Note: Before Oracle Enterprise Manager Ops Center 12*c* Release 2, only the Zones Virtualization Controller Agent and the basic Agent Controller were available. If you managed an Oracle VM Server for SPARC in 12*c* Release 1 or earlier releases, it is managed with the Zones Virtualization Controller Agent. See Changing the Type of Agent Controller for how to change the controller to the Oracle VM Server VC Agent.

Beginning with the 12.2.2 release, when you use the **Add and manage various types of assets via discovery probes** option and select the option to agent manage a control domain or a global zone asset, the agent automatically discovers existing logical domains and non-global zones. Oracle Enterprise Manager Ops Center discovers the server and operating system of the control domain or global zone, then installs and configures the Virtualization Controller agent on the control domain or global zone. If the asset has logical domains or non-global zones, the agent discovers them and they appear in the list of assets.

In most cases, Oracle Enterprise Manager Ops Center deploys the Agent Controller needed. When there might be a conflict, you are prompted to choose which type of controller to install.

When you deploy the Agent Controller on an Oracle VM Server for SPARC system, either through OS discovery or by changing the management access, you are prompted to define which of the following types of Agent Controller to install:

- **Zone VC Agent**: The global zone is reflected in the UI. Using this agent enables full zone monitoring and management actions.
- **Oracle VM Server VC Agent**: The Oracle VM Server, control domain and operating system are reflected in the UI. Using this agent enables full monitoring and management actions for the Oracle VM Server system.

Note: You cannot create zones on a global zone that is managed agentlessly. For an agentless managed zone, you can boot, shutdown, halt and delete the zone. An Agent Controller is required on the zone when you want to use the full range of OS update actions on the operating system.

For robust management, use the Oracle VM Server VC Agent to manage the domains. When you have an Oracle VM Server VC Agent installed on a managed system, you can use Oracle Enterprise Manager Ops Center or the Oracle VM Server for SPARC command line to perform configuration operations.

Metadata for all managed logical domains is stored in the Oracle VM Server's default local library. The agent runs on the control domain and monitors the configuration and reflects any changes on the configuration in its copy of the metadata. The Oracle VM Server VC Agent synchronizes the logical domain configuration defined on the control domain with the domain model view in Oracle Enterprise Manager Ops Center.

The following operations are synchronized:

- Configuration assignment of I/O resources
- Network configuration when adding, removing, or updating virtual switches
- Storage configuration when adding, removing, or updating some Virtual Disk Services and Virtual Disk Devices

- Logical domain configuration when adding, removing, or updating the following guest configurations: CPU, memory allocation and policies, networking, storage, and I/O resources definition.

The Logical Domain can be an I/O Domain with the following I/O resources allocated to it:

- one or more PCIe root complexes (using PCI bus split)
- one or more PCIe End Point (using Direct I/O)
- one or more Virtual Function (using SR-IOV)

Changing the Type of Agent Controller

You can install the agent during discovery, or at any time after discovery. You have the following agent management options:

- **Oracle VM Server for SPARC Virtualization Controller Agent**: Manages the logical domains that are running on the Control Domain. The Oracle VM Server, Control Domain and operating system are reflected in the UI. Using this agent enables full monitoring and management actions for the Oracle VM Server system.
- **Zones Virtualization Controller Agent**: Manages the zones that are running on the logical domains. The global zone is reflected in the UI. Using this agent enables full zone monitoring and management actions.
- **Agentlessly**: Limited management functionality is available with this method. Information is gathered by using SSH connection between the logical domains and the Proxy Controller.

To switch between the different agent controllers, you must use the Management Access point, agentlessly manage the asset and then manage again. Figure 15–1 is an example flow of selecting an operating system that is agent managed with a zone VC agent, switching to agentlessly managed, then switching back to agent managed and selecting the Oracle VM Server VC agent.

Figure 15–1 Flow to Change the Type of Agent Controller

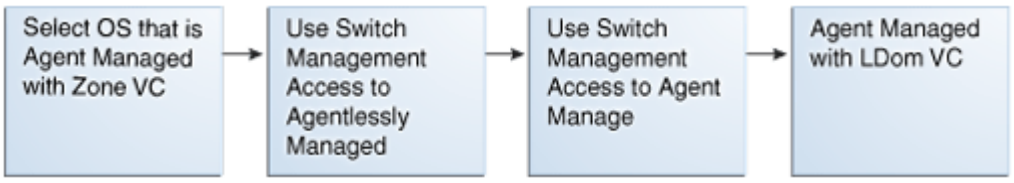

When you manage the asset, you are prompted to choose the type of Agent Controller when it is not obvious which type of agent you want to install.

To Change the Type of Agent Controller

1. Expand **Assets** in the Navigation pane.
2. Expand **Servers**, then select the operating system.

   The current management status appears in the Dashboard tab.

Figure 15–2 Agent Managed with an Oracle Solaris Zones Management Agent

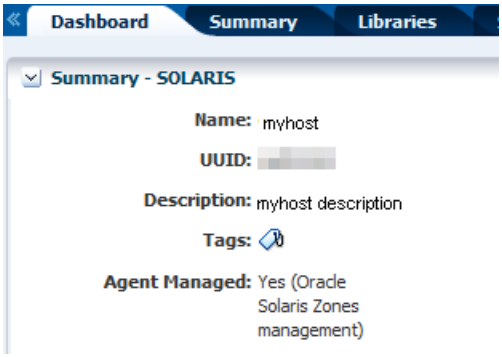

3. Click **Switch Management Access** in the Actions pane.
4. Add or select the credentials for the system, then click **Finish**.
   - (Optional) To create a new set of credentials, click **New** and complete the Create Credentials Wizard, then click **OK**.
   - (Optional) To select from a list of existing credentials, click **Select**, highlight the credentials from the list of available credentials, then click **OK**.

Figure 15–3 Switch Management Access

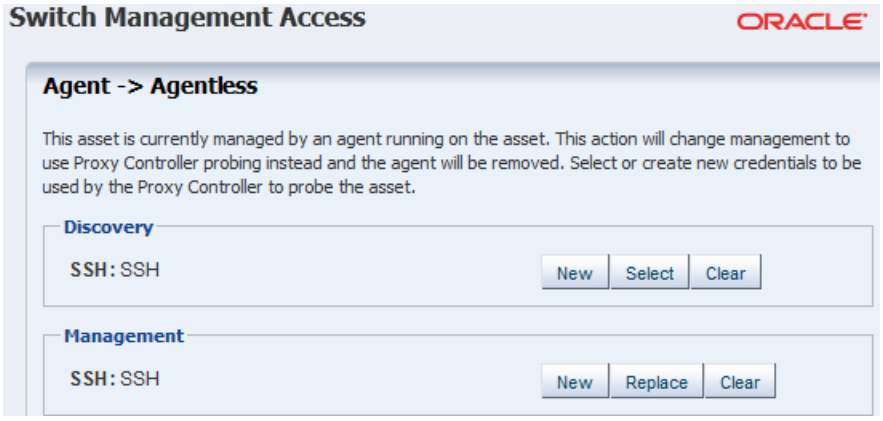

Wait for the job to finish. When the job finishes, the Asset is not agent managed.

5. Expand **Assets**, then **Servers** in the Navigation pane. Select the same operating system. The operating system is not agent managed. Click **Switch Management Access** in the Actions pane.

Figure 15–4 Switch Management Access

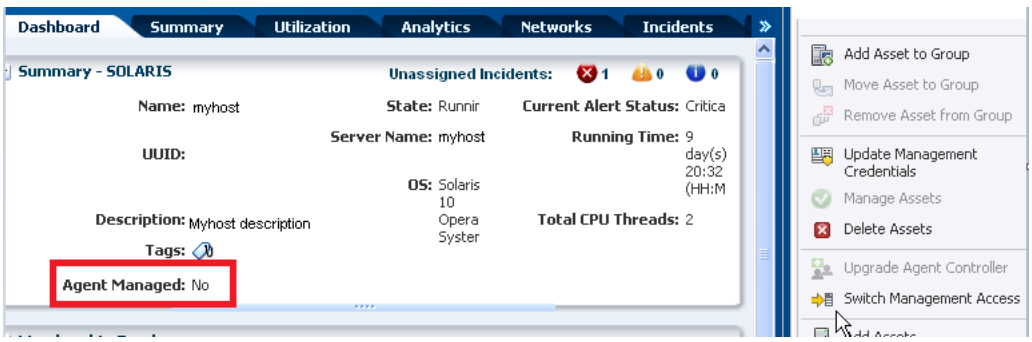

6. Select the check box to Enable Oracle VM Server for SPARC management to install that type of agent, then click **Finish**.

Figure 15–5 Switch Management Access From Agentless to Agent

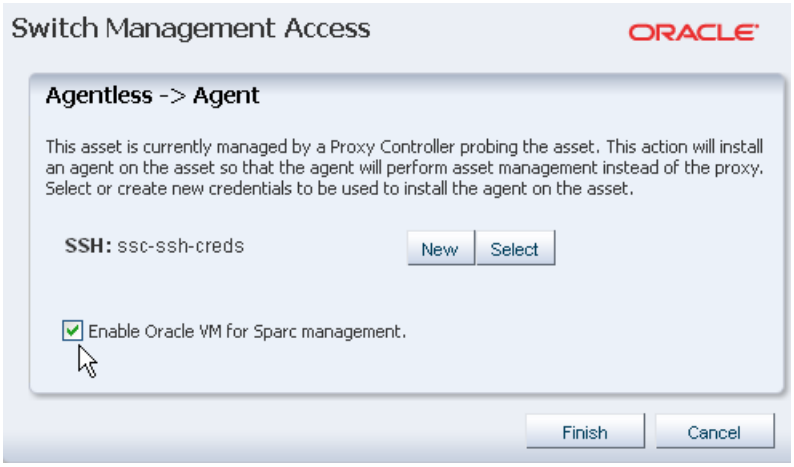

Wait for the job to finish. When the job finishes, the Asset is agent managed for Oracle VM Server for SPARC.

Introduction to Virtualization Management

Oracle Enterprise Manager Ops Center enables you to manages the lifecycle of virtualized operating systems and hardware, and provides centralized management of the virtualization infrastructure.

The following technologies are supported:

- **Oracle Solaris Zones**: Operating system virtualization for SPARC and x86 platforms.
- **Oracle VM Server for SPARC**: Hardware virtualization on a SPARC platform.
- **Oracle VM Server for x86**: Hardware virtualization on an x86 platform.

Operating system virtualization uses Oracle Solaris Zone partitioning technology to virtualize operating system services, and provide an isolated and secure environment for running applications. When you create a non-global zone, you produce an application execution environment in which processes are isolated from all other zones. This isolation prevents processes that run in a zone from monitoring or affecting processes that run in any other zones. See also global zone and non-global zone.

Hardware virtualization is a technology that creates multiple virtual systems on a single piece of physical hardware. When you create a hardware virtualized (HVM) guest, you must supply an ISO file in a repository to create the virtual machine. You can also virtualize hardware with paravirtualized drivers, PVHVM. PVHVM is identical to HVM, but has additional paravirtualized drivers for improved performance of the virtual machine. PVHVM improves the performance level of Microsoft Windows running in guests.

Oracle VM Server for SPARC is hardware virtualization technology that enables the creation of multiple virtual systems by a hypervisor in the firmware layer, interposed between the operating system and the hardware platform. This is designed to abstract the hardware and can expose or hide various system resources, allowing for the

creation of resource partitions that can operate as discrete systems, complete with virtual CPU, memory and I/O devices.

Oracle VM Server for x86 is hardware virtualization technology that runs on x86 platforms. It is a managed virtualization environment, or part of such an environment, that is designed to provide a lightweight, secure, server-based platform for running virtual machines. Oracle VM Server for x86 is based upon an updated version of the underlying Xen hypervisor technology, and includes Oracle VM Agent.

Roles for Oracle VM Server for SPARC Domains

All logical domains are the same and can be distinguished from one another based on the roles that you specify for them. The following are the roles that logical domains can perform:

- **Control domain**: The control domain is the first domain created when you install the Oracle VM Server for SPARC software. This is also called the primary domain and denoted as primary wherever applicable in the Oracle Enterprise Manager Ops Center UI. The Logical Domains Manager runs in this domain, which enables you to create and manage other logical domains, and to allocate virtual resources to other domains. You can have only one control domain per server.
- **Service domain**: A service domain has physical I/O devices and provides virtual device services to other domains. The following are examples of virtual device services: virtual switch, a virtual console concentrator, and a virtual disk server. In the Oracle Enterprise Manager Ops Center UI, the service domain list includes the primary domain, I/O domains, and root domains that can provide virtual device services. You can have more than one service domain, and any domain can be configured as a service domain. You should not run any applications in service domains.
- **I/O domain**: An I/O domain has direct access to a physical I/O device, such as a network card in a PCI EXPRESS (PCIe) controller. An I/O domain either uses the physical I/O devices to host its own applications or shares the physical I/O device with other domains in the form of virtual devices.

  An I/O domain can share physical I/O devices with other domains in the form of virtual devices when the I/O domain is also used as a service domain.
- **Root domain**: A root domain has a PCIe root complex assigned to it. This domain owns the PCIe fabric and provides all fabric-related services, such as fabric error handling. A root domain is also an I/O domain, as it owns and has direct access to physical I/O devices.

  The number of root domains that you can have depends on your platform architecture. For example, if you are using an Oracle Sun SPARC Enterprise T5440 server, you can have up to four root domains.
- **Guest domain**: A guest domain is a non-I/O domain that uses virtual services that are provided by one or more service domains. A guest domain does not have any physical I/O devices, but only has virtual I/O devices, such as virtual disks and virtual network interfaces.

Deploy and Manage Virtual Assets

Figure 15–6 is a high-level workflow of the tasks needed to deploy and manage Oracle Solaris Zones and virtual machines.

Figure 15–6 Workflow to Deploy and Manage Virtual Assets

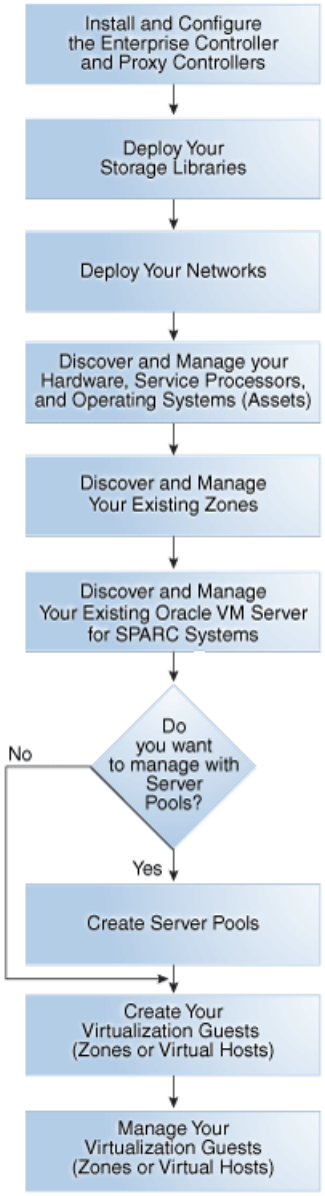

Note: When you have zones inside a logical domain, manage the Oracle VM Server for SPARC system *before* managing the zones.

You can discover and fully manage an existing Oracle VM Server for SPARC, one that you created without using the Oracle Enterprise Manager Ops Center software to provision the Control Domain. See *Adding an Existing Oracle VM Server for SPARC* in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` for the procedure to add an existing Oracle VM Server system to the user interface to begin managing the environment.

The software provides appropriate components to keep the physical and virtual systems up-to-date, such as latest firmware, packages, and patches. Oracle Enterprise Manager Ops Center reduces the complexity of deploying and maintaining a virtualization stack.

The following are some of the software features:

- **Asset Discovery**: Use discovery to locate hardware, service processors, operating systems and virtual systems in your data center and add the assets to the user interface.
- **Provisioning**: Provision operating systems on bare-metal systems and virtual systems.
- **Patching**: Update and upgrade your firmware and operating systems with the help of a comprehensive knowledge base. Patching is available for all of the components in the virtualization stack, including physical and virtual systems.
- **Monitoring**: Monitor the physical and virtual systems of the virtualization stack, including individual and aggregate resource utilization for a system.
- **Server pools**: Aggregate virtual resources to cater to specific purposes. Set the server pool policies for optimal use of the physical resources and to migrate virtual systems to another pool to reduce downtime or balance resources.
- **Migration**: Move virtual systems to different physical systems based on the resource usage, hardware failure, and other application requirements.

Introduction to Cloud Management

Oracle Enterprise Manager Ops Center's broad range of virtualization management capabilities work together to provide a comprehensive cloud management solution.

You can use the vDCs feature to enable the deployment and management process for a cloud-based infrastructure. A **cloud** is a defined set of physical resources that includes server pools backed by a virtualization infrastructure, storage, and networks. Use the UI to create, manage, and setup the clouds. The interface also provides options for a cloud user to create, run, and manage their virtual resources in the cloud.

Oracle Enterprise Manager Ops Center provides a cloud web service interface for Infrastructure as a Service (IaaS) functionality. The software also provides a Java Client API and a CLI to manage the cloud web service, enabling you to design applications and run instances on the cloud. See the *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide* for information on the API and CLI.

Deploy and Manage a Virtual Data Center

A virtual datacenter (vDC) is a consolidation of your physical resources that share the network and storage resources. A virtual data center has the following physical resources:

- Storage
- Networks
- Server pools

Server pools aggregate virtual resources for specific types of resources. Server pools enable you to optimize the use of physical resources. You can migrate virtual systems to another pool to reduce downtime or balance resources.

Figure 15–7 is a high-level workflow of the tasks needed to deploy and manage a virtual datacenter.

Figure 15–7 Workflow to Deploy and Manage a Virtual Datacenter

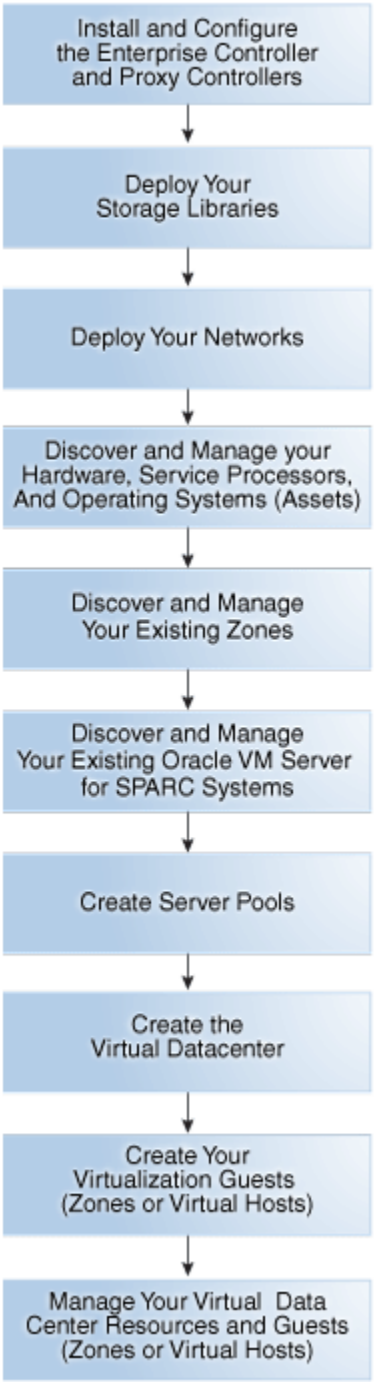

Related Resources for Getting Started with Virtualization

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- See the following chapters in the *Oracle Enterprise Manager Ops Center Feature Reference Guide*:
  - See Chapter 2, "Asset Management" for information about discovering and adding assets, and how to create and use groups.

  - See Chapter 16, "Storage Libraries for Virtualization" for information about the storage libraries you need for virtualization.
  - See Chapter 17, "Networks for Virtualization" for information about networks for virtualization.
  - See Chapter 18, "Oracle Solaris Zones" for information about managing incidents that result from monitoring.
  - Chapter 19, "Oracle VM Server for SPARC"
  - Chapter 20, "Oracle VM Server for x86"
  - Chapter 21, "Server Pools"
  - See Chapter 22, "Virtual Datacenters" for information about creating and managing a virtual datacenter.
- For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm and the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.
- See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant and for information on other administrative tasks.
- See the *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide* for information on the API and CLI.
- See *Oracle Enterprise Manager Ops Center Using Disconnected Mode* for how to obtain updates while in disconnected mode.
- For a list of the Oracle Solaris 11 documentation available in HTML and PDF formats, visit the Oracle Solaris 11 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-11-192991.html.
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html.
- For Oracle VM Server for SPARC, see http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html.
- See http://docs.oracle.com/cd/E19044-01/sol.containers/817-1592/ for Oracle Solaris Resource Management and Oracle Solaris Zones documentation.

16

Storage Libraries for Virtualization

The following information is included in this section:

Introduction to Storage Libraries for Virtualization

Oracle Enterprise Manager Ops Center manages and monitors storage servers and appliances, discovers and provisions storage capacity on these appliances through Storage Connect plug-in software, and then makes them available for use as guests storage.

Storage libraries are the storage resources for Oracle Solaris Zones, Oracle VM Server for SPARC, Oracle VM Server for x86, their server pools, and virtual datacenters. The storage libraries are used for storing the virtual host metadata and for virtual disk storage usage. The storage libraries must be associated with the virtualization hosts for guest storage usage.

Oracle Enterprise Manager Ops Center provides the ability to utilize the two following main types of storage:

- **Filesystem Storage**: File system based storage such as Network Attached Storage (NAS) Libraries and Oracle VM Storage Repository.
- **Block Storage**: Block-based storage arrays that support the Storage Array Network (SAN) protocols like Fibre Channel and iSCSI.

This chapter provides a brief description about different storage types, how they are represented in the UI and supported actions, and how they are associated with the virtualization hosts and used by the guests.

Discovering Storage Servers

Oracle Enterprise Manager Ops Center provides discovery profiles to discover your storage servers. If the storage vendor provides Storage Connect Plug-in for the storage server, then Oracle Enterprise Manager Ops Center checks for the plug-in and exposes the storage device features and attributes. With the discovery of the storage server, the exported file systems, LUNs, clones, and snapshots are populated and displayed in the UI. There are also periodic update of the storage server for any creation or deletion of exported LUNs, file systems, clones and snapshots.

For Oracle VM Server for x86, you must discover the storage resources from the Oracle VM Manager. To add storage to Oracle VM Servers, discover the storage resources from the corresponding Oracle VM Manager. Oracle Enterprise Manager Ops Center provides the UI options from the Oracle VM Manager to discover the storage resources such as Discover File System Storage, Discover SAN Storage and Discover iSCSI Storage. The same options are available from the discovery profile options.

The discovered servers are displayed under Storage in the Assets tree. You can select the storage and the center pane displays the details about the server.

You can manage storage servers like Oracle ZFS Storage Appliance and Oracle Exadata Servers. For Oracle ZFS Storage Appliance, there are more options to launch the appliance UI from the Oracle Enterprise Manager Ops Center UI, and to manage the shares and services of the server.

Oracle Exadata Server are discovered as part of discovering the SuperCluster Engineered Systems.

See Chapter 11 for detailed information about discovering, managing, and setting up storage hardware in Oracle Enterprise Manager Ops Center.

Storage Types

Oracle Enterprise Manager Ops Center provides support for the following storage types:

- **File System Storage**: The NAS storage libraries are based on the NFS file systems. You can create the storage library by providing the NFS mount point on the discovered storage asset or any other hosts.

  Oracle VM Storage Repository are based on NFS or OCFS (Oracle Cluster File System). The repository is especially for Oracle VM Server for x86 systems. Whereas, the NAS storage are used for all other systems. Create the repository by providing the Oracle VM Manager, Oracle VM Server, and the NFS file systems. For OCFS, you must select the LUNs from the list.

- **Block Storage**: Block storage libraries are the storage arrays that support the SAN protocols Fiber Channel and iSCSI. The block storage are also known as SAN libraries. The SAN libraries are groups of Logical Unit Numbers (LUNs). The LUNs supported by one or more storage arrays that are managed by Oracle Enterprise Manager Ops Center form the Dynamic Block Storage Library. In the dynamic block storage library, the user can create, edit or delete the LUNs.

  The LUNs supported by one or more storage arrays that are not managed by Oracle Enterprise Manager Ops Center form the Static Block Storage Library. Also,

for storage servers that are discovered and managed with existing LUNs form the static block storage as the user cannot create, edit or delete the LUNs.

- **Local Library and Local Devices**: The file system in the virtualization host are termed as local library. Local library is the default library that is displayed in selecting the storage library. The storage devices that are attached to the virtualization host are defined as local devices. You cannot manage the local devices. You can use to store the data on the local devices.

Using the Storage Libraries

You associate the storage libraries to virtualization hosts like Oracle Solaris Zones, Oracle VM Server for SPARC, and Oracle VM Server for x86. Server pool is a group of homogenous virtualization hosts. Associating a storage library to a server pool results in associating the storage library to all the virtualization hosts in the server pool.

The storage libraries are provided to store ISO images, guest metadata and for disk storage for the guest requirements. The storage libraries store the metadata of the guest such as zones, logical domains, and virtual machines. The metadata storage in the File System storage such as NAS enables to migrate the guests between compatible servers. You cannot store the metadata in SAN storage libraries. You use the SAN libraries only for disk storage purposes of guests.

Storage Library Setup

Prepare your storage resources to be accessed and used by the systems managed in Oracle Enterprise Manager Ops Center. Create appropriate storage libraries depending on the type of storage which are exposed as file system or block storage disks. You can then associate the libraries to the virtualization hosts during provisioning or assign them later.

See Chapter 16 and the *Deploy Storage Libraries Workflow* in the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` for procedures about creating storage libraries.

Roles for Storage Libraries for Virtualization

Table 16–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 16–1 Storage Tasks and Roles

Task	Role
Create Storage Libraries	Storage Admin
Associate Storage Libraries to Virtualization Hosts	Virtualization Admin or Storage Admin
Add Storage to Virtual Hosts	Virtualization Admin

Actions for Storage Libraries for Virtualization

You can perform the following actions, depending on the requirements:

- Create Storage Libraries

- Edit Library Attributes
- Associate Libraries
- Disassociate Libraries
- Create Local Library
- Remove Storage
- Enable Sharing
- Disable Sharing

Location of Storage Information in the User Interface

Table 16–2 shows where to find information.

Table 16–2 Location of Library Information in the BUI

Object	Location
To see storage libraries	Expand Libraries in the Navigation pane, then scroll to the Storage Libraries. Select the type of library.
To associate storage libraries to virtualization hosts	Expand All Assets in the Navigation pane, then select the virtualization host and click Associate Libraries in the Actions pane.
To add storage to virtual hosts	Expand All Assets in the Navigation pane, select the virtualization host and then the virtual host. Click Add Storage in the Actions pane.
To associate storage libraries to server pools.	Expand Assets in the Navigation pane, then Server Pools in the Resource Management Views. Select the server pool and click Associate Library in the Actions pane.

Storage Libraries for Oracle Solaris Zones

The file systems of the zones are implemented as ZFS file systems and a dedicated zpool is created for each zone. The storage allocated to the zone are pooled in the zpool and used by all the file systems of the zone.

For Oracle Solaris Zones, you can use the following types of storage libraries:

- Local Library and Local Devices
- NAS Storage Libraries
- Static Block Storage Libraries
- Dynamic Block Storage Libraries

Apart from the local library and local devices, associate the other libraries to the global zone. The libraries associated with the global zone are automatically available for all the non-global zones under it. The local storage is always available by default to all the zones.

Unmanaged Storage

When you discover and manage existing zone environments, the underlying storage of the zone is considered to be unmanaged and not recognized by Oracle Enterprise Manager Ops Center. You can use the option Move Storage option to move the zone metadata and storage to managed storage. See Moving Zone Storage for more information about the procedure.

Migration Capability

For migrating the zones, the non-global zone must have shared storage. It is required to store the non-global zone metadata in a shared storage such as NAS storage and the virtual disk storage on SAN libraries. Only then, the migration option is enabled for the zone.

For any unmanaged storage, you must use scripts for migrating the unmanaged storage from the source to the target global zone. See Migrating Zones for more information about using scripts and other requirements.

Storage Libraries for Oracle VM Server for SPARC

For Oracle VM Server for SPARC, you can assign the following storage libraries:

- NAS Storage Library
- Static Block Storage Library
- Dynamic Block Storage Library

The local storage libraries are available by default. The storage resources when associated with Oracle VM Server for SPARC becomes available for the logical domains. A virtual disk server (vds), *primary-vds0*, is created in the control domain by default. The vds provides virtual disk services to the logical domains to access the storage disks that are not directly assigned to them.

Assigning a PCIe bus or PCIe HBA to a root domain or I/O domain results in exclusive storage resource to the domains. A vds is created in the root domain and I/O domain so that it can provide virtual disk services to guest domains.

Virtual Disk Multipathing

You can provide redundant access to the logical domain storage. The alternate path to access the same back-end storage can be provided while creating logical domains or adding storage to logical domains. For each virtual disk of the logical domain, a multipathing group is created and you must specify a group of virtual disk servers (vds) of other domains as the alternate path to the back-end storage access.

A multipathing group is created only when there is more than alternate path to access the back-end storage. You must select an active alternate path when there is a failure in a service domain.

When there is only one alternate path to the back-end storage, then unless you specify a name for the group, the multipathing group is not created.

You can either enter a name for the multipathing or a name is created in the format of *logical domain name_mpGroup_devID*. The devID is the disk index. For example, if the name of the logical domain is *ldom*, then the multipathing group name is in the format of *ldom_mpGroup_1*.

When you want to migrate guest domains that have multipathing configured for the back-end storage, then the target server must also have the I/O domains or root domains that provide multipathing to the back-end storage. This is essential for live migration of the guest domains. To successfully migrate a guest domain, the target Oracle VM Server must have root or I/O domains that have the same virtual disk service name as that of the source Oracle VM Server.

For example, when *vds1* and *vds2* are the alternate paths for a guest domain virtual disk, then the target Oracle VM Server must also have the same virtual disk servers, *vds1* and *vds2*.

When you create guest domains, you can enable automatic recovery of the guest domains and authorize the recovery of the domains without I/O redundancy. This results in guest domains being recovered on servers when the server pool does not have any servers that do not have any I/O domains or root domains in them.

Opaque Storage

When you discover and manage logical domains created using CLI, the logical domain metadata is stored in the local storage library. The virtual disks of the logical domain are on the storage servers that are not managed by Oracle Enterprise Manager Ops Center, and they are defined as opaque storage. You must enable the storage disks as Shared to allow migration or automatically recover guests with opaque file systems.

Moving Metadata

When you have the logical domain metadata stored in local storage library, migration option is disabled. You must move the metadata storage from local storage to shared storage such as NAS storage library. Use the **Move Metadata** option to move the metadata of a logical domain to a shared storage.

Associating Storage Libraries

When you associate storage libraries with Oracle VM Server for SPARC, you have the option to select to which domain the library is associated. You can associate any storage library with the control domain, I/O domain or root domain provided the domains can access the library. The storage library that is associated with the domain adds to the virtual disk server (vds) of that domain. For example, when a FC card is installed only in the primary or the control domain, then the FC SAN library must be associated with the control domain only as it is the only domain that can access it. The FC SAN library storage adds to the control domain vds and enables the guest disks to use the storage.

Adding Storage to Logical Domains

You can add storage to the logical domains from the storage libraries associated with the server pool or the Oracle VM Server for SPARC. The control domain, I/O domains, and the root domains provide the virtual disk service to the logical domains. Use the option **Add Storage** to add storage to the logical domains.

Storage Libraries for Oracle VM Server for x86

Oracle Enterprise Manager Ops Center provides option to create Oracle VM Storage Repositories. This type of storage library stores virtual machine metadata, templates, assemblies, ISO images, and virtual disks for the Oracle VM Server for x86.

When you create the storage repository on a LUN, it is a block-based repository. When you create the storage repository on a NFS file server, it is a NFS-based storage repository.

When you discover an Oracle VM Manager, Oracle Enterprise Manager Ops Center provides the option to discover the storage servers attached to it. Create discovery profiles for the storage servers and save them.

When you select to discover a storage server, existing discovery profiles are searched for the selected storage type. You can select from an available profile. When there are no profiles available, then the wizard to discover the corresponding storage server appears.

When you discover a storage resource, existing resources such as virtual disks, templates, ISO images, and virtual machine metadata are also discovered and displayed.

The following types of libraries setup are supported for storage of Oracle VM resources:

- **File system Storage**: Oracle Enterprise Manager Ops Center provides an option to create Oracle VM Storage Repositories. A storage repository is a logical disk space made available through a file system on top of physical storage hardware. The supported types of file system are NFS and OCFS. When the storage repository is created on an NFS file server, it is a NFS based storage repository. The NFS file server consists of NFS file systems. When the storage repository is created on a LUN, it is a LUN-based repository. The OCFS file system is created on the storage server. Create storage repositories on these file servers to be used by Oracle VM Servers to store resources. The resources include virtual machine metadata, templates, assemblies, ISO images and virtual disks.

  To create Oracle VM storage repositories, you must have an Oracle VM Server discovered. Only NFS-based repositories can be shared by multiple server pools.

- **Static Block Storage Libraries**: The LUNs from the storage arrays that are not managed by Oracle Enterprise Manager Ops Center form the static storage libraries. The LUNs are addressed by iSCSI or Fibre Channel protocols. This forms the iSCSI and SAN static storage libraries. Also, add LUNs exported from managed storage arrays. The LUNs can belong to one library at a time.

- **Dynamic Block Storage Libraries**: The storage servers that are discovered and managed in Oracle Enterprise Manager Ops Center are presented under dynamic block storage libraries. The dynamic block storage libraries contains the exported LUNs from the storage array servers.

Storage Libraries for Virtual Datacenters

The virtual datacenter (vDC) inherits the storage resources allocated for the server pool. The virtualization type of the server pool defines all the possible types of storage resources available for the vDC.

When you create a vDC, you allocate the storage resources to be used for all the accounts created in the vDC. The allocated storage resources are used as root disks of the virtual servers (vServers) and volumes. Volumes are additional storages that you can attach to the vServers. For volumes, the storage is allocated from the storage libraries that are associated with the server pools. A cloud user can also import external volumes into the account.

The following storage resource that can be designated as volumes for vDCs based on different virtualization types:

- Oracle VM Server for x86 based vDC
  - Oracle VM Storage Repositories
  - Dynamic Block Storage Libraries
- Oracle Solaris Zone based vDC
  - NAS Libraries
  - Block Storage Libraries that include both Dynamic and Static
- Oracle VM Server for SPARC based vDC

  - NAS Libraries
  - Block Storage Libraries that include both Dynamic and Static

Related Resources for Storage Libraries

For instructions on performing actions or to learn more about the role of this feature, go to one of the following resources.

- For end-to-end examples, see the *Deploy Storage Libraries Workflow* in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm.

17

Networks for Virtualization

The following information is included:

Introduction to Networks for Virtualization

Oracle Enterprise Manager Ops Center provides extensive management support for your datacenter network infrastructure. It supports both Ethernet and InfiniBand network protocols. Network management in Oracle Enterprise Manager Ops Center provides a full lifecycle management of network domains, networks, and fabrics. It provides the following services:

- Discover and manage switches
- Define Ethernet fabrics
- Create and define networks
- Create and manage network domains
- Automatic network discovery during asset discovery
- Provision to create private networks on demand
- IP address allocation that includes reserve and release a subnet member
- DHCP management for host interface configuration
- UI support for complex network configurations of virtualization deployments

This chapter provides a brief description about the network management in Oracle Enterprise Manager Ops Center, the different types of network infrastructure that can be setup and managed for virtualization technologies like Oracle Solaris Zones, Oracle VM Server for SPARC and Oracle VM Server for x86.

The prerequisites and how networks are connected to or assigned to the virtualization hosts, virtual host and the OS are described in this chapter.

For more detailed information and procedures about networking, refer to the Networking chapter in *Oracle Enterprise Manager Ops Center Feature Reference Guide*.

Roles for Networks for Virtualization

Table 17–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 17–1 Network Tasks and Roles

Task	Role
Create Networks	Network Admin
Create Network Domains	Network Admin
Create Private Networks	Network Admin
Define Networks	Network Admin
Assign Networks	Network Admin
Define Ethernet Fabric	Network Admin
Add Fabric	Network Admin
Assign Fabric	Network Admin
Assign Network	Network Admin
Delete Network Domain	Network Admin
Remove Network	Network Admin
Delete Network	Network Admin
Remove Fabric	Network Admin
Delete Managed Fabric	Network Admin
Assign VLAN ID Range	Network Admin
Edit Attributes	Network Admin
Attach Network	Virtualization Admin
Connect Guests	Virtualization Admin
Associate Network Domain	Virtualization Admin
Detach Networks	Virtualization Admin

Actions for Networks for Virtualization

You can perform the following actions from the virtualization hosts, server pools, and guests:

- Attach Network

- Connect Guests
- Associate Network Domain
- Detach Networks
- Disconnect Guests from Network
- Modify Physical Connectivity

For other actions of network related, refer to Chapter 7, "Networks" for more information.

Location of Network Information in the User Interface

Table 17–2 shows where to find information.

Table 17–2 Location of Network Information in the BUI

Object	Location
To see fabrics	Expand Networks in the Navigation pane, then select Fabrics in the filter.
To see all network domains	Expand Networks in the Navigation pane, then select Networks in the filter. Network Domains are listed.
To see all network	Expand Networks in the Navigation pane, then select Networks in the filter. Network Domains are listed. Expand a Network Domain to view all networks.
To attach networks to virtualization host	Expand Assets in the Navigation pane and select the Virtualization host. Click Attach Network in the Actions pane.
To connect guests to network	Expand Assets in the Navigation pane and select the guest. Click Connect Network in the Actions pane.
To associate network domains to server pool	Expand Assets in the Navigation pane, then Server Pools in the Resource Management Views. Select the server pool and click Associate Network Domain in the Actions pane.
To attach networks to server pool	Expand Assets in the Navigation pane, then Server Pools in the Resource Management Views. Select the server pool and click Attach Networks in the Actions pane.

Manage Networks

Networks are managed in the following way in Oracle Enterprise Manager Ops Center:

- Discovering an asset automatically discovers the network.
- Discovering and managing a switch automatically discovers all the fabrics in the switch.
- Defining the Ethernet fabrics and assigning VLAN IDs to the fabrics.
- Defining networks by providing network address, gateway, VLAN ID, fabrics, and network services.
- Creating networks by providing network address, gateway, fabric, and network services. Oracle Enterprise Manager Ops Center automatically allocates the VLAN IDs or P-Keys.
- Creating network domains that are administrative containers for networks. The network domains handle the relationship between the physical fabrics and networks constructed on the fabrics.

Physical Fabrics Management

Networks are built on the physical fabrics that provide network resources such as links and IP subnets. The physical fabrics can be fully managed, host managed, or unmanaged in Oracle Enterprise Manager Ops Center.

Depending on the network infrastructure in your datacenter, you can use Oracle Enterprise Manager Ops Center to manage the physical fabrics in the following way:

- **Fully managed fabrics**

  When you discover a physical switch in Oracle Enterprise Manager Ops Center, all the fabrics that the switch supports are discovered and managed. A physical fabric can be partitioned to support many logical fabrics. Each port on an Ethernet switch can support 128 fabrics through its VLAN ID. Each partition on an InfiniBand switch can support 32000 partition keys. You can create VLANs or partitions as required. The fully managed fabrics facilitates to create dynamic private network for each VLAN ID or partition key.

- **Host managed fabrics**

  The host of the switch is a managed asset in Oracle Enterprise Manager Ops Center. You must manually enable the VLAN IDs or the partition keys on the switch ports connected to the host.

  The host managed fabrics facilitates to create dynamic private network for each VLAN ID or partition key.

- **Unmanaged fabrics**

  The switches are not managed in Oracle Enterprise Manager Ops Center. The VLAN IDs or partition keys are not available to manage the fabrics through a host. The networks are defined or discovered while discovering an asset in Oracle Enterprise Manager Ops Center. You cannot create dynamic private networks on unmanaged fabrics. Instead, you can assign some of the managed networks on these fabrics as private, called as static private networks.

Networks and Network Domains

You can create, define, or discover the networks in Oracle Enterprise Manager Ops Center as follows:

- Create and Define Public Networks
- Create Network Domains
- Create Private Networks
- Dynamic Private Network Creation

Create and Define Public Networks

Use the **Create Network** or **Define Network** option in the UI to create public networks on the fully managed and host managed fabrics. When you use Create Network option, Oracle Enterprise Manager Ops Center automatically assigns the VLAN IDs or P-Keys.

When you use Define Network option, you must specify the VLAN IDs for host managed Ethernet fabrics. For InfiniBand networks, the P-Keys are automatically assigned. For unmanaged fabric, use Define Network option to create public networks on the fabrics. You do not require to specify any VLAN IDs for Ethernet fabrics.

Create Network Domains

Network domain is a container for managed networks that handles the relationship between the physical fabrics that support the networks and the virtualization hosts or server pools that use the networks. The Oracle Enterprise Manager Ops Center software always has a **Default Network Domain** and all public networks are members of that domain. You can create a user-defined network domain. When you create a network domain, you assign the fabrics and associate the networks that are already known to the system.

Create Private Networks

Use the Create Private Network option to create private networks on fully managed and host managed fabrics. The **Create Private Network** option is useful for environment where you want to isolate applications any services from the public network. Private networks are created within a specific user-defined network domain for a specific purpose.

Dynamic Private Network Creation

To facilitate dynamic private network creation from virtual datacenter, associate user-defined network domains with the server pools of the virtual datacenter.

Select fully managed and host managed fabrics in the network domain to dynamically create private networks on demand. When you use unmanaged fabrics, then you must assign some of the existing managed networks assigned to the fabric. You must choose the managed networks that are not in use and do not route to other networks.

Properties of a Network

Figure 17–1 is an example of the network characteristics that appear in the **Details** tab.

Figure 17–1 Network Details Tab

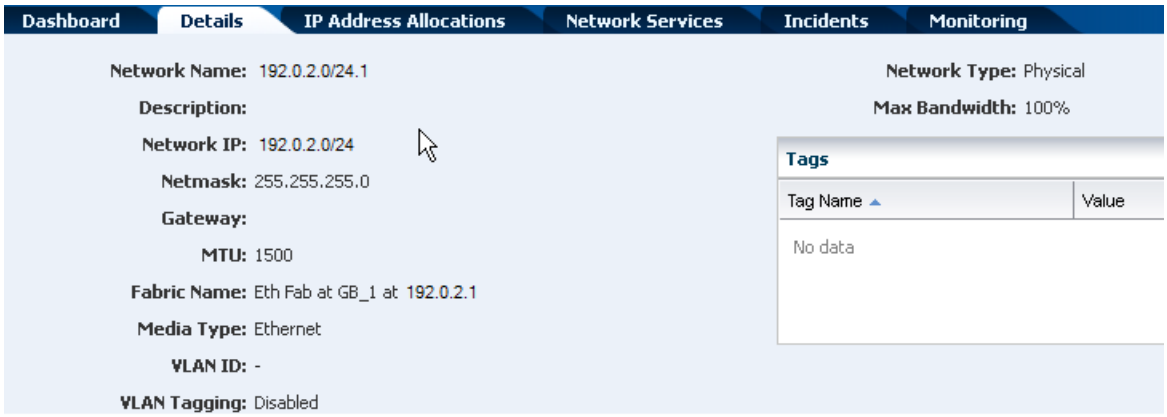

You cannot change the network IP address or the network type. Use the **Edit Network Attributes** action to change the network name and description, default gateway, MTU size, and to change the static IP routes. To change the MTU size, see the instructions for the Maximum Transmission Unit (MTU).

Use the **Edit Managed IP Ranges** action to change the range of IP addresses that are available from the selected network. You specify the range with the starting IP address and the ending IP address. You have the option to exclude a specific IP address from the range. When you attach the network to a virtualization host, server pool, or virtual datacenter, the IP address is not available. You cannot exclude an IP address that is in

use, which can be difficult to determine. For example, in a virtual data center, an account is assigned a range of IP addresses for its exclusive use. While the account exists, the IP addresses are in use, regardless of whether there is network activity.

VLAN and VLAN Tags

For fabrics based on Ethernet protocol, the ability to use VLAN tags is an attribute of each network. Use the **Edit Network Attributes** action to add or change the VLAN capability only for networks on fully managed fabrics.

When one CIDR supports both tagged and untagged networks, you can distinguish them by the default User Friendly Name (UFN), as shown in Figure 17–2. Oracle Enterprise Manager Ops Center appends the VLAN ID or tag to the UFN. For an untagged network with no VLAN ID, the `[UNTAG]` string is appended.

Figure 17–2 VLAN Tags

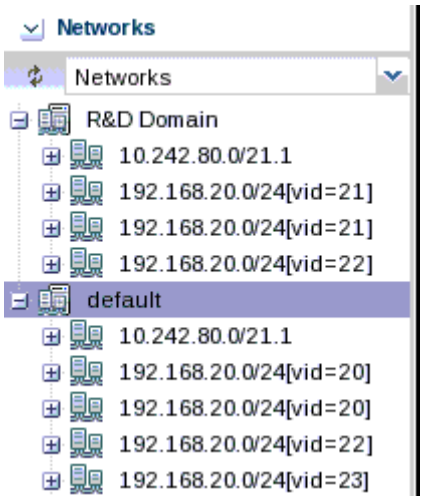

IP Multipathing Groups

Using IP Multipathing (IPMP), two or more physical network interface cards (NIC) form a group that use one IP address. If one NIC fails, the other NIC in the group maintains network access.

A network interface can be a physical network interface card (NIC) or, for an Oracle Solaris OS asset, it can be an IPMP group or link aggregation. You can implement both methods on the same network because they work at different layers of the network stack.

For information about how IPMP groups work in Oracle Solaris 11, see *Network Interfaces and Network Virtualization* at `http://www.oracle.com/technetwork/documentation/solaris-11-192991.html`. For Oracle Solaris 10, see *IP Services* at `http://docs.oracle.com/cd/E26505_01/html/E27061/index.html`.

Note: IPMP groups are supported only for IPv4 protocol.

IPMP provides increased reliability, availability, and network performance for systems with multiple physical interfaces because IPMP detects a physical interface failure and migrates network access to another member transparently.

Using IPMP, you can configure two or more physical interfaces into an IPMP group. If an interface in the group fails or is removed for maintenance, IPMP migrates the failed interface's IP addresses to another member of the group. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections.

The association between an IPMP group and a network must be unique. You can associate an IPMP group with only one network and you can associate a network with only one IPMP group or individual NICs.

In an IPMP group, you define whether each interface is a failover or a standby interface. The actions of each type differ if the current network interface fails, as follows:

- Network access changes from the failed interface to the failover interface in the IPMP group and uses the failover interface data address. You must provide the data address for an interface that is defined as failover.
- Network access changes from the failed interface to the standby interface in the IPMP group but does not change its data address. The data address of the failed interface migrates to the standby interface.

Link-based failure detection in an IPMP group is always enabled if your interface supports this type of failure detection. You can set up probe-based failure detection by providing a test address for each interface in the group.

You can create a single IPMP group while provisioning an operating system. If you create IPMP groups manually, Oracle Enterprise Manager Ops Center identifies and displays the groups on the UI. See Creating IPMP Groups for information and procedures for creating IPMP groups.

Link Aggregation

A network interface can be a physical network interface card (NIC) or, for an Oracle Solaris OS asset, it can be an IPMP group or link aggregation. You can implement both methods on the same network because they work at different layers of the network stack.

In an aggregated link, two or more NICs form a group and all members of the link aggregation provide network access at the same time. In addition to the high availability and load balancing that an IPMP group provides, an aggregated link can provide increased throughput when the network ports are also aggregated.

When interfaces have been aggregated, they are treated as a single network interface. Oracle Enterprise Manager Ops Center includes any link aggregations in the list of available NICs as if the link aggregation were an individual interface. To assign a network with a link aggregation to an Oracle VM Server or global zone, select the link aggregation from the NIC list. You can view the link aggregation details on the Oracle VM Server's or global zone's Network tab as described in Link Aggregation.

Link aggregation is a standard defined in IEEE802.3ad. An aggregated link consists of several interfaces on a system configured as a single, logical unit. Link aggregation increases the speed and high availability of a connection between a server and a switch. The most common protocol used to manage link aggregation is LACP (Linked Aggregation Control Protocol).

For information about how link aggregation work in Oracle Solaris 11, see *Network Interfaces and Network Virtualization* at http://www.oracle.com/technetwork/documentation/solaris-11-192991.html. For Oracle Solaris 10, see *IP Services* at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html

In Oracle Solaris 10 and by default in Oracle Solaris 11, the type of link aggregation you create is a trunk aggregation, which has these requirements:

- All the members of the aggregated link are connected to the same switch.
- The members of the aggregated link are of the same type. For example, NICs with the `e1000g` interface cannot be mixed with NICs that use the `bge` interface.
- The required driver is `GLDv3`.

Oracle Solaris 11 supports an alternative to trunk aggregation called Datalink Multipathing Aggregations (DLMP) This type of aggregation overcomes the limitations of trunk aggregation for network virtualization because DLMP aggregation works with more than one switch and provides the benefits of the link layer of the network stack to the aggregation.

In trunk aggregation, every port is associated with every datalink in the link aggregation. In a DLMP aggregation, every port is associated with every datalink in the link aggregation and every port is associated with the primary network interface and any of its VNICs that are configured to use the link aggregation.

Note: In the current release, Oracle Enterprise Manager Ops Center can display the details of both trunk and DLMP aggregation and displays them for selection when attaching a network. However, it is not possible to create a DLMP link aggregation.

For a link aggregation created in Oracle Solaris 11 OS, the MTU size for one of the members of the aggregation must be at least 9216 bytes to allow Oracle VM Servers and logical domains to use VLAN tagged networks. To change the MTU size, see the Maximum Transmission Unit (MTU) instructions.

Networking for Virtualization and Virtual Datacenter

Oracle Enterprise Manager Ops Center provides systems and users with efficient, controlled and secure sharing of the networking resources. The virtualization properties available in different Oracle Solaris OS version and the virtualization technology are implemented and available through the software UI.

The UI provides options to assign the network connection to the managed assets. You can also select the network interfaces through which the network connection is made to the virtualization host, virtual host and the OS. The attachment and connection varies, depending on the virtualization technology.

Networking for Server Pools

A server pool must have at least one network. When a server pool has more than one network, all virtualization hosts in the server pool are associated with the same set of networks. When you add a virtualization host to a server pool, the virtualization host is provided access to all the networks defined for the pool. This ensures that all virtual hosts have network access, even when you migrate a virtual host from one virtualization host to another one within the pool.

For zones and Oracle VM Server for SPARC server pool, it is recommended to create server pool that has homogenous network connection. Refer to Chapter 21, "Server Pools" for more detailed information.

Figure 17–3 is an example of network connections to two virtualization hosts in a server pool. This server pool has two virtualization hosts and two network associations.

Figure 17–3 Network Connections for a Server Pool

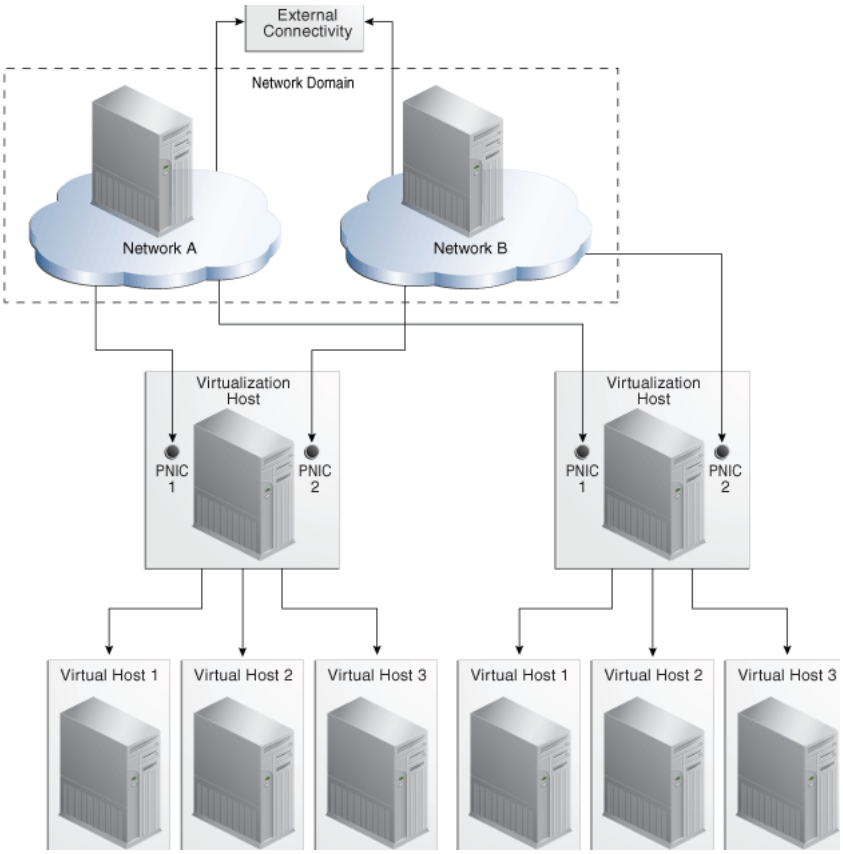

Networking for Zones

When you attach networks to zones, VNICs are created. Virtual Network Interface Cards (VNICs) are pseudo interfaces created on top of datalinks. It has an automatically generated MAC address.

You can define the mode of the network to be attached as Shared IP or Exclusive IP. In Shared IP mode, the global zone shares its network interface with one or more zone. You must define the network interface when you assign the network to the global zone. In Exclusive IP mode, a dedicated network interface is allocated to the zone. You can choose the network interface when you assign the network to a zone.

When a network is assigned as shared on a global zone, you can assign the network as exclusive on another global zone. For a global zone, a network can be attached in either shared or exclusive mode only. For non-global zones, a network that is used in a shared mode for one zone cannot be used in exclusive mode for another zone.

While you attach networks to a global zone, you can deploy IP Multipathing (IPMP) to obtain better network performance or link aggregation to provide increased reliability, availability, and network performance for systems with multiple physical interfaces.

Table 17–3 identifies the differences in attaching the network for Oracle Solaris 10 OS and Oracle Solaris 11 OS global zone in Oracle Enterprise Manager Ops Center.

Table 17–3 Differences in Network Connection for Global Zone

Oracle Solaris 10 OS	Oracle Solaris 11 OS
You can attach network in Shared IP or Exclusive IP mode.	Networks are always attached in Exclusive IP mode.
You cannot make multiple connections to a network.	You can make multiple connections to a network.
You can deploy IPMP or Link Aggregation for better network performance.	You can deploy only Link Aggregation in Oracle Solaris 11 OS.

In Oracle Solaris 11 OS, the network is always attached in exclusive IP mode, this is because a VNIC is created when the zone boots, and deleted when the zone is halted.

When you connect networks to the global zones, you can also select the tagging mode for the networks configured with VLAN ID. You can select Tagged or Untagged mode for the network connection.

Zones Server Pool

Server pool for zones reflect the networking properties of the Oracle Solaris OS version of the global zones in the pool. You can connect to a network only once for Oracle Solaris 10 OS. Whereas, you can make multiple network connections for Oracle Solaris 11 OS.

For zones server pool that contains a mixture of Oracle Solaris 10 and Oracle Solaris 11 OS, you cannot make multiple connections to a network.

Also, create zones server pool that are homogenous in network tagging mode. It can prevent any network outages for the zones created on the members of the server pool.

Before you attach a network to a server pool, verify that each virtualization host in the server pool has a physical network interface to the network so that all members of the pool can continue to share the network resources of the server pool.

Networking for Oracle VM Server for SPARC

You can attach networks to the Oracle VM Server using the physical interfaces or `etherstub` device that can belong to the control domain, I/O domain, or root domain. When you attach networks to the Oracle VM Server, you can select the service domain and the physical network interfaces available from that domain.

You can also specify the tagging mode for attaching networks configured with VLAN ID. You can select Tagged or Untagged mode for the network connection.

Attaching networks to Oracle VM Server result in the creation of a virtual switch for each network connection. This is not applicable for SR-IOV enabled networks. See SR-IOV Enabled Networks for more information about attaching SR-IOV enabled networks.

You can make multiple connections to a network in the control domain. For each network connection, a virtual switch is required. If there is an already existing virtual switch for the physical interface or etherstub device, you can re-use the virtual switch.

If there is no virtual switch for the physical interface or etherstub device, you can either provide a user-friendly name for the virtual switch or a virtual switch is automatically created with a default naming pattern. For an example network 1.1.1.0/24, the virtual switches take the name as 1.1.1.0_24, 1.1.1.0_24_1, 1.1.1.0_24_2 and 1.1.1.0_24_3. This ensures that the switches have unique names.

When a network connection is made to the server, the virtual switch created is incremented. When you create and start a logical domain, you define the virtual switch that connects to the logical domain. Each virtual switch must be connected to a NIC.

When you connect to the physical interfaces from I/O domains and root domains, the virtual switch is created in the control domain. You cannot define the IP address allocation for the network connection. Instead, you can define the IP address in the OS of the guest domain as required. You can define the IP address only when you use the network interfaces from the control domain or primary.

You can create IPMP groups and aggregate links in the control domain.

SR-IOV Enabled Networks

An SR-IOV enabled network interface means that there are virtual functions created on the physical functions of the PCIe Endpoint device and you can assign the virtual functions to the logical domains. Oracle Enterprise Manager Ops Center does not create any virtual switch when you connect to a network using SR-IOV enabled network interface. When you select **SR-IOV** option while attaching networks to the Oracle VM Server, only the interfaces on which the virtual functions are created are available for network configuration.

Guest domains that are assigned with SR-IOV enabled networks cannot be migrated. SR-IOV enabled networks are available only from control domain and root domain.

SR-IOV enabled networks on root domain are available only in the following conditions:

- Oracle Solaris 11 Update 1 OS (SRU 4.5) is necessary for dynamic attach of networks.
- Available only from Oracle VM Server for SPARC 3.1 version.
- Refer to Oracle VM Server for SPARC Release Notes at `http://docs.oracle.com/cd/E38405_01/html/E38409/index.html` for hardware and firmware requirements for SR-IOV feature.

Attach Networks to Oracle VM Server for SPARC Server Pool

Before you attach a network to a server pool, verify that each virtulization host in the server pool has a physical network interface to the network so that all members of the pool can continue to share the network resources of the server pool.

The following options are available when you attach networks to the server pool:

- You can select the service domain which provides the network interface for the network connection.
- If the Oracle VM Server is already connected to the network, you can keep the existing connection or make a new connection.
- If there are any virtual switches available for the network interface, you can re-use the virtual switches.
- If there are no virtual switches, you can either provide a name for the virtual switch or a virtual switch is automatically created with a default naming pattern.

- You can select the tagging mode for the networks configured with VLAN ID.

It is recommended to maintain server pool networks attached either in tagged or untagged mode to the server pool members. You can maintain server pool with mixed configuration. There is likely occurrence of network outage in the logical domain OS when you try to migrate the logical domain between servers that have different tagging modes. To avoid such outage, maintain the server pool members with homogenous network condition. Refer to Network Tagging Mode Conditions for more information about selecting tagging modes for the network connection.

Connect Networks to Logical Domains

You can connect networks to logical domains in running state. You can make multiple connections to a network. For each connection, you require a virtual switch or virtual function of SR-IOV enabled networks to connect the logical domain to the network.

You can re-use a virtual switch to make multiple connections to a network from the logical domain or use the same virtual switch to connect to a network for different logical domains. A virtual network device or vnet is defined when you connect the logical domain to a network through a virtual switch. For each network connection, a vnet is created. Oracle Enterprise Manager Ops Center tries to re-use the vnets. This reduces the number of vnets created for the network connections.

You can connect network root domains and I/O domains only when their operating systems are managed in Oracle Enterprise Manager Ops Center.

When you connect networks to logical domains, you can define the following parameters for the connection:

- Select the service domain that will provide the network services.
- Select the network mode as Tagged or Untagged for an VLAN ID network.
- Select SR-IOV enabled network connection.
- Select the virtual switch or the virtual function for SR-IOV enabled network function through which the logical domain is connected to the network.

You need an untagged network connection for provisioning OS on the logical domain. If the network is already configured with an VLAN ID, then select Untagged option while connecting the logical domain to the network.

When the network is connected to logical domain using Oracle Enterprise Manager Ops Center, by default 10 alternate MAC addresses are created.

Networking for Oracle VM Server for x86

During the installation of Oracle VM Server, the network interface used for the management is configured as a bonded interface. The bond is created with one interface and named as `bond0`. You can create additional bonds to add redundancy and load balancing of your network environment.

Attach networks to the Oracle VM Server or to the server pool that consists of a group of Oracle VM Servers on an Oracle VM Manager. Configure the network interfaces or the bonds to the network to be attached. You can assign different roles or functions to the networks attached to the Oracle VM Server.

The following are the network roles available for an Oracle VM Server:

- **Server Management**: Manages the Oracle VM Servers in a server pool. The Oracle VM Manager has one Server Management network.

- **Live Migrate**: Migrates the virtual machines from one Oracle VM Server to another in the server pool, without changing the state of the virtual machine.
- **Cluster Heartbeat**: Verifies that the Oracle VM Servers in the server pool are running.
- **Virtual Machine**: Monitors the network traffic between the virtual machines in a server pool.
- **Storage**: Transfers between virtual machines and virtual disks.

The management network created during the installation of Oracle VM Server has the following roles:

- Server Management
- Cluster Heartbeat
- Live Migrate

You can add and remove the roles of this management network, except for the Server Management role.

Depending on the available network interfaces on the Oracle VM Server, you can attach networks to Oracle VM Server and assign different roles to the networks. For example, you can attach the network in which your storage servers are placed and assign the Storage role to that network. You can assign a network with Live Migrate to be used only for migration.

Networking for Virtual Datacenters

Each virtual datacenter uses server, storage, and network resources in a dynamic way, allocating and releasing resources when necessary.

The virtual datacenter inherits its network resources from the network domain that supports the server pool. These networks form the public external networks for the virtual datacenter. These networks can then be assigned to the accounts in the virtual datacenter. When the user of an account creates a private vNet, either a dynamic private network is created or the static private network is made available for use in that account.

For a complete description of networks for virtual datacenters, see Creating vNets and Setting Up Network Resources in Chapter 22.

Related Resources for Networks

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- For Oracle Solaris 11 network interfaces and network virtualization, see http://www.oracle.com/technetwork/documentation/solaris-11-192991.html
- For Oracle Solaris 10 IP services, see http://www.oracle.com/technetwork/documentation/solaris-10-192992.html

18

Oracle Solaris Zones

The following information is included:

Introduction to Oracle Solaris Zones

Oracle Solaris Zones, also known as Oracle Solaris Containers, are used to virtualize operating systems and provide an isolated and secure environment for running software applications. A zone is a virtualized operating system environment created within a single instance of the Oracle Solaris operating system.

Think of a zone as a box with flexible, software-defined walls. One or more applications can run in this box without interacting with the rest of the system. Because zones isolate software applications or services, applications that are running in the same instance of the Oracle Solaris OS are managed independently of each

other. For example, you can run different versions of the same application in separate zones.

Zones require a machine that is running an Oracle Solaris 10 or later release.

Global and Non-global Zones

The global zone is the default operating system and has control over all of the processes and has system-wide administrative control. The global zone oversees the CPU, memory, and network resource allocation of all of the non-global zones. A global zone always exists, even when no other zones are configured.

Non-global zones, or simply zones, are configured inside the global zone. Zones are isolated from the physical hardware by the virtual platform layer. A zone cannot detect the existence of other zones.

Types of Non-Global Zones

You can create different types of non-global zones for different purposes:

- **Sparse Root Zone**: Contains a read/write copy of a portion of the file system that exists on an Oracle Solaris 10 global zone. Other file systems are mounted read-only from the global zone as loop-back virtual file systems. As part of creating a sparse root zone, the global administrator selects which file systems to share with the sparse root zone and the default read-only file systems: `/usr`, `/lib`, `/sbin`, and `/platform`. All packages that are installed on the global zone are available to the sparse root zone; a package database is created and all files in the mounted file system are shared with the zone.

  Note: Sparse root zones are not available beginning with Oracle Solaris 11. You can create sparse root zones only in Oracle Solaris 10.

- **Whole Root Zone**: Contains a read/write copy of the entire file system that exists on the global zone. When a whole root zone is created, all packages that are installed on the global zone are available to the whole root zone; a package database is created and all files are copied onto the whole root zone for the dedicated and independent use of the zone.
- **Branded Zone**: The non-global zone runs the same operating system software on the global zone. The branded zone facility is used to create non-global branded zones that contain operating environments different from that of the global zone. For example, you can install Oracle Solaris 8, 9, or 10 in a branded zone.

  Note: Oracle Solaris 11 Immutable Zones, Kernel Zones, and Zones on Shared Storage (ZOSS) are not supported in Oracle Enterprise Manager Ops Center 12.2.

  Oracle Solaris 11 shared IP zones support is limited to basic monitoring.

Zones and Virtual Machines

Oracle Enterprise Manager Ops Center supports three types of virtualization:

- Oracle Solaris Zones: operating system virtualization
- Oracle VM Server for SPARC: hardware virtualization on a SPARC platform

- Oracle VM Server for x86: hardware virtualization on an x86 platform

You can create and manage zones within virtualized hardware on both SPARC and x86 platforms, including managed zones in Control Domains and I/O Domains.

Roles for Oracle Solaris Zones

The following table lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 18–1 Oracle Solaris Zones Tasks and Roles

Task	Role
Create, manage, update, and delete zones	Virtualization admin
Provision and manage virtualization host	Virtualization admin
Discover and manage virtualization hosts	Asset admin
Create and manage zone profiles and plans	Profile and plan admin
Create and manage IPMP groups	Network admin
Create and manage Link Aggregation	Network admin
Set monitor thresholds	Asset admin

Actions Available for Oracle Solaris Zones

Oracle Enterprise Manager Ops Center provides the following solution for managing your zones environment in a data center:

- Create zones using profiles and deployment plan.
- Discover and manage existing zone environments.
- Perform zones functions such as booting, rebooting, shutting down, cloning, migrating and delete zones from the software UI.
- Manage zone configuration such as file systems, storage, and networks from the UI.
- Aggregate your NICs or create IP Multipathing (IPMP) groups.
- Create zones server pool for balancing available resources and provide the infrastructure support for virtual datacenter creation.
- Use UI and native CLI interchangeably to create and manage zones.
- Complete support for managing Oracle Solaris 11 zones.
- Options to upload scripts to manage the unmanaged file system and zone dependencies.

Location of Oracle Solaris Zones Information in the User Interface

Figure 18–1 shows how zones appear beneath the global zone in the Assets section of the user interface. The zones appear with a different icon. Badges on the icon indicate the status of the zone. The green triangle badge indicates that the zone is running. The black circle with a white vertical line indicates that the zone is shut down.

Figure 18–1 Zone Display in the UI

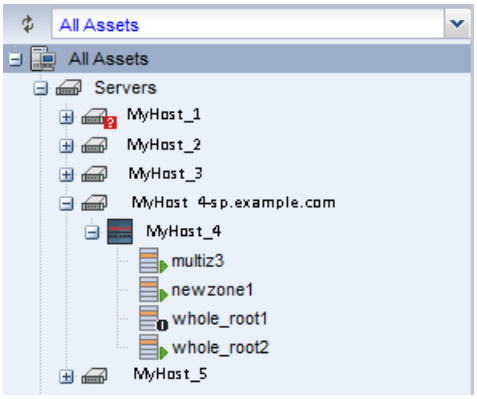

Table 18–2 lists where to find the different information about zone in the UI.

Table 18–2 Location of Zone Information in the UI

To See	Location
Zones	Expand Assets in the Navigation pane. The operating system under which the zones are created are listed with the icons representing it.
Zone resources	Expand Assets in the Navigation pane, then select Oracle Solaris OS. Select a zone listed under it. The center pane lists the zone information such as Storage, Networks, Analytics, and Summary.
Zone actions	Expand Assets in the Navigation pane, then select Oracle Solaris OS. Select a zone listed under it. The Actions pane lists the various zone actions such as Add Storage, Migrate Zone, Replicate Zone, Connect Network, Add File System, Move Storage, Boot, Halt, Shutdown, and Reboot Zone.

Preparing Your Global Zone

The global zone is the default Oracle Solaris 10 or 11 operating system installed on a system. When you provision systems, you can use the **Add Assets** or **Find Assets** option in Oracle Enterprise Manager Ops Center to discover the global zone. See Chapter 13, "Operating System Provisioning" for information on provisioning a system with an Oracle Solaris operating system.

Each global zone has an `/etc/patch/pdo.conf` file that specifies the number of processes that are forked to execute the patch utilities in parallel on a zoned system. The file contains a `num_proc=` entry indicating the number of processes to be forked. Oracle Enterprise Manager Ops Center requires the number of CPUs on the system to be *1.5.

Beginning with Oracle Enterprise Manager Ops Center 12.2.2.0.0, the software checks this file every 12 hours, during an agent refresh, and makes the following adjustments:

- If the file does not exist, the software creates the file with the following entry: `num_proc=`*Number of CPUs on the system*1.5*.
- If the file exists, but the value in the entry is incorrect, where `num_proc=`*Value is not equal to Number of CPUs on the system*1.5*, the software removes this entry and appends to the contents of the file an entry: `num_proc=`*Number of CPUs on the system*1.5*.

If the file exists with a correct entry: `num_proc=`*Number of CPUs on the system*1.5*, the file is not changed or updated. For more information about this file, see the Oracle Solaris Zones pdo.conf man page.

You must prepare your global zone with the network and storage resources that the non-global zones will use, including associating the required storage libraries with the global zone and attaching networks to the global zone. You can also aggregate the interfaces or create IPMP groups in the global zone before you connect them to the zones. The use of IPMP and link aggregated interfaces results in enhanced network availability for the zones.

This section covers the following topics:

Associating a Storage Library with a Global Zone

Global and non-global zone metadata and the data that is the output of operations is saved in a SAN or NAS storage library.

You can associate the following types of storage libraries to a global zone:

- Filesystem storage: NAS libraries
- Block storage: SAN and Dynamic storage libraries

Libraries associated with the global zone are available to all of the zones in the hierarchy of the global zone. A list of available LUNs appears when you create a new zone. When you associate a storage library with the global zone, Oracle Enterprise Manager Ops Center discovers the available LUNs. A list of available LUNs appears in the wizard when you create a new zone in the global zone. When you select the LUN for the zone, the LUN is reserved for that zone and is not available for other zones. When the SAN storage library is associated with a zones server pool, the LUNs are available to all global zones in the pool.

When using a SAN storage library, the metadata and data are stored on LUNs that are managed by the SAN storage library that is associated with the global zone. The number of LUNs determines the number of zones that the library can support. SAN LUNs are associated with Fibre Channel or iSCSI target groups. Fibre Channel targets use a dedicated optical network and iSCSI targets use the IP network. In both cases, the targets in the target group expose the LUNs as a storage resource for the zones.

When you use LUNs backed by Fibre Channel disks, the block storage is static. You cannot change the size of the LUNs, create LUNs, or delete LUNs in the UI. When the LUNs are not in use, you can add LUNs to the library or remove LUNs.

For information about how to create SAN and NAS storage libraries., see the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm`.

To Associate a Library With the Global Zone

1. Select the global zone in the Assets section.
2. Click **Associate Libraries** in the Actions pane.

The Associate Library window lists the libraries that are not associated with the global zone.

3. Select one or more libraries to add.
4. Click **Associate Libraries**.

   A job is submitted to associate the libraries with the global zone.

To Disassociate Libraries From Global Zone

When you disassociate a library from the global zone:

- The libraries are not available to new zones.
- The libraries remain associated to existing non-global zones in that global zone.

1. Select the global zone in the Assets section.
2. Click the **Libraries** tab in the center pane.

   The associated libraries with the global zone are listed.
3. Select a library from the list.

   The Disassociate Library icon is enabled.
4. Click the **Disassociate Library** icon to display the Disassociate Library window.
5. Click **Disassociate Library** to confirm unmounting the library.

Managing Global Zone Networks

You can attach one or more networks to global zone using the Attach Networks option. When your networks are grouped as network domains in Oracle Enterprise Manager Ops Center, select the domain and choose the corresponding networks to be attached.

When you assign a network to a stand-alone global zone, you must define the mode of the network, either Shared IP or Exclusive IP:

- In Shared IP mode, the global zone shares its network interface with one or more zone. You must define the network interface when you assign the network to the global zone.
- In Exclusive IP mode, a dedicated network interface is allocated to the zone. You can choose the network interface when you assign the network to a zone.

Beginning with Oracle Enterprise Manager Ops Center 12.2.2.0.0, the type of IP (Shared or Exclusive) that is assigned appears in the non-global zone's Summary tab.

A network that is assigned as shared on a global zone can be assigned as exclusive on another global zone. For a global zone, a network has only one mode. For non-global zones, a network that is used in a shared mode for one zone cannot be used in exclusive mode for another zone.

You can attach networks that are configured with VLAN ID or P-key. When you use networks with VLAN ID, you can select to attach the network in tagged or untagged mode.

While you attach networks to a global zone, you can deploy IP Multipathing (IPMP) to obtain better network performance or link aggregation to provide increased reliability, availability, and network performance for systems with multiple physical interfaces.

For stand-alone global zones with Oracle Solaris 10 OS, network can be attached in shared or exclusive IP mode. Whereas with Oracle Solaris 11 OS, the network is always

attached in exclusive IP mode. This is because a virtual NIC or VNIC is created when the zone boots and deleted when the zone is halted.

This scenario is different for attaching networks to global zones that are in a server pool. See Attaching Networks for more information about attaching networks in a zones server pool.

Some of the networking conditions that must be followed while attaching the network in tagged or untagged mode:

- You can select networks without VLAN ID. The UI does not provide the option to select Tagged or Untagged mode.
- You can select to associate and configure the networks with VLAN ID in Tagged mode.
- You can select to associate and configure the networks with VLAN ID in Untagged mode.
- You can select to configure the networks in mixed tagging mode in the server pool. For example, you can attach the network N1 with VLAN ID = 100 in tagged mode with the server S1 and in untagged mode for server S2. Refer to Mixed Network Tagging Mode Configuration for more detailed information.
- You can attach networks whose VLAN ID is similar to another network already connected to the servers. For example, a server S1 is already connected to network N1 with VLAN ID = 100, then while creating the server pool with S1 as the member of the pool, you can attach a network N2 with VLAN ID =100.
- You can edit the VLAN ID of a network when you are attaching the network in Tagged mode for the first time.
- When you can edit the VLAN ID of the network, you cannot enter -1 as the value for the VLAN ID.
- If the selected network with a VLAN ID is already connected to the selected assets in Tagged mode, then you cannot edit the VLAN ID and make another connection.
- You cannot make multiple network connections to the global zone over the same network in both tagged and untagged modes. The mode can be either in tagged or untagged mode only. For example, if you attach network N1 with VLAN ID =100 for the first time to server S1 in Tagged mode, then you cannot make another connection to the same network N1 in Untagged mode. Every other connection with network N1 must always be in Tagged mode for server S1.
- If the selected members of the pool are already connected to network N1 with VLAN ID =100, then you cannot select the same network with different VLAN ID to be connected for the server pool.

IPMP Groups

IPMP groups provide network failover for your global and non-global zones. You can configure one or more physical interfaces into an IPMP group in the global zone and extend that functionality to the non-global zones. After configuring the IPMP group, the system monitors the interfaces in the IPMP group for failure. If an interface in the group fails or is removed for maintenance, IPMP migrates, or fails over, the failed interface's IP addresses. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections. The network access changes from the failed interface to the standby interface in the IPMP group and the data address of the failed interface migrates to the standby interface. See IP Multipathing Groups and Creating IPMP Groups for more information about IPMP groups.

Note: You can create IPMP groups when you attach network in exclusive IP mode on Oracle Solaris 10 and 11 operating systems or you can create an IPMP group without using the attach network option. IPMP groups are not available for shared IP mode networks.

When you attach networks to a global zone in a shared IP mode, you can create an IPMP group and configure the following characteristics:

- Define the number of members for the IPMP group.
- Select the NICs that are part of the IPMP group. The NICs selected are placed in the IPMP group.
- Select the type of interface:
  - Active interface enables you to provide the data address and choose whether failover must be enabled for the interface for uninterrupted access to the network. This is applicable for Oracle Solaris 10 OS and is inherent in Oracle Solaris 11 OS.
  - Standby interface enables you to provide a test address when you want probe-based detection.

Multiple shared IP zones can use the IPMP group. When you create a zone or connect a network to a zone, the IPMP group appears in the NIC list. Select the IPMP group from the list to connect to the selected network.

Link Aggregation

Several interfaces in a system can be aggregated into a single logical link. The aggregation is done as per the standard defined in IEEE802.3ad. When interfaces have been aggregated, they are treated as a single network interface. Oracle Enterprise Manager Ops Center displays the link aggregation in the list of available NICs as if it were an individual interface. You can assign a network with a link aggregation to a non-global zone, and select the link aggregation from the NIC list.

When you attach a network to a global zone, you can aggregate the physical interfaces and attach to the network. Link aggregation is available for both shared and exclusive IP network stack.

To Attach a Network to a Global Zone

When you attach a network to a global zone, you can choose to create an IPMP group or aggregate link. See Creating IPMP Groups to create IPMP groups on the global zone without using the attach network option.

The following network options are available:

- **Exclusive network**: For zones that use an exclusive IP network, each zone has a dedicated network interface. You can configure the interfaces when you connect a zone to the exclusive network. An IPMP group requires an exclusive network.
- **Shared network**: The global zone shares the network interface with the zones that use a shared IP network. You must configure the network interfaces of the global zone.

Perform the following steps to attach a network to a stand-alone global zone:

1. Select the global zone in the Assets section.
2. Click **Attach Network** in the Actions pane.

The Attach Network Wizard is displayed.

3. Select the network domain and the corresponding networks that are not yet connected to the asset are displayed in the network list. Select a network from the list.

   You can attach one or more networks to the global zone. Click **Next** to configure the networks.

4. For each selected network, select the following information:

   - Select the IP stack as **shared** or **exclusive**.
   - Select the Mode as Tagged or Untagged for networks configured with VLAN ID. If you are attaching the network in Tagged mode for the first time, then you can edit the VLAN ID of the network, provided the VLAN ID is not used by any of the networks attached to the asset.
   - Select the network configuration as **IPMP**, **Link Aggregation**, or **None**. For an exclusive IP stack, only the Link Aggregation network configuration is available.
     - When you select **IPMP** network configuration, you are taken to Step 6 for IPMP configuration.
     - When you select **Link Aggregation** configuration, you are taken to Step 5 to configure the aggregation.
     - When you select **None**, you are taken to Step 7 to configure the network interfaces for shared IP stack or to the Summary step.
   - Enter the number of connections for each network.

     The number of connections depends on the Oracle Solaris OS version. For Oracle Solaris 10 OS, you cannot make multiple connections. For Oracle Solaris 11 OS, you can make multiple network connections.

   Click **Next**.

5. Provide the following information for link aggregation:

   - The link aggregation name.
   - Select the load balancing policy.

   Click **Next** to configure the link aggregation. Specify the following information for link aggregation:

   - LACP mode and timer
   - MAC address policy and the MAC address if required.

   Click **Next** to go to the Step 7 to configure the interfaces.

6. Specify the IPMP group details.

   - IPMP group name
   - Number of members of the IPMP group
   - Select whether the probe-based failure detection must be enabled.

   Click **Next** to configure the interfaces.

7. Configure the interfaces for networks that do not have network configuration:

   - Specify the NIC and the IP address for the network connection.

- If required, modify the network tagging mode specified for the network connection.
- You can select **System Allocated for the NIC** and **Assign by DHCP** for the system to take care of the NIC and IP address allocation.
- Select **Do Not Allocate IP** for the IP address, when you do not want to allocate IP address for the network connection.

For IPMP group, select the interface that will act as Standby interface. You must have at least one active interface.

Click **Next** to specify the data addresses for IPMP group, otherwise go to Step 9.

8. Enter the data addresses that must be used for the active network interfaces. Also, specify whether failover must be enabled for the interface. You must have more than one interface to enable failover.
9. Review the information provided for attaching a network to the global zone and click **Finish** to attach the network.

To prevent a failure of the attach network job, ensure that you enter the correct information for different versions of Oracle Solaris OS.

Creating IPMP Groups

You can directly create IPMP groups on the global zone without using the attach network option. From the network tabs, you have the option to create and manage the IPMP groups in the selected Oracle Solaris OS.

Figure 18–2 shows the options that are available to create and manage IPMP groups.

Figure 18–2 IPMP Group Option

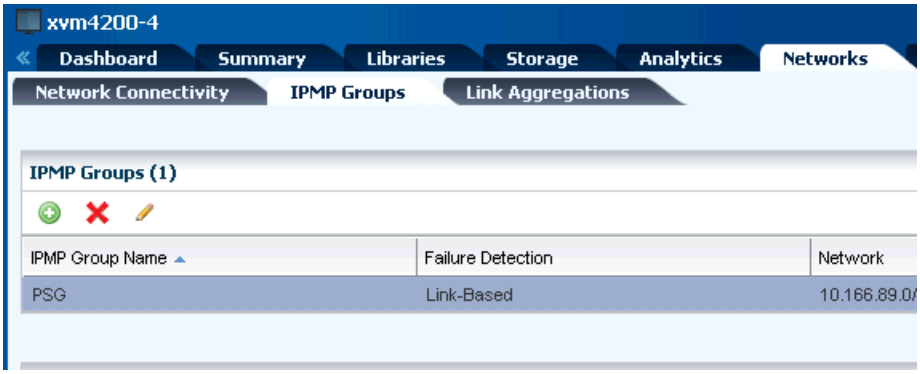

IPMP provides physical interface failure detection, transparent network access failover, and packet load spreading for systems with multiple interfaces. Oracle Enterprise Manager Ops Center provides option to create IPMP groups. You can configure one or more interfaces into an IPMP group. The group functions like an IP interface with data addresses to send and receive network traffic. When an underlying interface in the group fails, the data addresses are redistributed among the remaining underlying active interfaces in the group. Thus, the group maintains network connectivity despite an interface failure. With IPMP, network connectivity is always available, provided that a minimum of one interface is usable for the group. IPMP also provides load spreading for the outbound network traffic across the network interfaces in the group.

To create an IPMP group, you must the define the following parameters for the group:

- The active and the standby interfaces of the group. By default, an interface added to an IPMP group is active. You can configure as many standby interfaces as you want for the group.
- The link-based failure detection is enabled by default. You must select whether you want to enable Probe-Based failure detection. For probe-based failure detection, you must provide the test address to track the interface status.
- You must assign the data addresses for the physical interfaces in the IPMP group. Data traffic flow use the data addresses that are hosted on the IPMP interface and flow through the active interfaces of that group.

Refer to Creating IPMP Groups in Chapter 7, "Networks" for using the tagged and untagged mode for the networks that are configured using different media type. The tagging mode also varies for different Oracle Solaris OS version.

Creating Link Aggregation

You can also create link aggregation directly on the Oracle Solaris OS. to create link aggregation, you must define the following details:

- Load balancing policy
- LACP mode and timer
- MAC address policy and if required, the MAC address

Refer to Creating Link Aggregation section for more information about the procedure to create a link aggregation.

Modifying and Detaching a Network from the Global Zone

You can modify the network connection that are attached to the network except for the management network. The modify and unbind network options are available in the Network tab.

Figure 18–3 shows the icons the represent the following options that are available to manage networks of a global zone: Connect Networks, Disconnect Networks, and Modify Physical Connectivity.

Figure 18–3 Network Options

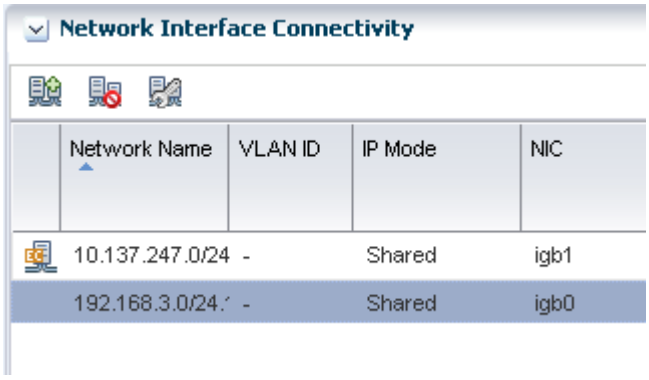

To Modify Physical Connectivity

1. Click the **Network** tab of the selected global zone in the Assets section.
2. .Click the **Modify Physical Connectivity** icon

Figure 18–4 shows the Modify Physical Connectivity window that displays the network details that can be modified.

Figure 18–4 Modify Physical Connectivity

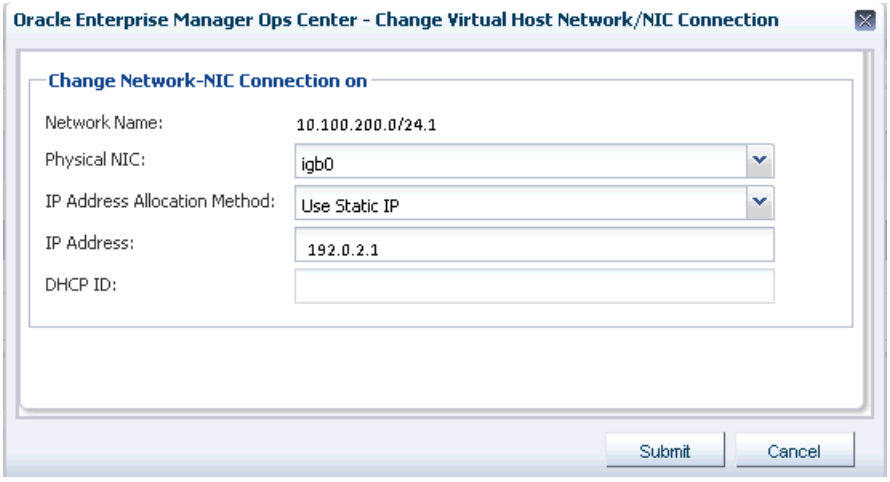

You can change the permanent IP address or change the allocation method.

Note: When the network interface is an IPMP group, you cannot modify the network interface.

Unbind Network Connection

When you unbind a network from a global zone, the global zone's non-global zones are also disconnected from the network. Select the global zone and click the **Unbind Network** icon to remove the network.

Discovering and Managing Existing Zones

When you discover a global zone that has existing zones, these zones are also automatically discovered and displayed on the UI. The option to deploy the agent on the global zone is inherited to the zones.

The following steps provide an outline of the procedure to discover and manage existing zones:

1. Use the option **Add Assets** to discover a global zone.
2. In the Add Assets Wizard, select whether you want to deploy the agent on the global zone or you want agentless management.
3. All the non-global zones in the selected global zone are automatically discovered.
4. When you want to deploy the agent on the global zone, the agent is installed on all the zones that are in running state. When you want to deploy the agent for the zones in shutdown state, boot them and use the option Switch Management Access. Using this option you can change the state from agentless management to managed by agent mode.
5. You can always use the option **Switch Management Access** to switch between managed by agent and agentless management mode.

You cannot create zones on a global zone that is managed agentless. For an agentless managed zone, you can boot, shutdown, halt and delete the zone. The Zones Virtualization Controller Agent is required on the zone when you want to use the full range of OS update actions on the OS. See Table 12–3 for information about what functions are supported for agent and agentlessly managed zones and operating systems. See Virtualization Agent Controllers and Changing the Type of Agent Controller for more information about agents.

When you use the native Oracle Solaris CLI to create a zone, select the global zone in the UI and then click the **Refresh** icon to display the new zone in the UI. You can interchangeably use the CLI and the UI to perform zone functions. To display the changes in the zone configuration and state, select the zone in the UI and click the **Refresh** icon. If you do not refresh, the software will refresh the state in 12 hours.

You can also use the **Find Asset** to discover an asset with Service Tags. See Chapter 2, "Asset Management" for more information about discovering and managing an asset in Oracle Enterprise Manager Ops Center.

As shown in Figure 18–5, the discovered zones appear in the UI and the following actions are enabled for the zone: Shutdown Zone, Halt Zone, Reboot, Edit Attributes, Edit Tags, Edit Configuration, Replicate Zone, Move Storage, and Delete Zone.

The Oracle Enterprise Manager Ops Center UI fully supports Oracle Solaris 11 zones configured with an exclusive IP mode for the network. You can discover and monitor Oracle Solaris 11 zones that are configured with shared IP, but zone support is similar to a zone without an agent. The zones appear in the UI, but active management is not supported and advanced network configurations are not monitored.

Note: For Oracle Solaris 11 zones that are configured with shared IP, you cannot perform migration, connect networks or add storage resources to the zone and this zone should not be part of a server pool.

Figure 18–5 Enabled Actions

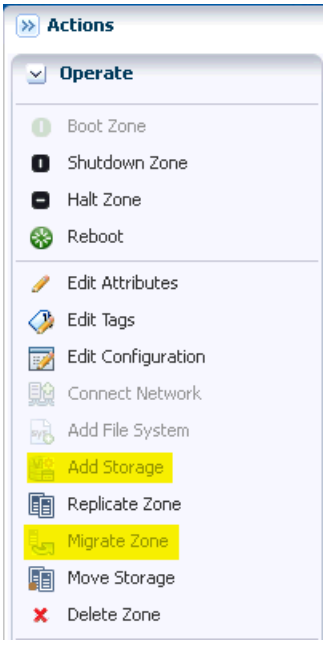

In Figure 18–5 the Migrate Zone and Add Storage options are not enabled because the zone storage source is unmanaged. To enable these actions, you must move the storage source to managed using the option Move Storage.

See Actions Available for Oracle Solaris Zones for more information about performing zone management operations.

Deleting or Unmanaging a Global Zone

When you delete or unmanage the global zone, all the non-global zones under it are also automatically unmanaged.

You do not need to unmanage the non-global zones before unmanaging the global zone. See Chapter 2, "Asset Management" for more information about deleting an asset.

Unmanaged Zone Storage Information

When you discover and manage existing zones, Oracle Enterprise Manager Ops Center handles the zone metadata and networks except for the zone storage. The metadata of the existing zones is stored in the local storage of the global zone. Whereas, the zone's storage is considered to be unmanaged storage source.

Even for the zones that are created using Oracle Enterprise Manager Ops Center, when you add a storage or file system manually using the native CLI, the storage becomes an unmanaged storage or file system.

See Moving Zone Storage for more information about how to move the unmanaged storage to managed storage.

Outline of Zone Creation

Using Oracle Enterprise Manager Ops Center, you create a zone profile that captures the zone configuration. Use the profile in a deployment plan to create one or more zones simultaneously. The benefit of using a profile is that you can create multiple zones with consistent configuration.

Note: When you create or migrate a zone using Oracle Enterprise Manager Ops Center, you cannot use Oracle Solaris Live Upgrade and an alternate boot environment (ABE) to update the zone.

Create a zone profile that captures all the zone configurations. Then, create a deployment plan with the zone profile. The deployment plan is a single step plan which covers only the zone profile. During the application of the deployment plan, select the number of zones to create. In the plan deployment, you can correct the storage and network details as required.

Figure 18–6 Zone Creation Process

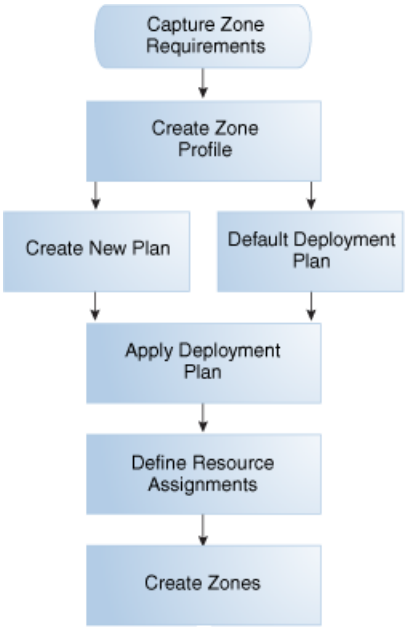

Determining Zone Requirements

Before you create a zone profile, determine the requirements for the zone.

Note: When you use MAC addresses, you must have free alternate MAC addresses on available on the global zone before you can create zones. The addresses must be available before you apply the deployment plan to create the zones.

The requirements vary according to the Oracle Solaris OS version and the type of zone. See the following sections for os-specific zone requirements:

- Requirements for Zones on Oracle Solaris 10 OS
- Requirements for Zones on Oracle Solaris 11 OS

Requirements for Zones on Oracle Solaris 10 OS

You can create sparse root, whole root, and branded zones on Oracle Solaris 10 OS. While creating the zone profile, select the appropriate options and provide the required resources for a successful zone creation.

Requirements for a Whole Root Zone

A whole root zone contains a read/write copy of the global zone's file system and has the following requirements:

- Minimum size of the file system is 5 GB.
- Minimum size of the virtual disk for the storage library is 6 GB.

Requirements for a Sparse Root Zone

A sparse root zone contains a read/write copy of a portion of the global zone's file system and shares the global zone's `/usr`, `/lib`, `/sbin`, and `/platform` directories in read-only mode. A sparse root zone has the following requirements:

- Minimum size of the file system is 1 GB.
- Minimum size of the virtual disk for the storage library is 1 GB.

Requirements for a Branded Zone

A branded zone emulates the user environment of earlier versions of Oracle Solaris. For example, you can create a branded zone to run Oracle Solaris 8 applications on your Oracle Solaris 10 system. Use the following procedure to prepare to create a branded zone:

1. Download the images for the operating system, as shown in Table 18–3.
2. Import the images into one of the software libraries associated with the global zone.

During the process of creating a branded zone, you install the images in the global zone. For instructions, see the README files in the Oracle Solaris 8 or 9 Containers bundle.

Table 18–3 Packages for Branded Zones

Global Zone OS	Branded Non-Global Zone
Oracle Solaris 10 10/08 OS or later	For Oracle Solaris 8, install SUNWs8brandk from Oracle Solaris 8 Containers 1.0.1. For Oracle Solaris 9, install SUNWs9brandk from Oracle Solaris 9 Containers 1.0.1. The Oracle Solaris 10 installation automatically installs the following required packages for branded zones: ▪ Oracle Solaris 9 branded zones: SUNWs9brandr and SUNWsbrandu packages ▪ Oracle Solaris 8 branded zones: SUNWs8brandr and SUNWs8brandu packages
Releases prior to Oracle Solaris 10 10/08 OS	For Oracle Solaris 8, install SUNWs8brandr and SUNWs8brandu from Oracle Solaris 8 Containers 1.0.1. For Oracle Solaris 9, install SUNWs9brandr, SUNWs9brandu, and SUNWs9brandk from Oracle Solaris 9 Containers 1.0.1.

Requirements for Zones on Oracle Solaris 11 OS

You can install zones in Oracle Solaris 11 OS either using the IPS packages in the default Oracle Solaris 11 Package Repository or using an image of an installed system running the Oracle Solaris release.

To install zones using images, you must have an Oracle Solaris 10 image in flash archive format or from an existing Oracle Solaris 11 OS image in gzip format. See `http://www.oracle.com/technetwork/documentation/solaris-11-192991.html` for how to prepare your installed Oracle Solaris 10 or 11 systems, and create the `flar` or `gzip` image. Upload these images to the storage libraries and install branded zones in Oracle Solaris 11 OS.

Requirements for Installing Zones Using the Repository

To install zones from the repository, configure your Oracle Solaris 11 Software Update Library in the Enterprise Controller to synchronize with the Oracle Solaris 11 Package Repository. When you install from repository, the default software group, `solaris-small-server` group is used. This occupies less space. The zones are whole root type only.

> **Note:** To provision Oracle Solaris 11 and Oracle Solaris 11 zones, the Enterprise Controller and Proxy Controller must be installed on an Oracle Solaris 11 operating system. See Enterprise and Proxy Controller Requirements for OS Provisioning for more information on operating system requirements for Oracle Solaris 11 actions.

Requirements for Oracle Solaris 10 Branded Zone

You can migrate an Oracle Solaris 10 OS into an Oracle Solaris 11 environment. Create an archive of the Oracle Solaris 10 instance that you would like to migrate. In the Enterprise Controller software library, import an ISO image of Oracle Solaris 10 OS that has the same architecture (SPARC or x86) of the instance to be migrated. Then, import the `flar` archive with the parent as the ISO image imported previously.

For Oracle Solaris 11 OS, create an archive in the format of `gzip` file using `cpio` command. Import the archive into Enterprise Controller software library with the reference to the parent ISO image.

You can migrate only systems that have Oracle Solaris 10 10/09 or later versions. To migrate earlier versions, install the kernel patch 141444-09 (SPARC) or 141445-09 (x86/x64), or later version, on the original system.

Since zones do not nest, existing zones in the original system are detected and a warning is issued that nested zones are not usable and that the disk space can be recovered.

To use the Oracle Solaris 10 package and patch tools in your Oracle Solaris 10 Container, install patches 119254-75 (SPARC) and 119255-75 (x86/x64) on your source system before the image is created. The P2V process works without the patches, but the package and patch tools do not work properly within the `solaris10` branded zone.

To use Oracle Solaris 10 zones on your system, the `system/zones/brand/brand-solaris10` package must be installed on the system running Oracle Solaris 11 OS.

Zone Configuration Parameters

When you create a zone profile, you can provide zone configuration details such as CPU model, memory caps, priority value, network mode, and storage resources for the zone.

- **Zone name**: Specify a unique name for the zone. Do not use names that start with global or SUNW.
- **Autoboot**: Specify whether the zone must boot immediately after you create it and whenever the global zone boots.
- **Shared CPU**: A zone with a shared CPU gets its CPU resources by the number of shares you allocate to it from the resource pool, which is also used by other zones. The new zone is added to the Fair Share Scheduling automatically. You have the option to set a maximum value for the CPU resources. The CPU cap limits the

amount of CPU resources that can be used by one zone. For example, a CPU cap value of 1 means 100% of a CPU.

- **Dedicated CPU**: A zone with a dedicated CPU gets exclusive use of the available CPU resources. You specify the minimum and maximum number of CPUs available to the new zone. A temporary resource pool is created and dedicated to the zone.

  Note: This parameter is available when you select dedicated CPU. Set the priority of the zone. Assign an importance value for the zone so that when there are not enough CPU resources to satisfy all zones, the zone with the greater importance value receives a larger share of the available CPU resources.

- **Priority of recovery**: When you enable automatic recovery for the zone, the priority of recovery value decides which zone must be migrated first during a global zone failure in a server pool. Set the priority of recovery between 0 to 100.
- **Memory Caps**: Set the maximum value for physical, swap and lock memory resources.
- **Naming Service**: Specify the name service that the zone uses to communicate with network objects. You can select the DNS, NIS, NIS+ or LDAP naming service. To specify the name service, you require the domain name and the IP address of the name server.
- **Shared IP Address or Exclusive IP Address**
  - A zone with a shared IP address uses its global zone's IP layer configuration and state. The zone has a logical network interface to the IP address.
  - A zone with an exclusive IP has its own dedicated IP layer configuration and state. The zone has its own set of network interfaces. You must configure the network interfaces using the same network configuration methods applied to all Oracle Solaris OS configurations.

    For Oracle Solaris 11 OS, only exclusive IP mode is supported for the network in the Oracle Enterprise Manager Ops Center UI.

Creating a Zone Profile

Oracle Enterprise Manager Ops Center provides option to create profile that captures the zone configuration. Use the zone profile in a deployment plan and apply to create one or more zones simultaneously.

To Create a Zone Profile

1. Expand the **Plan Management** section in the Navigation pane.
2. Expand Profiles and Policies and click **Oracle Solaris Zone**.
3. Click **Create Profile** in the Actions pane.

   The Create Profile – Oracle Solaris Zone Wizard is displayed.
4. Enter a name and description for the profile identification.

   Select whether you want to create a deployment plan automatically using this profile.
5. Select the OS version of the zone.

Ensure that you apply the plan on correct target to create the zone. See Table 18–4 for more information.

Table 18–4 Oracle Solaris Zones Targets

Zone OS Version	As	On Target
Oracle Solaris 10	Whole root zone	Oracle Solaris 10
	Sparse root zone	
Oracle Solaris 10	Branded zone	Oracle Solaris 11
Oracle Solaris 11	Whole root zone	Oracle Solaris 11
Oracle Solaris 8	Branded zone	Oracle Solaris 10
Oracle Solaris 9	Branded zone	Oracle Solaris 10

Click **Next** to specify the zone identity.

6. You can create one or more zones using the profile. To identify the zones, provide a zone prefix name and a number to start the series.

   Each zone created uses the prefix name appended with the number that is incremented. For example, if the prefix name is Myzone and the number to start from is 1, then the zones are created with names Myzone1, Myzone2, and Myzone3.

   Enter the description and tags for the zones. This is common for all the zones created using this profile.

   Click **Next** to specify the zone installation source.

7. According to the OS version selected, the page displays the following information:

   - Oracle Solaris 10

     Select whether you want to create whole root, sparse root, or branded zone. For branded zone, provide the following information:

     – Select the architecture as x86 or SPARC

     – Select the branded zone image. You must have created a flash archive image of an installed Oracle Solaris 10 OS and uploaded to the software library in Oracle Enterprise Manager Ops Center. These images are listed in the Branded Zone Image list.

     – Enter the prefix for the Host ID and the starting number.

     – Select the machine type as sun4u or sun4v.

   - Oracle Solaris 11

     You must select the installation source for installing Oracle Solaris 11 zones. You have the following options to select:

     – Install from repository. You must have configured your Oracle Solaris 11 Software Update Library in Oracle Enterprise Manager Ops Center in synchronization with the Oracle Solaris 11 Package Repository. This option installs the `solaris-small-server` software group by default. This is also referred to as the whole root zone for Oracle Solaris 11 OS.

     – Install from selected image. You must have created `gzip` archive image of an installed Oracle Solaris 11 OS and uploaded to software libraries in Oracle Enterprise Manager Ops Center. Select an image from the list. You can create the gzip archive formats using `cpio` or `zfs` command options.

- Oracle Solaris 8 and 9

  You can install Oracle Solaris 8 and 9 as branded zones on Oracle Solaris 10 OS only.

  Download the images of Oracle Solaris 8 or 9 a and upload them to a software library. Refer to Requirements for Zones on Oracle Solaris 10 OS for more information about requirements for branded zones.

Click **Next** to specify the zone configuration details.

8. Specify the CPU type, either Shared or Dedicated. Set the memory thresholds and verify that the locked memory threshold value is less than or equal to the physical memory threshold.

   You can enable automatic recovery option for the zone. Set the priority of recovery value between 0 to 100. This value decides which zone is recovered first when the global zone fails.

   Click **Next** to specify the zone file system.

9. The zone is created with a default root file system which is the zone path. You cannot delete this file system or change its read and write access. You can add more file systems from original zone's list of file systems. Set the size and access to the file system.

   For each file system added, specify whether it is managed or unmanaged. When it is unmanaged, enter the mount point for the file system.

   The Reserved size is the size of the file system that the user can reserve. The Quota size is the maximum size that the file system can utilize.

   Click **Next** to configure the zone's storage.

10. Select the storage library for the zone and its metadata.

    You can store the zone metadata in the local or NAS library only. To migrate a zone, you must store the zone metadata in a NAS storage library.

    All the libraries that are available in the Oracle Enterprise Manager Ops Center are listed. Select the library type and the corresponding libraries are listed. Select the library and the virtual disk for the zone. You can select library types such as NAS, SAN, Local, Local Devices, and Dynamic Storage.

    Note: When you specify a local storage library, you cannot migrate the zone in the future.

    Caution: An Oracle Solaris Zone running with a zone path on a NFS share is not a supported configuration. When you specify this type of storage, do not use the zone for production or non-experimental workload.

11. When the library is local or NAS storage, specify the virtual disk name and size of the disk. For SAN library, select a LUN from the available list. You cannot change the size of the LUN and the size of the local devices.

    For Dynamic Storage, select the Dynamic library that is available for the discovered storage servers in Oracle Enterprise Manager Ops Center. You can select the existing LUNs or create new LUN. Select Create LUN in the drop-down

list of the column LUN/Virtual Disk Name and select the volume group. The size of the volume group is automatically displayed.

Click **Next** to specify the zone networks.

12. Select the networks that you want to connect to the zone. Also, specify the number of connections to the zone for each network. The actual binding of the networks takes place during deployment of a plan with this profile.

    Click **Next** to specify the zone setup parameters.

13. Specify the following setup parameters:

    - Language, time zone, terminal-type and root password for the zone.
    - Provide a domain name for the NFSv4 Domain Name or accept the default value dynamic to allow the naming service that you specify in Step 14 to determine the NFSv4 domain at run time.
    - Set the boot properties for the zone. You can set the zone to boot after it is created or whenever the global zone boots.
    - For Oracle Solaris 10 version, you can select whether to install the agent on the zone or not. When the agent is not installed, the global zone access the zone through zlogin.

    Click **Next** to specify the naming services.

14. Specify the naming service for the zone: DNS, NIS, NIS+ or LDAP naming service. You can also choose not to specify a naming service.

    - **DNS**: Enter the domain name of the DNS server and the IP address of the DNS server You can enter up to three IP addresses in the Name Server field. To specify additional domains to search, enter up to six domain names for the Domain Name Search List. The total length of each entry cannot exceed 250 characters.
    - **NIS and NIS+**: Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the Specify an NIS Server option to provide the NIS server host name and its IP address. When you do not have the NIS server information, select the Find an NIS Server option.
    - **LDAP**: Enter the domain name of the LDAP server. Specify the name of the LDAP Profile. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.
    - **NONE**: Select this option so that no name server is configured.

    Click **Next** to view the summary of the selected parameters for creating a zone profile.

15. Review the information and click **Finish** to create the zone profile.

The zone profile is created with version 1 and a corresponding deployment plan also with version 1. Apply the deployment plan to create one or more zones of consistent configuration.

Creating and Deploying Zone Plans

Apply the zone deployment plans on the required number of targets. When you apply the deployment plan, you must provide the resource assignments for storage and network. When there are resources that are provided from the profile are not available

or not accessible, it is flagged in red and you must re-assign the resources to continue further.

Use zone deployment plans to modify the parameters such as storage, networks, and zone name. The zone deployment plan is a single step plan which collects details of the number of zones to be created. You can also create your own zone deployment plans.

To Create a Zone Deployment Plan

1. In the Plan Management section, expand **Deployment Plans** and click **Create Oracle Solaris Zones**.
2. Click **Create Plan from Template** in the Actions pane.
3. Enter a name and description for the plan.
4. Select the failure policy.
5. In the Deployment Plan Steps, select the **Oracle Solaris Zone** profile.
6. Enter the number of zones to create.
7. Click **Save** to save the deployment plan.

When you modify the zone profile, you can choose to update the deployment plan with the correct version of the zone profile.

Zone deployment allows you modify the parameters such as storage, networks, and zone name. When you apply the zone deployment plan, resource assignments that are not available is marked in red and it must be corrected.

To Apply Zone Deployment Plan

1. Select the zone deployment plan and click **Apply Deployment Plan** in the Actions pane.
2. Select the target asset from the list and click **Add to Target List.**

   You can add more than one asset to the list.
3. Select how to apply the plan.

   You can either apply the plan with minimal interaction or override the profile values. When you select to override the profile values, you are taken through each step of the profile. Otherwise, you are directed to provide only the required resource assignments. In this procedure the minimal interaction is taken into consideration. Click **Next**.
4. In the Specify Storage Step, correct the storage resources defined in the profile if required.

   For example, when the selected storage libraries in the profile are not associated with the target global zone., then you must modify the storage resources accordingly.

   Click **Next** to specify the zone networks.
5. Designate the IP stack for the zone as Shared or Exclusive. Select the IP stack and the network list is updated accordingly.

   For Shared IP network:

   - Select a network from the list.
   - The address allocation method is Use Static IP only. This is fixed for shared IP networks.

- Enter the IP address. When you create more than one zone, enter the IP addresses either in the form of range, separated by comma or both. For example, 192.0.2.1 - 192.0.2.3, or 192.0.2.1,192.0.2.2, 192.0.2.3.

For Exclusive IP network:

- Select a network from the list.
- Enter the number of times the zone connects to the network. This is applicable only for exclusive IP networks.

Note: For Oracle Solaris 11 OS, you can connect to networks configured as exclusive IP stack only.

- Select the type of address allocation as Use static IP or DHCP allocated.
- When the allocation is static, then the IP range for the zones is proposed. The range depends on the number of zone and number of connections for each zone. You can modify the IP range as required.

Click **Next** to define the network resource assignment.

6. Modify zone host name and network assignment as required. You can also add or remove the network assignment to a zone.

   For each network connection of the zone, the NIC and the individual IP address is displayed. Each zone's host name is also editable. Modify the zone resource assignment as required.

   Click **Next** to schedule the job.

7. Select to run the zone creation job now or schedule to a later time.

   Click **Next** to view the summary.

8. Review the information provided for creating zones. Click **Apply** to start the zone creation tasks.

Zpool and File System of Zones

Creating zones with Oracle Enterprise Manager Ops Center results in the following operations which run automatically in the background:

- The file systems of the zones are implemented as ZFS file systems.
- A zpool is created for each zone with the name of the zone metadata.
- The storage that is allocated to the zone is pooled in a zpool and used by all the file systems.

For each deployed zone you can view the following information: file systems, zpools with reservations and quotas, the storage library used, and the file system usage. The Storage tab for a zone shows the zpools with reservations and quotas. The Analytics tab for a zone contains information on file system usage. Use the **Move Storage** action to change the reservation and quotas.

To View the File System and Storage Added to a Zone

1. Select the zone in the Navigation pane.
2. Click the **Storage** tab.

As shown in Figure 18–7, the Storage tab has two tables: a File Systems table and a ZPool and Storage table. The file system table displays the file system, the real path of the file system on the global zone, the amount of Reserved GB, the quota in GB, and the Access permissions for the zone.

Figure 18–7 Storage Tab for a Zone

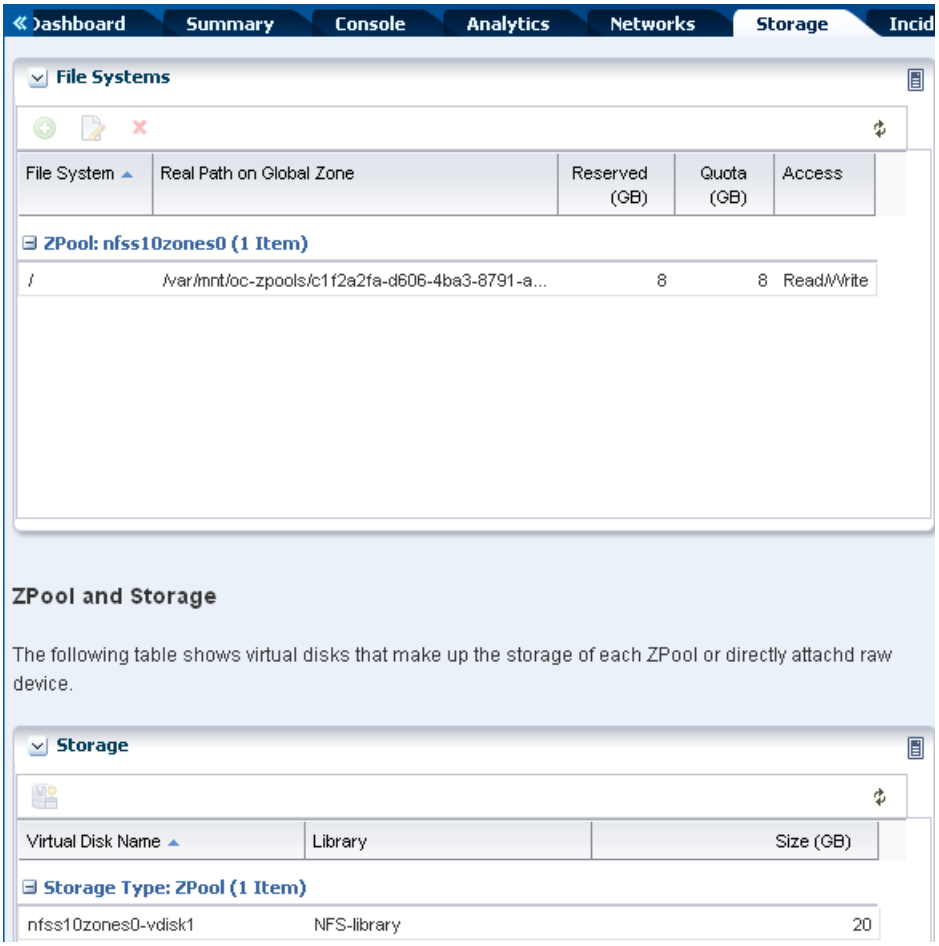

To Change the Default Reservation and Quota

1. Select the zone in the Navigation pane.
2. Click the **Storage** tab.
3. Click **Move Storage** in the Actions pane.

To View the File System Usage for a Zone

1. Select the zone in the Navigation pane.
2. Click the **Analytics** tab, then click the **Metrics** subtab.
3. Click **File Systems** in the Navigation pane.

Figure 18–8 Zone File System Metrics

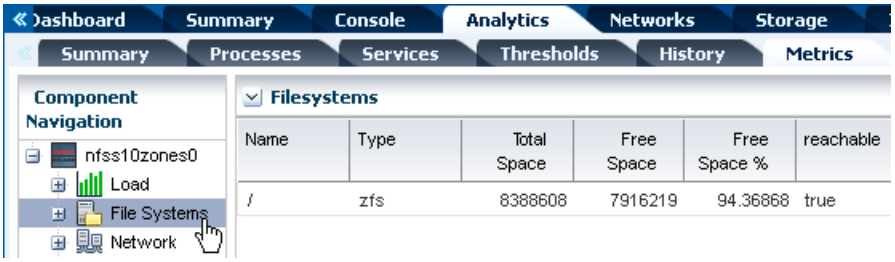

Modify Zone Configuration

You can modify the configuration of a zone to change the CPU or memory resources, scheduler, and the identity of the zone. When you modify the CPU and memory configuration, you must reboot the zone for the changes to take effect.

You can modify the following properties of a zone:

- **CPU Model**: The CPU model can be shared or dedicated. For a shared CPU, you can modify the number of CPU shares allocated to the zone. For a dedicated CPU, you can edit the minimum and maximum dedicated CPU numbers. You can also set the relative importance of the zone compared to other zones when contending for CPUs.
- **Memory Caps**: You can specify caps on various aspects of memory used by the zone. You can modify the caps set for physical, swap, and locked memory.
- **Scheduler**: For shared CPU model, the scheduler is assigned to Fair Share Scheduler (FSS). For a dedicated CPU model, you can set the following scheduler attributes:
  - Fair Share Scheduler (FSS)
  - Fixed Priority (FX)
  - Interactive (IA)
  - Real-time (RT)
  - Timer Sharing (TS)
- **LightWeight Processes (LWP)**: You can set the maximum number of LWPs simultaneously available to a zone.
  - **Message IDs**: Set the maximum number of message queue IDs.
  - **Semaphore IDs**: Set the maximum number of semaphore IDs.
  - **Shared Memory IDs**: Set the maximum number of shared memory IDs.
  - **Shared Memory**: Set the maximum amount of shared memory.
- **Automatic Recovery**: Set the value of priority of recovery. When the server fails, the zone with highest priority is recovered first. See Automatic Recovery for more information on how automatic recovery works.

Modify Zone Attributes

Use the option **Edit Attributes** to modify the description and tags of the zone. You can also add new tags to the zone.

Creating and Deploying Zones on a Logical Domain

Oracle Enterprise Manager Ops Center has a consolidated view of a managed Oracle VM Server for SPARC Control Domain and associated logical domains and I/O resources. You can create zones, on the domain using the domain's I/O resources that are not already in use by other assets.You can use the Oracle Enterprise Manager Ops Center UI to manage, update, and delete the zones.

Note: When you create a zone on a logical domain, the logical domain cannot be a root domain that is dedicated to provide I/O resources exclusively to an Oracle VM Server for SPARC Server Pool.

To manage zones on an Oracle VM Server for SPARC logical domain, you must deploy the Zone VC Agent on the logical domain and the Oracle VM Server for SPARC VC Agent on the associated Control Domain. See Virtualization Agent Controllers in Getting Started with Virtualization for information about Zone VC and Oracle VM Server VC Agents.

When you create a zone on logical domain, storage resources are assigned to the zone and are no longer available to use. The storage resource is not available to create other zones or to create a virtual server in exclusive mode. The zone's alternate MAC address cannot be assigned to a new VNIC or used to create a new zone.

When you delete the zone, the storage resources that are assigned to the zone are put back into the available storage resources and the alternate MAC address is available.

Managing Zones

After creating the zones, you can perform various operations such as boot, reboot, shut down, and halt on the zone. Oracle Enterprise Manager Ops Center provides all these options to be performed from the UI. The UI updates the status of the zone when you perform these operations from the CLI.

Select the global zone Summary tab in the center pane. The actions available for the zone are displayed as in Figure 18–9.

Figure 18–9 Zone Management Functions

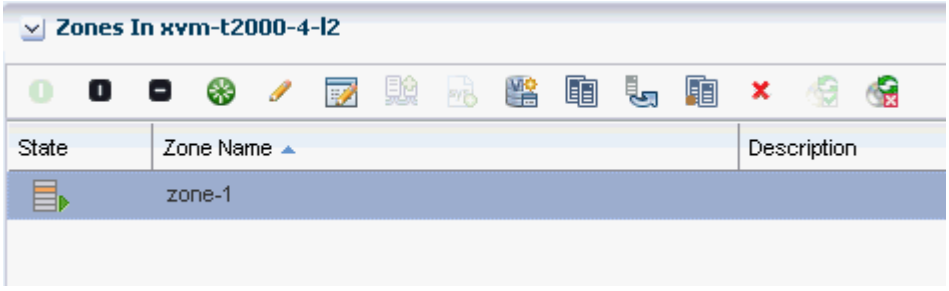

The same actions are available in the Actions pane when you select the zone.

Figure 18–10 Zone Actions Pane View

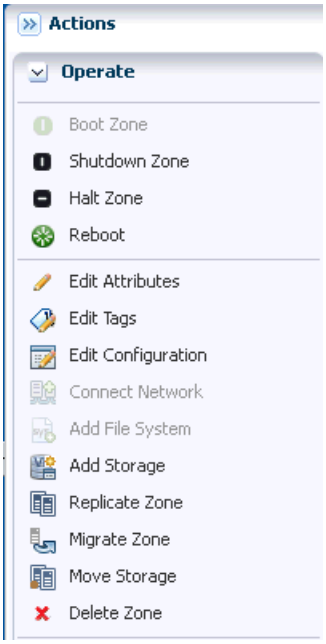

Simple zone operations that can be performed from the UI are as follows:

Boot Zone

Booting a zone places the zone in the running state, using the current configuration. This option is enabled only when the zone is in the ready or installed state. The zone boots whenever the global zone boots depending on the autoboot properties set during zone creation.

Reboot Zone

You can reboot a zone that is in the running state. The zone is shut down and then booted. This is different from the `zoneadm reboot` command in which the zone is first halted and then booted.

Halt Zone

When you halt a zone, it removes both the application environment and the virtual platform of the zone. Halting a zone changes the zone's state to Installed, all processes are killed, devices are unconfigured, network interfaces are unplumbed, file systems are unmounted, and the kernel data structures are destroyed.

Shutdown Zone

Shut a zone down in a graceful manner so that it is in a state that can be restarted.

Delete Zone

When you delete a running zone, the zone is halted, uninstalled, then deleted from the global zone. The following changes are also made:

- Zone root file system is deleted.
- Other file systems that were added to the zone are deleted.
- Zone metadata is deleted from the storage library.
- The zpool for the zone is deleted and the storage is made available.

- Exclusive IPs that were assigned to the zone are available for re-use.

Replicating Zones

Use this option to copy an existing zone so that you can provision a new zone on the same system efficiently. The process of cloning a zone is similar to the process of creating a zone because you can accept each of the original zone's specification or change it before you create the copied zone.

To Clone a Zone

1. Select the zone that you want to clone in the Assets section.
2. Click **Replicate Zone** in the Actions pane.

   The Replicate Zone Wizard is displayed.
3. Enter a different zone name and description for the zone.
4. Select a library from the list of libraries that are associated with the global zone to store the cloned zone's image and metadata. Click **Next**.
5. The CPU shares that are allocated in the original zone are displayed. You can edit the changes for the cloned zone. Click **Next**.
6. Accept or change the attributes of the original zone: language, time zone, terminal type, host name and root password. Click **Next**.
7. Set the boot properties for the cloned zone. You can set the properties so that the new zone boots after it is created and whenever the global zone boots. Click **Next**.
8. Configure the file systems for the cloned zone. The new zone has a default root file system which is the zone path. You cannot delete this file system or change its read and write access. You can add more file systems from original zone's list of file systems. Accept or change the size and access to the file system. The Reserved size is the size of the file system that the user can reserve. The Quota size is the maximum size that the file system can utilize. Click **Next**.

   Note: Make sure that the cloned zone has the same or more size for its file system than the size of the original zone's file system. When the cloned zone's file system is smaller than the original zone's file system, the clone operation cannot complete and the job fails. Do not modify the root file system of the new zone to a size less than the source zone root file system. Do not modify the system file system if it is defined in the source zone.
9. Accept or change the storage library. The library can be either a NAS storage or Fibre Channel library. When the library is NAS storage, specify the virtual disk name and size of the disk. For SAN library, select a LUN from the available list of LUNs in the library. The size of the selected LUN is displayed. You cannot change the size of the LUN.
10. The accumulated size of the storage is displayed as Currently Accumulated Storage. The required storage is displayed as Recommended Storage Size. Click the **Add** icon to configure more storage resources to the zone if the Currently Accumulated Storage is less than the Recommended Storage Size. When you have defined the Currently Accumulated Storage size as at least equal to the Recommended Storage size, click **Next**.

11. Assign at least one network to the zone. Select the zone IP type as Shared or Exclusive.
    - For Shared IP networks:
      a. Select a network from the list of networks that use the Shared IP mode and are assigned to the global zone.
      b. Select a NIC from the list of shared Network Interface Cards (NIC).
      c. Specify the management interface for the N IC. When the network has a defined IP range, the Auto Allocate IP option is displayed with the zone's IP address from the range. When the network does not have a defined IP range, the Use Static IP option is displayed. Enter an IP address for the zone in the Zone IP field.
      d. Click **Next**.
    - For Exclusive IP networks
      a. Select a network from the list of networks that are assigned to the global zone and are not used by other zones.
      b. Select a NIC from the list of the selected network's NICs that are not bound or assigned to other networks.
      c. Specify the management interface for the NIC. When the selected network has a defined IP range, the Auto Allocate IP option is displayed in the Management Interface and the zone's IP address is populated with an IP address from the defined range.

         When the network does not have a defined IP range, you must either provide the IP address or specify a DHCP server to provide one. To provide an IP address, select the **Static IP** option and enter the IP address in the Zone IP field. To designate a DHCP server, select the **Assigned by ext. DHCP** option. The Zone IP field contains the Automatically Allocated value.
    - Click **Next**.
12. Click **Finish** to launch the job for replicating a zone.

Adding Storage to Zones

You can add storage to zones dynamically. The storage is added to the zpool that is created for the zone. You cannot delete the storage from the zpool.

The storage libraries that are associated with the global zone are available for the zone. You can assign the following types of libraries to a zone:

- File system storage: Use NAS libraries
- Block storage: Use SAN or Dynamic storage libraries
- Local storage: Add the local storage on the global zone and local devices that are attached to the system

See Chapter 16, "Storage Libraries for Virtualization" and Chapter 5, "Software Libraries" for more information about setting up your storage servers and software libraries.

To Add Storage to Zones

1. Select the zone in the Assets section.

2. Click **Add Storage** in the Actions pane. As an alternative, click the **Storage** tab and click the **Add Storage to Zone** icon under Zpool and Storage.

   The Add Storage to the Zone window is displayed.

3. Click the **Add** icon to add storage.

   The storage libraries associated with the global zone are displayed in the Library list. The library can be a local storage, NAS, local device, Dynamic Storage, or SAN library.

4. Select a library from the list.

5. The selection of virtual disks varies according to the library selected.

   - Local and NAS library: Enter a virtual disk name and specify the size of the disk.
   - For SAN and iSCSI libraries: Select a LUN from the list. The size of the LUN displays automatically.
   - For Dynamic Storage Library: Either select a LUN from the list or create new LUN. When you create a new LUN, select the volume group from the list and specify the size of the new LUN.

   The Additional Storage Specified shows the additional storage added to the zone.

6. Click **Add Storage** to add the specified storage to the zone.

Moving Zone Storage

Some management and monitoring functions of the zone might be disabled depending on the managed resources of the zone. For example, the zone migration option is disabled when the zone file system is not on a shared storage library in Oracle Enterprise Manager Ops Center.

Note: A shared storage library in Oracle Enterprise Manager Ops Center is one that is accessible by the server and operating system. It is not related to Zones on Shared Storage in Oracle Solaris 11.1.

To enable the storage management and monitoring functions of the zone, you must move the zone's existing storage to a shared storage library. This provides the following options to you to manage the storage and zpools:

- Convert the local storage to a shared storage library so that you can migrate zones.
- Convert the experimental NFS file systems into a local or a shared storage library.
- Combine small ZFS pools per zone into a single ZFS pool.
- Split a single ZFS pool hosting zones into ZFS pools dedicated to each zone.

Oracle Enterprise Manager Ops Center provides the Move Storage option to convert the storage of the zones. You can use this option to manage your unmanaged storage resources of the zone.

Note: You cannot move managed storage to unmanaged storage and you cannot alter the root file system of the zone.

When you discover and manage existing zone environments in Oracle Enterprise Manager Ops Center, the storage is considered to be unmanaged. You can move this storage from unmanaged state to managed using this option.

To Move Zone Storage

1. Select the zone for which you want to move the storage.
2. Click **Move Storage** in the Actions pane.

   A warning message is displayed that the zone is shutdown during this process. You can use this option to relocate the zone metadata to a shared storage library and copy the selected unmanaged file system to the dedicated zpool of the zone.

Figure 18–11 Warning Message for Move Storage

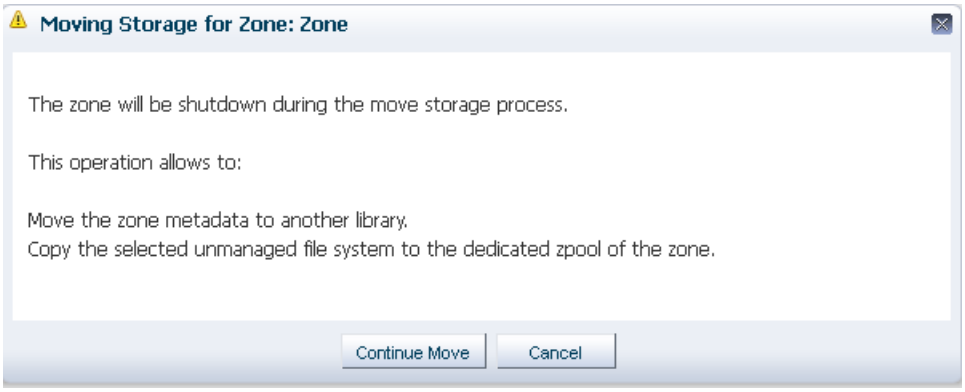

3. Click **Continue Move** to proceed with the continue the process.

   The Move Storage Wizard is displayed. The zone details are displayed.
4. You can edit the description and tags of the zone.
5. If required, you can change the library in which you have stored the zone metadata. All the libraries associated with the global zone are listed. Select the library from the list. Click **Next**.
6. Select the file systems that you want to be managed.
   - When there are unmanaged file systems in the zone, you can select the option Managed and provide the Reserved and Quota size of the file system. These are added as new ZFS file systems on the existing zpool for the zone.
   - When you want to relocate the root file system to a managed storage, a dedicated zpool is created for the zone and you must configure the virtual disks for the file systems.
   - You can also modify the size of the file systems except for the root file system.
   - You can keep some unmanaged file systems and relocate only those you require.
   - You cannot relocate managed storage to unmanaged storage.

   Click **Next** to configure the virtual disk storage when you move the root file system from unmanaged to managed state.
7. Select the library and the virtual disk for the zone's storage.

   The libraries associated with the global zone are listed. Select NAS, SAN or Dynamic Storage library from the list.

For NAS library, provide a virtual disk name and enter the size of the disk.

For SAN library, select a LUN from the list.

For Dynamic Storage library, either select a LUN or create new LUN from the list. When you create a new LUN, select the volume group and enter the size of the LUN.

Click **Next** to view the summary.

8. Review the information and click **Finish** to change the zone storage from unmanaged to managed.

Adding File Systems to Zones

You can add file systems to zones. The zone must be in a shut down state to add file systems. The storage source for the file system can be managed or unmanaged storage source. For unmanaged storage source, you must provide the mount point of the storage.

Provide the Reserved and Quota size for the file system. Boot the zone for the changes to take effect. The file system is added to the existing zpool of the zone.

To Add a File System to a Zone

1. Select the zone in the Assets section.
2. Click **Add File Systems** in the Actions pane.

   The Add File Systems window is displayed.
3. Click the **Add** icon to add file system.
4. Enter the file system.
5. Select whether the storage is managed or unmanaged.

   When it is not managed, enter the mount point of the storage source.
6. Enter the Reserved and Quota size for the new file system.
7. Click **Add File Systems**.
8. Boot the zone for the changes to take effect.

Connect and Disconnect Networks

You can connect and disconnect networks from non-global zones. The networks attached to the corresponding global zone are available for the zones. The shared IP zones can connect to only networks that are specified for shared IP mode. The exclusive IP zones can connect to only exclusive IP networks.

Select the network and click the connect or disconnect icon as required.

Enabling Automatic Recovery for Zones

Use the options Enable Automatic Recovery and Disable Automatic Recovery to set the recovery option of created zones. To set automatic recovery for the zone, select the option Enable Automatic Recover. Edit the zone configuration to set the priority of recovery. The zone with highest value is recovered first. See Recovering Zones to manually recover the zones. See Automatic Recovery for more information about how automatic recovery works.

Migrating Zones

In Oracle Enterprise Manager Ops Center, zone migration is a cold migration because the zone is shut down, all applications are stopped, the migration occurs, and then the zone is restarted.

To enable migration action for a zone, the zone storage must be on a shared storage library in Oracle Enterprise Manager Ops Center that is accessible by the server and operating system. When a zone uses local storage, use the Moving Zone Storage option to change the storage from local to shared.

Beginning with Oracle Solaris 10 10/08, the `zoneadm attach` command updates the zone to match the destination global zone during migration. This option is not available for branded zones. To migrate branded zones, both the source and destination global zones must have the same patch level.

When you migrate a non-global zone with a network attached to the global zone, the software adds a router entry for the network on the destination zone. If the non-global zone has several different types of networks (such as management, public, and private) attached to the global zone, the software adds a router entry on the destination zone for each network. You can edit a system property to disable this feature. When you disable the feature, a default router is not configured when creating or migrating a zone.

You must have Ops Center Admin permissions to disable the property. To disable the system property, go Administration->Configuration->Virtualization' in the UI. Set

the property 'ZoneDefaultAddRouter' to false.

This section describes the procedure for the following types of migration:

- Migrating an Oracle Solaris system into a new non-global zone. This is referred to as physical to virtual (P2V) conversion. See Migrating a Physical Oracle Solaris System into a Zone.
- Migrating a non-global zone from one global zone to another global zone, when all components are managed by Oracle Enterprise Manager Ops Center. See Migrating Zones to a Different Machine.
- Migrating a non-global zone that has dependencies that are not managed by Oracle Enterprise Manager Ops Center. This process uses a script to perform the migration. To migrate unmanaged file systems, see Script to Migrate a Zone With Dependencies.

Note: When you migrate a zone, the Alternate Boot Environment (ABE) is not supported.

Disabling the Automatic Router Assignment

When you migrate a non-global zone with a network attached to the global zone, the software adds a router entry for the network on the destination zone. If the non-global zone has several different types of networks (such as management, public, and private) attached to the global zone, the software adds a router entry on the destination zone for each network. You can edit a system property to disable this feature. When you disable the feature, a default router is not configured when creating or migrating a zone.

Migrating a Physical Oracle Solaris System into a Zone

A physical to virtual (P2V) conversion moves an existing Oracle Solaris system into a new non-global zone on the target system's global zone.

Figure 18–12 shows the steps for a P2V conversion.

Figure 18–12 Workflow for Migrating an Existing Oracle Solaris System Into a Zone

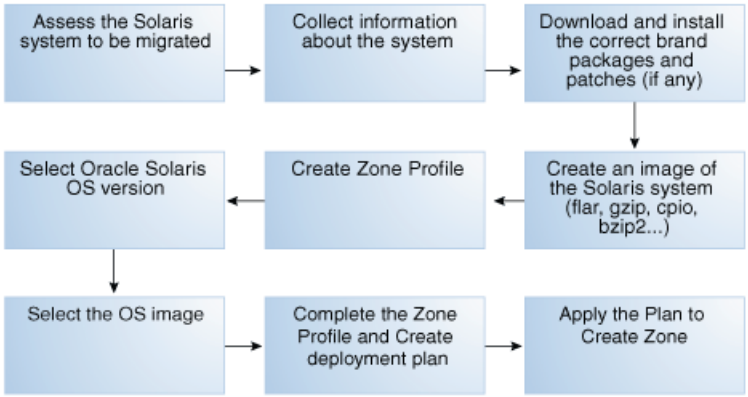

Ensure the following before migrating the zone:

- Assess the system to be migrated and collect information.

  See http://docs.oracle.com/cd/E19683-01/817-1592/ for information about collecting information about the source system.
- The system image to be installed in the non-global zone must not be newer than the target global zones's operating system release or the installation fails.
- The destination global zone must be running at least Oracle Solaris 10 8/07 OS.
- The supported branded zones are Oracle Solaris 8 and 9. For the branded zones to be supported on the destination global zone, remove the following brand packages from the global zone:
  - `SUNWs8brandu`
  - `SUNWs8brandr`
  - `SUNWs9brandu`
  - `SUNWs9brandr`
- For branded zone migration, the target and source global zones must have the same patch levels. The `zoneadm attach` option to update the branded zone to match the target global zone patches and packages is not available.
- When migrating a branded zone, you must provide an Address Allocation Method when you specify the network interface.

For instructions to download and install the correct brand packages, refer http://docs.oracle.com/cd/E22645_01/index.html

Migrating Zones to a Different Machine

Using Oracle Enterprise Manager Ops Center you can migrate one or more zones simultaneously. You can either migrate the zone to an individual global zone or to a

zones server pool. When you migrate zones to a server pool, the target global zone depends on the server pool placement policy.

When you migrate a zone, perform a trial run before submitting the migration job to verify that the target global zone has the correct configuration to host the non-global zones.

The target global zone must have same or later versions of the following operating system packages and patches that are installed on the non-global zone.

- Packages that deliver files under an `inherit-pkg-dir` resource
- Packages where `SUNW_PKG_ALLZONES=true`

Other packages and patches, such as those for third-party products, can be different.

When the source and target global zone do not have the same patches and packages during migration, either update the zones patches and packages to match the target global zone or migrate without updating the zone patches and packages.

> **Note:** You cannot migrate branded zones if the source and target global zones have different patch levels. The option to update on attach capability is also not available for branded zones migration.

When the target global zone has later versions of the zone-dependent packages or patches, update those packages in the non-global zone before the migration to match the target global zone. When the target global zone supports the update on attach capability, it checks the non-global zone for packages that must be updated and only those packages are updated. The rest of the packages, and their associated patches, can vary from zone to zone.

You cannot downgrade the patches and packages of the zones to a lower version. In such cases, the update option fails.

The patches that must be backed out of the zone before the update are also listed. You must remove the patches manually and run the migration job again.

You can change the Name, Description, Tags, and NIC details of the non-global zone but, when you change other configuration, a warning message indicating that to change the zone configuration you must execute the `sys-unconfigure` command on the zone.

After you migrate a zone, you cannot use an alternate boot environment (ABE) to upgrade the zone.

Compatible Global Zones for Migration

When you click the Migrate Zone option in the Actions pane, Oracle Enterprise Manager Ops Center checks all global zones for compatibility with the source global zone. When there are compatible global zones, the Migrate Zone Wizard is displayed.

Otherwise, the following pop-up window is displayed.

Figure 18–13 Message Displayed When No Global Zones are Compatible

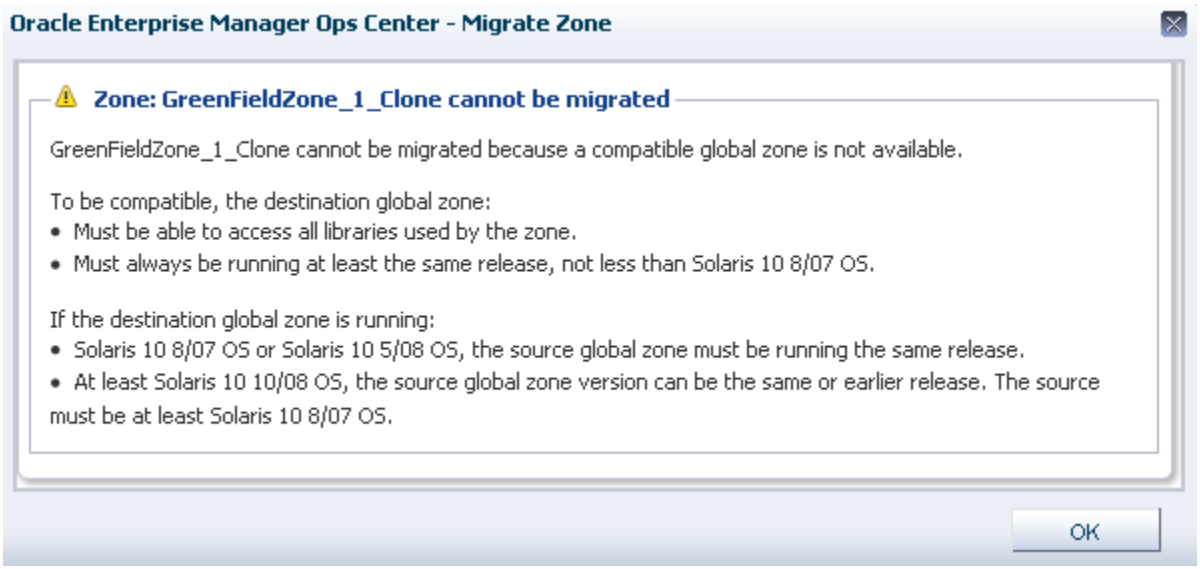

Table 18–5 shows the compatible global zones, depending on the operating system release.

Table 18–5 Compatible Global Zones

Source Global Zone	Compatible Target Global Zone
Oracle Solaris 10 8/07	Oracle Solaris 10 8/07
	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 5/08	Oracle Solaris 10 5/08
	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 10/08	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 5/09	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11

Table 18–5 (Cont.) Compatible Global Zones

Source Global Zone	Compatible Target Global Zone
Oracle Solaris 10 10/09	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 9/10	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 8/11	Oracle Solaris 10 8/11
Oracle Solaris 11	Oracle Solaris 11
	Oracle Solaris 11 Update 1
Oracle Solaris 11 Update 1	Oracle Solaris 11 Update 1

Verify the following conditions:

- The source and target global zones are compatible.
- The source and target global zones have access to all the libraries associated with the non-global zone.
- The non-global zone's metadata is not stored in the source global zone's local library.
- The non-global zone's data is not stored in the source global zone's local library. The zone must use a shared storage library.
- The non-global zone is in the running state.
- For a server pool, it must have at least one compatible global zone for migration.
- When you migrate the zone, a warning message indicates that when you change the zone configuration except for Name, Description, Tags, and NIC details, you must execute the `sys-unconfigure` command on the zone. Also, when the NIC names are changed for a zone that uses exclusive IP mode, the `/etc/hostname`. *itf* and `/etc/dhcp`.*itf* file are renamed accordingly.

When the target global zone does not support backout on attach capability, you must remove or downgrade the patches and packages manually before continuing with the migration. Create an update profile that includes the patches that must be removed. Run an update job with this update profile. Repeat the migration job on the zone after the removal of the patches and packages.

To Migrate a Zone

1. Select the zone in the Assets section.
2. Click **Migrate Zone** in the Actions pane.

   A warning message indicating that when you change the zone configuration except for Name, Description, Tags, and NIC details, you must execute the `sys-unconfigure` command on the zone. Also, when the NIC names are changed for a zone that uses exclusive IP mode, the `/etc/hostname`.*itf* and `/etc/dhcp`.*itf* file are renamed accordingly.
3. Click **Continue Migration**.

   The Migrate Zone Wizard is displayed.
4. Select an individual global zone or server pool as the target.

The compatible global zones list the number of zones that are running, total CPUs, and available dedicated CPUs. The server pools list the average usage of CPU and memory.

Click **Next**. The zone migration test starts.

5. Review the migration test result. Select an update option to continue the migration on the target global zone or server pool:

   - Update the patches and packages of the zone to match the target and then migrate the zone.
   - Migrate the zone without updating its patches and packages.

   Click **Save Test Result As** to save the migration test result. When you want to change the target, click Previous and select another target to run the migration test.

   Click **Next** to review the zone identification.

6. If the zone name exists in the target global zone, you must change the zone name.

   Click **Next** to specify the zone setup.

7. Select the language, time zone, and terminal type for the zone. The host name is the zone name as defined.

   The dynamic value for NFSv4 domain name enables the domain name to be derived dynamically from the naming service configuration. To hard code the value for NFSv4 domain, provide a domain name.

   Leave the password fields empty to use the existing password. Click **Next** to define the network interfaces.

8. You must have at least one network interface for the migration to continue. The network interfaces that are not accessible to the target global zone are displayed in yellow. Specify a new network interface for the inaccessible networks or click **Do Not Connect** to that network.

   a. Select the network which is marked in yellow color. The network interface details are displayed under Network Interface.

   b. Select a new network interface from the Network list or click **Do Not Connect**. The selected network can be either shared or dedicated. Enter the required network interface information for the selected network.

   Click **Next** to specify the naming service.

9. Specify the naming service for the zone: DNS, NIS, NIS+ or LDAP naming service, or choose to not specify a naming service. Click **Next**.

   - **DNS**: Enter the domain name of the DNS server and the IP address of the DNS server You can enter up to three IP addresses in the Name Server field. To specify additional domains to search, enter up to six domain names for the Domain Name Search List. The total length of each entry cannot exceed 250 characters.
   - **NIS and NIS+**: Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the Specify an NIS Server option to provide the NIS server host name and its IP address. When you do not have the NIS server information, click **Find an NIS Server**.
   - **LDAP**: Enter the domain name of the LDAP server. Specify the name of the LDAP Profile. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.

- **NONE**: Select this option when you do not want to configure a naming service.

10. Review the summary of the migration job. When there are no patches and packages to be backed out, continue with the migration. When there are patches or patches that must be removed, a warning is displayed.
    - When the target global zone supports backout on attach capability, a warning is displayed that the list of patches have been removed or downgraded.
    - When the target global zone does not support backout on attach, the patches and packages on the zone must be removed or downgraded manually so that the source zone matches the destination global zone. Cancel the migration or select a new target global zone.
11. Click **Finish** to submit the migration job.

To Migrate Multiple Zones

1. Select the global zone from which you want to migrate the zones.
2. Click **Migrate Zones** in the Actions pane.

   The Migrate Zones Wizard appears. The list includes the zones running in the global zone.
3. Select one or more zones from the list. Click **Next**.
4. Select an individual global zone or server pool to be the destination for the zone migration.

   The table displays the list of eligible global zones and server pool to which you can migrate the zones. The target global zone in the server pool depends on the server pool placement policy.
5. Select an update option to continue with migration.

   The source and the target global zones might not be in the same patch level. Either select to update the patches and packages of zone to match the target global zone or continue migration without updating the zone.
6. Review the details and click **Finish** to migrate the zones.

Script to Migrate a Zone With Dependencies

You can migrate zones even when the zone has dependencies that are not managed by Oracle Enterprise Manager Ops Center, such as when the zone uses storage that is not part of the zpool.

Scripts enable you to extend the zone migration feature to include the migration of dependencies, such as storage or other resources that are not managed by Oracle Enterprise Manager Ops Center. For example, unmanaged file systems are lost when you migrate a zone. To avoid this, use a script to migrate the file systems. See Example 18–1 for a sample script.

Develop your own scripts to migrate the dependencies and place them on the source and target global zone before migration.

The script for migrating the dependencies is executed in the following way:

- The migration job checks for a script placed on both the source and target global zones. The migration job is aborted when the script is found only on one of the global zones.

- When you create zones server pool, ensure that you upload the script to handle the unmanaged storage. The script is placed in all the global zones in the server pool and thus the migration of zone dependencies are taken care.
- The job checks whether the script has zero on exit. A non-zero exit is a failure and the migration job fails.
- The script is called on the global zone eight times when the migration job is executed.

Note: A new variable is available in Oracle Enterprise Manager Ops Center 12.2.2. You can use the OEMOC_AUTOMATIC_RECOVERY variable to determine if a script runs in the context of a migration or automatic recovery. Set the variable as follows:

- For migration, set the variable to false.
- For automatic recovery, set the variable to true.

Table 18–6 Script Call to Environment Variables

Call to the Script	Environment Variables
On the source global zone, to verify the script existence and user dependencies	OEMOC_ZONENAME=source zonename OEMOC_PHASE=VERIFY OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=SOURCE
On the target global zone, to verify the script existence and user dependencies	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=VERIFY OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=DESTINATION
On the source global zone, before the zone is shutdown	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_RUNNING OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=SOURCE
On the source global zone, after the zone is shutdown.	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=SOURCE
On the source global zone, after the zone is detached	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_SHUTDOWN_ DETACHED OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=SOURCE

Table 18–6 (Cont.) Script Call to Environment Variables

Call to the Script	Environment Variables
On the target global zone, before the zone is attached	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=POSTOPERATION_SHUTDOWN_ DETACHED OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=DESTINATION
On the target global zone, before the zone is started.	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=POSTOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=DESTINATION
On the target global zone, after the zone is started	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=POSTOPERATION_RUNNING OEMOC_OPERATION=MIGRATION OEMOC_JOBID=ID of the job running the script OEMOC_TARGET=DESTINATION

During rollback phase, the same sequence is executed in reverse order. The environment variables values change accordingly:

- PREOPERATION becomes POSTROLLBACK
- POSTOPERATION becomes PREROLLBACK

Script Requirements

The script must be executable and follow these conventions:

- Good error checking and clean-up within the script
- Standard exit code conventions
- Non-zero exit indicates that an error has occurred
- Informational messages are in stdout
- Error messages are in stderr
- The script must be named as guest-operations and placed in the `/var/opt/sun/oc/public` directory on both the source and target global zone.

See the Sample Script for migrating the additional inherited file systems of an adopted zone. The sample script is based on the assumption that `/opt/` file system is the additional inherited file system. The script re-configures `/opt/` file system on the target global zone before rebooting the zone. The script reconfigures the user-inherited file systems during POSTOPERATION_NOTRUNNING phase

Example 18–1 Sample Script

```
#!/bin/sh

do_migration_action() {

 echo "executing migration action for zone ${OEMOC_ZONENAME}"
```

```
 if [ ${OEMOC_PHASE} != "POSTOPERATION_NOTRUNNING" ]; then
       exit 0
 fi

 #
 # for migration action, before starting the zone on the target GZ
 # add to /opt inherited filesystem.

 /usr/sbin/zoneadm -z ${OEMOC_ZONENAME} detach >/dev/null 2>&1
 if [ $? -ne 0 ] ;then
       echo "unable to detach ${OEMOC_ZONENAME}"
       exit 1
 fi

 echo "add inherit-pkg-dir" >/tmp/tmp-$$.txt
 echo "set dir=/opt" >>/tmp/tmp-$$.txt
 echo "end" >>/tmp/tmp-$$.txt

 /usr/sbin/zonecfg -z ${OEMOC_ZONENAME} -f /tmp/tmp-$$.txt >/dev/null 2>&1
 if [ $? -ne 0 ] ;then
       echo "unable to add inherit filesystem for ${OEMOC_ZONENAME}"
       exit 1
 fi

 /usr/sbin/zoneadm -z ${OEMOC_ZONENAME} attach -u >/dev/null 2>&1
 if [ $? -ne 0 ] ;then
       echo "unable to attach ${OEMOC_ZONENAME}"
       exit 1
 fi
}

echo ${OEMOC_ZONENAME}
echo ${OEMOC_PHASE}
echo ${OEMOC_OPERATION}

if [ ${OEMOC_OPERATION} = "MIGRATION" ] ;then
       do_migration_action
fi
exit 0
```

Recovering Zones

When the global zone crashes or must be halted, its non-global zones can be migrated to another global zone. However, this zone recovery procedure is not the same as zone migration because Oracle Enterprise Manager Ops Center cannot get access to the information in the global zone and therefore cannot perform compatibility checks. The procedure to relocate non-global zones from the source global zone to the target global zone is a forced attachment of the non-global zone.

The zone recovery procedure uses the command-line interface for the Oracle Enterprise Manager Ops Center.

See *Oracle Enterprise Manager Ops Center Command Line Interface Guide* for instructions on getting access to the CLI and the available CLI commands.

Before You Begin

- Verify that the zone metadata is on the NAS storage library.

- Verify that both the source and target global zones have the same network connectivity.
- Verify that the source and the target global zones have the shared storage library.
- Verify that the target global zone has access to the same libraries associated with the non-global zone.
- Verify that the non-global zone's metadata and operational data is not stored in a local library.
- Verify that JDK version 6 or 7 is used. In the Enterprise Controller command prompt, enter the following command:

```
export JAVA_HOME=/usr/jdk/latest
```

To Recover Zones

1. Connect to the Oracle Enterprise Manager Ops Center CLI using the following command:

```
/opt/SUNWoccli/bin/oc
```

2. Connect to the local Enterprise Controller.

```
xvmSh > connect

localhost >
```

3. Enter the virtualization mode.

```
localhost > virtualization
localhost [virtualization] >
```

4. List the available global zones.

```
localhost [virtualization] > list_hosts

   Name      |    Type   |   Health   | Reachable |
ObjectName                          |                   UUID       |
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
  gzhost36    |     zone    |      OK     |    False    |
com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost           |
2b7c71ac-70ab-48a2-a2f2-ac291e580c39
  gzhost44    |     zone    |      OK     |    True    |
com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost-4          |
3b6c61ab-50ab-34a1-b2d2-bd253e632c45
```

5. List the zones that are running in the source global zone. For example, when the source global zone is gzhost36, then enter the following command.

> **Note:** Do not use the user-friendly name of the zone in the commands.

```
localhost [virtualization] > list_guests -C
com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost

     Name     |      Type     |   State   | Migratable |
ObjectName                                  |
-------------------------------------------------------------------------------
```

```
-------------------------------------------
 test      |         ZONE        |   RUNNING   |    True    |
com.sun.hss.domain:type=Server,name=NORM-07e91405-8313-43ec-9671-dc320989866e
```

6. Select the destination global zone and start the zone on it. For example, when the destination global zone is gzhost44, execute then execute the following command to start the test zone on it.

```
localhost [virtualization] >startup -Z <source global zone ObjecName> -D
<target global zone ObjectName>

localhost [virtualization] >startup -Z
com.sun.hss.domain:type=Server,name=NORM-07e91405-8313-43ec-9671-dc320989866e
-D com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost-4

submitted job : <Ecname>-1.17
```

A job is submitted. You can view the status of the job in the Jobs pane.

After the zone is migrated to the target global zone, the zone is a managed zone of the target global zone. The zone might continue to be displayed under source global zone. This is because the status of the source global zone is not updated. After the source global zone is rebooted, the zone does not appear as one of its managed zones.

Zones Server Pool

You can create server pools for zones in Oracle Enterprise Manager Ops Center. Pooling your virtualization hosts provides the capability for load balancing the virtualization servers, high availability and minimize power consumption.

See Chapter 21, "Server Pools" for information about creating and managing zones server pools.

Updating Zones

Oracle Enterprise Manager Ops Center enables you to update the global and non-global zones. You can also patch zones that are running on a supported configuration. The installation of the patches on the zones depend on the package parameters and the attribute set for the patch commands. This section describes the parameters for installation of the packages and patches. The concepts involved in updating global and non-global zones, and the procedures to update the zones are described in this section.

Note: When you use Oracle Solaris Live Upgrade to update the OS in a zone, you cannot use Oracle Enterprise Manager Ops Center to manage the zone. Alternate Boot Environment (ABE) is not supported for zones that were created using Oracle Enterprise Manager Ops Center.

Install Packages and Patches on Zones

A patch is a collection of files and directories that replace existing files and directories that are preventing proper execution of the software.

You can install packages and patches on a zone. The `patchadd` and `pkgadd` commands operate in the background to install a patch and package respectively. However, the installation of packages on zones also depends on the parameters SUNW_PKG_

ALLZONES, SUNW_PKG_HOLLOW, and SUNW_PKG_THISZONE. These parameters control whether a package can be installed on global zones or non-global zones. The actions for the parameters are as follows:

- **SUNW_PKG_ALLZONES**: If the value is true, the package is installed on all zones, both global and non-global.
- **SUNW_PKG_HOLLOW**: If the value is true, the package information is propagated to the non-global zones, but the package is not installed.
- **SUNW_PKG_THISZONE**: If the value is true, the package is installed only in that zone.

Configure patchadd and pkgadd Commands

In Oracle Enterprise Manager Ops Center, the `patchadd`, `pkgadd`, `patchrm`, and `pkgrm` commands are implemented without the `-G` switch by default. To install updates or packages only on the current zone, enable the `-G` switch by editing the `.uce.rc` file.

Note: Two files in the `/SUNWuce/agent/bin` directory have similar names. Verify that you are editing the `.uce.rc` file. Do *not* edit the `uce.rc` file.

Editing the .uce.rc File

1. Open the `.uce.rc` file in the `/SUNWuce/agent/bin` directory in the managed system.
2. Add the following lines to the `.uce.rc` file:

```
( all ) (invisible.__is_patchadd_g_specified, false)
( all ) (invisible.__is_patchremove_g_specified, false)
( all ) ( invisible.__is_pkgadd_g_specified, false)
```

3. Set the `-G` parameter to true for the action that you want to perform.
4. Save and close the file.
5. For this change to take effect, restart the services using the following commands:

```
svcadm disable -s update-agent
svcadm enable -s update-agent
```

Updating a Global Zone

In Oracle Enterprise Manager Ops Center, when a package or patch is installed, the `patchadd` and `pkgadd` commands are implemented in the background as shown in the following example:

```
patchadd <patchid>
pkgadd <pkgname>
```

Change the way that these commands are implemented by enabling the `-G` switch. You can enable the `-G` switch to cause the patch or package to be installed to the target zone only if the package parameter SUNW_PKG_THISZONE is set to true. See Install Packages and Patches on Zones for information about configuring the `patchadd` and `pkgadd` commands on the managed systems.

See the following scenarios when you are updating a global zone. The result for each scenario determines whether the update job is successful, depending on the package information.

Table 18–7 Updating a Global Zone Scenarios

SUNW_ PKG_ ALLZONES	SUNW_ PKG_ THISZONE	SUNW_PKG_ HOLLOW	Impact	Impact with -G Configuration
False	False	False	The package is installed on the global zone, and all the non-global zones	The package is installed only on the global zone.
True	False	False	The package is installed on the global zone and all the non-global zones.	The -G switch cannot override the SUNW_PKG_ ALLZONES parameter, and the package is installed on all the zones.
True	False	True	The package is installed on the global zone and the package information is made available on all the non-global zones.	The -G switch cannot override the SUNW_PKG_ ALLZONES parameter, and the package is installed on all the zones.
False	True	False	The package is installed only on the global zone.	The package is installed only on the global zone.

Patches are sets of updates to packages. When you install a patch, the patch is installed on the global zone and the non-global zones, depending on the package parameters as shown in the previous table.

Note: Use caution while enabling the -G option on a host with sparse zones. Packages that are inherited from the global zone that are not SUNW_ALL_ZONES cannot be patched within a sparse zone.

Updating a Non-Global Zones

A user with the virtualization administrator role can install packages and patches on non-global zones. The patchadd and pkgadd command must be used without options. Do not configure the -G switch to the commands while updating the non-global zones.

See the following scenarios when you are updating a non-global zone. The results of each scenario determine whether the update job is successful, depending on the package information.

Note: The -G switch does not have an effect on installing packages or patches in a non-global zone.

Table 18–8 Updating Non-Global Zones Scenarios

SUNW_ PKG_ ALLZONES	SUNW_PKG_ THISZONE	SUNW_PKG_ HOLLOW	Impact
False	False	False	The package is installed only on the target non-global zone.

Table 18–8 (Cont.) Updating Non-Global Zones Scenarios

SUNW_PKG_ALLZONES	SUNW_PKG_THISZONE	SUNW_PKG_HOLLOW	Impact
True	False	False	The package installation fails.
True	False	True	The package installation fails.
False	True	False	The package is installed only on the target non-global zone.

Note: When the patch is installed only on the non-global zone, ensure that autoboot property is set to true for the zone. Otherwise, single user mode patches fail to apply as the zone does not come up after the reboot.

Patches are sets of packages that must be installed. When one of the packages has the SUNW_PKG_ALLZONES parameter set to true, then the patch installation fails. For a successful patch installation, ensure that none of the packages have SUNW_PKG_ALLZONES parameter set to true.

Note: Packages that deliver to read-only inherit directories do not install on sparse root zones. These packages must be installed from the global zone with the `-G` switch disabled. When a package has the parameter SUNW_PKG_THISZONE=true, it does not appear as installed from the sparse zone and the software might not function correctly. In this case, a whole root zone must be used. Packages with SUNW_PKG_THISZONE=true must not deliver to read-only inherit directories.

Zone Parallel Patching

To view the number of zones that you can patch in parallel when you update the global zone, go to the global zone's Summary page. When the agent is installed on the Oracle Solaris OS, the number of zones that you can patch in parallel is calculated as 1.5 times the number of CPU cores in the server. For example, if you have a 2 CPU core machine, then you can patch three (3) zones in parallel. This is set and displayed on the UI.

Oracle Solaris 10 and 11 OS version displays this information in the Summary page. When the update capability is not enabled on the OS, then the zone parallel patching is also not enabled.

Related Zone Operations

Refer to the following chapters in this guide for operations that are common for many resources:

- See Chapter 16, "Storage Libraries for Virtualization" for setting up your storage resources.
- See Chapter 17, "Networks for Virtualization" for setting up your network infrastructure.

- See Chapter 8, "Plans and Profiles" for information about managing your zone profiles and deployment plans.
- See Chapter 21, "Server Pools" for information about creating and managing zones server pools.
- See Chapter 12, "Operating System Management" for information about monitoring your resources, setting up boot environments, and Agent Controllers.
- See Chapter 2, "Asset Management" for discovering and registering your assets.

Related Resources for Oracle Solaris Zones

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- For end-to-end examples, see the workflows and how to documentation in the Deploy How To library at http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm and the Operate How To library at http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm.

  See the following example workflows for preparing for, deploying, and using zones:

  - *Deploy Software Libraries Workflow*
  - *Deploy Storage Libraries Workflow*
  - *Deploy Networks Workflow*
  - *Deploy Hardware Workflow*
  - *Deploy Operating System Workflow*
  - *Deploy Oracle Solaris 10 Zones Workflow*
  - *Deploy Oracle Solaris 11 Zones Workflow*
  - *Operate Zones Workflow*
- See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about how to manage the software, including user and role management.
- For a list of the Oracle Solaris 11 documentation available in HTML and PDF formats, including the *Oracle Solaris 11.1 Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* and *Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management* documentation, visit the Oracle Solaris 11 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-11-192991.html.
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation website at http://www.oracle.com/technetwork/documentation/solaris-10-192992.html.
- The complete Oracle Solaris 10 documentation set is located at http://docs.oracle.com/cd/E19253-01/index.html.
- See http://docs.oracle.com/cd/E23824_01/html/E24456/transzone-1.html for information about creating Solaris Flash archive images of an Oracle Solaris 10 operating system.
- See http://docs.oracle.com/cd/E19082-01/819-6990/index.html for information about network interfaces and virtualization, and administration of your network interfaces.

- For a list of the Oracle Solaris 8 and 9 documentation, visit the Legacy Solaris Documentation website at http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html.

19

Oracle VM Server for SPARC

The following information includes:

Introduction to Oracle VM Server for SPARC

Oracle VM Server for SPARC technology enables server virtualization on SPARC platforms. You can create and manage multiple virtual machine instances simultaneously on a single SPARC machine. Each virtual machine, or guest, can run a different operating system.

Oracle VM Server for SPARC technology is virtualization of SPARC servers. This technology is part of a suite of methodologies for consolidation and resource management for SPARC Chip Multi Threading (CMT) systems. Using this technology, you can allocate the various resources of the system such as memory, CPU threads, and devices, into logical groupings and create multiple discrete systems. These discrete systems have their own operating system, resources, and identity within a single system. By careful architecture, an Oracle VM Server for SPARC environment can help you achieve greater resource usage, better scaling, and increased security and isolation.

When Oracle VM Server for SPARC software is installed, a domain called the control domain is created. From this control domain, you create virtual machines called logical domains that each run an independent OS. A logical domain is a virtual machine with resources, such as CPU threads, memory, I/O devices, and its own operating system. The control domain manages the logical domains. Each logical domain can be created, destroyed, reconfigured, and rebooted independently of other logical domains.

Logical domains that are created manually for an Oracle VM Server for SPARC provisioned through the UI are also automatically discovered and managed on the UI.

You can also provision the OS on the manually created logical domains using the OS provisioning profile and plans in Oracle Enterprise Manager Ops Center.

Note: When the Enterprise Controller and Proxy Controller are installed on an Oracle Solaris 11 operating system, you can perform more tasks than with other operating systems. For example, the Enterprise Controller must be installed on Oracle Solaris 11 to provision Oracle Solaris 11, Oracle Solaris 11 zones, and newer versions of Oracle VM Server for SPARC. In some cases, the Proxy Controller must also be installed on Oracle Solaris 11. See Enterprise and Proxy Controller Requirements for OS Provisioning for more information on operating system requirements for Oracle Solaris 11 actions.

For more information about Oracle VM Server for SPARC, refer to `http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html`.

Domain Types and Representation on the UI

You can create different types of logical domains. Depending on how the physical resources are assigned to the logical domain, distinguishes a logical domain from another domain. The different types are:

- **Control Domain**: The control domain is the first domain created when you install the Oracle VM Server for SPARC software. This is also called the primary domain and denoted as `primary` wherever applicable in the Oracle Enterprise Manager Ops Center UI. Only from the control domain, you can create logical domains.
- **I/O Domain:** An I/O domain has direct access to a physical I/O device, such as the PCIe Controller. An I/O domain either uses the physical I/O devices to host its own applications or shares the physical I/O device with other domains in the form of virtual devices.
- **Root Domain**: A root domain is also an I/O domain which has PCIe root complex assigned to it.

- **Guest Domain**: A guest domain is a non-I/O domain that uses virtual devices, such as virtual disks and virtual network interfaces, provided by one or more service domains.
- **HA Guest Domain**: A guest domain with redundant network and storage resources.
- **Service Domain**: A service domain provides virtual device services to other domains. This means that the domain has physical I/O devices. In the Oracle Enterprise Manager Ops Center UI, the service domain list includes the primary domain, I/O domains, and root domains that can provide virtual device services. It is generally recommended to not run any applications in service domains.

When you create logical domains using Oracle Enterprise Manager Ops Center, ensure to select the appropriate domain subtype. The selection of the domain subtype defines the step for selection of PCIe Endpoints for the I/O domain and PCIe buses for the root domain.

In Oracle VM Server for SPARC, an I/O domain can be created either by assigning PCIe Endpoints or SR-IOV Virtual Functions.

In Oracle Enterprise Manager Ops Center, an I/O domain is defined as Physical I/O domain, and created by assigning PCIe Endpoints only.

The subtype for the I/O domains is defined as Physical IO Domain in the UI. The I/O domain in this documentation refers to the Physical I/O Domain hereafter.

In the Oracle Enterprise Manager Ops Center UI, for I/O domains, the center pane displays the **I/O Resources** tab. The I/O Resources tab displays the details of the PCIe Endpoints assigned to the domain.

For root domains, the **I/O Resources** tab displays the PCIe buses assigned to the domain and the corresponding Endpoint devices in the bus.

Refer to View I/O Resources section for more detailed information.

Roles for Oracle VM Server for SPARC

The following table lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 19–1 Oracle VM Server for SPARC Tasks and Roles

Task	Role
Provision and manage virtualization host	Virtualization admin
Create, manage, update, and delete guests	Virtualization admin
Create, manage, and delete I/O domains and root domains	Virtualization admin
Discover and manage virtualization hosts	Asset admin
Create and manage profiles and plans	Profile and plan admin
Create and manage IPMP groups	Network admin
Create and manage Link Aggregation	Network admin
Set monitor threshold	Asset admin
Create credentials	Security admin

Actions for Oracle VM Server for SPARC

Using Oracle Enterprise Manager Ops Center, you can perform the following tasks:

- Provision Oracle VM Server for SPARC.
- Manage Oracle VM Server, including, rebooting, shutting down, changing name servers and NFS4 domains.
- Monitor the performance of Oracle VM Server.
- Create and provision logical domains, that includes I/O domains, root domains and guest domains.
- Manage logical domains, including editing, migrating, starting, rebooting, and shutting down domains.
- Monitor the performance of logical domains.
- Discover and manage existing Oracle VM Server for SPARC systems and logical domains.
- Create server pools to maximize capacity and for automatic load balancing.
- Migrate guest domains.
- Manage automatic recovery of guest domains in a server pool.

Oracle Enterprise Manager Ops Center uses special Agent Controllers, called virtualization controller agents, or VC Agents, to manage Oracle VM Server systems and Oracle Solaris Zones. See Virtualization Agent Controllers for more information about VC Agents.

Location of Oracle VM Server for SPARC Information in the User Interface

Oracle VM Server for SPARC assets are visible from the All Assets section in the user interface (UI). The server hardware, control domain, the control domain OS, the logical domains and its OS are displayed as shown in Figure 19–1.

Figure 19–1 Oracle VM Server for SPARC System in the UI

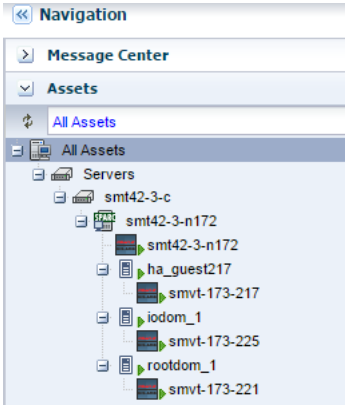

Table 19–2 lists where to find different information for Oracle VM Server for SPARC in the UI.

Table 19–2 Location of Oracle VM Server for SPARC Information in the UI

To See	Location
Control Domain	Expand **Assets** in the Navigation pane. The Oracle VM Server for SPARC asset is represented by the SPARC icon which is the control domain.
Logical Domains	Expand **Assets** in the Navigation pane. Select the control domain and all the logical domains created in that server are listed.
Virtual Services	Expand **Assets** in the Navigation pane and select the control domain. The center pane displays the **Virtual Services** tab that displays the Virtual Disks, Virtual Network Switches, and Virtual Console Concentrators details.
I/O Resources	Expand **Assets** in the Navigation pane and select the control domain. The center pane displays the **I/O Resources** tab that displays the Buses/Endpoint Devices and SR-IOV Services.
Other options	Expand **Assets** in the Navigation pane and select the control domain. The Actions pane lists the different options available for managing Oracle VM Server for SPARC. The options are Associate Libraries, Attach Network, Reboot and Edit Attributes.

Discovering Existing Oracle VM Server for SPARC Environments

The following information is covered in this section:

- Supported Logical Domain Configurations
- Limitations of Discovering an Existing Oracle VM Server for SPARC System with Logical Domains
- Discover and Manage Existing Oracle VM Server for SPARC Servers

You can use Oracle Enterprise Manager Ops Center to manage Oracle VM Server for SPARC systems created outside of Oracle Enterprise Manager Ops Center. There is no difference between the logical domains created using Oracle Enterprise Manager Ops Center and the domains created outside of Oracle Enterprise Manager Ops Center except for some limitations that are described in Limitations of Discovering an Existing Oracle VM Server for SPARC System with Logical Domains.

You can manage the following types of existing Oracle VM Server for SPARC servers:

- Oracle Solaris 11 running Oracle VM Server for SPARC server. The control domain can have non-global zones.
- Oracle Solaris 10 running Oracle VM Server for SPARC server. The control domain can have non-global zones.

To manage an existing Oracle VM Server for SPARC environment, use the Oracle Enterprise Manager Ops Center discovery feature to display the domains (control, root, I/O, and guest domains) in the user interface. The information about the I/O resources, and SR-IOV configurations appear in the UI.

Note: If you manually configured an Oracle VM Server for SPARC environment that you managed with an earlier version of Oracle Enterprise Manager Ops Center, the Zones VC Agent is installed. See Virtualization Agent Controllers for information about VC Agents and Changing the Type of Agent Controller for how to change the Agent Controller to the Oracle VM Server VC Agent. Information about agent controllers and functionality is also located in Chapter 12, "Operating System Management".

For an end-to-end example of how to discover and manage an existing Oracle VM Server for SPARC system, see *Oracle Enterprise Manager Ops Center Adding an Existing Oracle VM Server for SPARC*.

Supported Logical Domain Configurations

You can discover and agent manage an existing fully configured Oracle VM Server for SPARC system that has logical domains and associated network and storage resources that are configured on the control domain.

The logical domain configuration might be one of the following types with I/O resources or virtual resources:

- A root domain with one or more PCIe root complexes
- A physical I/O domain with one or more PCIe Endpoint devices
- A guest domain with virtual resources

Limitations of Discovering an Existing Oracle VM Server for SPARC System with Logical Domains

When you discover and deploy an agent on an existing Oracle VM Server for SPARC system that has associated logical domains, the system and logical domains appear in the Assets section of the UI. Some actions are disabled when the agent is not able to manage a resource.

Disabled Migration Action

The following are some reasons why the migration management action is disabled:

- Migration is disabled for all discovered guest domains because the metadata of the domain is in the default local library. To resolve, move the metadata to a shared library for all guest domains that you want to migrate. See Moving Metadata to Another Library for how to move the metadata.
- Migration is disabled for all guest domains that use storage on local devices or local files. To resolve, use shared FC or iSCSI LUNs for the guest domain's storage. Or enable shared action for the local file system storage

Note: You cannot migrate a guest domain until you move all of its virtual disks to shared storage and move the domain's metadata to a shared library. See Migration Requirements for more details.

Monitoring Information

The monitoring information might be incomplete due to missing managed resources. The following are examples of situations of missing resources:

- The networking view of logical domains shows only the networks that you discover with the control domain Virtualization Agent Controller.
- The networking view of a logical domain is incomplete when the domain is provisioned without an Agent Controller running on it.
- The storage view of a logical domain might be incomplete if its backing storage is not part of a library that is managed by Oracle Enterprise Manager Ops Center.

Agentless Management

Agentless management does not deploy the Agent Controller on the control domain, which restricts the management capabilities to the following:

- Perform basic state change command such as shutdown, start, reboot, and destroy.
- Edit the number of vCPUs, the number of Crypto units, and the memory size.
- Provision the Oracle Solaris operating system.
- Access the guest's serial console.
- View the network configuration and storage configuration for each guest.

With an agentlessly managed domain, you are unable to change the network and storage configurations, and therefore, unable to migrate the guests to another managed Oracle VM Server.

Discover and Manage Existing Oracle VM Server for SPARC Servers

You can discover and manage an Oracle VM Server for SPARC system that you manually provisioned. The underlying logical domains are automatically discovered and managed in Oracle Enterprise Manager Ops Center. In the Assets section, the discovered logical domains are displayed in order and grouped under the control domain.

You can manage the following versions of Oracle VM Server for SPARC systems in Oracle Enterprise Manager Ops Center:

- 1.2
- 1.3
- 2.0
- 2.1
- 2.2
- 3.0
- 3.1

> **Note:** To discover the server, the Logical Domains Manager must be running in the Oracle VM Server for SPARC system.

The following tasks are covered:

- To Agent Manage an Existing Oracle VM Server for SPARC System
- To Agentlessly Manage an Existing Oracle VM Server for SPARC System

See Chapter 15, "Getting Started with Virtualization" for information about agent controllers.

To Agent Manage an Existing Oracle VM Server for SPARC System

1. (Optional) Create SSH credentials. If you have not created SSH credentials, you can create them before creating an OS Discovery profile, or you can create them as part of the OS Discovery profile. See Using Management Credentials to create SSH credentials.
2. Create an OS Discovery profile. See Creating a Discovery Profile for more information.
   a. Expand **Plan Management** in the Navigation pane.
   b. Click **Discovery** in Profiles and Policies.
   c. Click **Create Profile** in the Actions pane.
   d. Enter a name and description for the discovery profile. Expand Operating Systems and select **Solaris, Linux OS**. Click **Next**.
   e. (Optional) Click **Next** to skip tags. See Using Tags for information about using tags in a Discovery profile.
   f. (Optional) Click **Next** to skip adding IP Ranges. You are prompted to provide IP addresses and host names during discovery.
   g. Click **Select**, select the SSH credentials that you previously created, then click **OK**. To create new credentials, click **New** and create the credentials as described in Using Management Credentials.

      Select the **Enable Oracle VM for SPARC management** check box to deploy the Agent Controller during discovery. Click **Next**.

Figure 19–2 Discovery Credentials for Oracle VM for SPARC Management

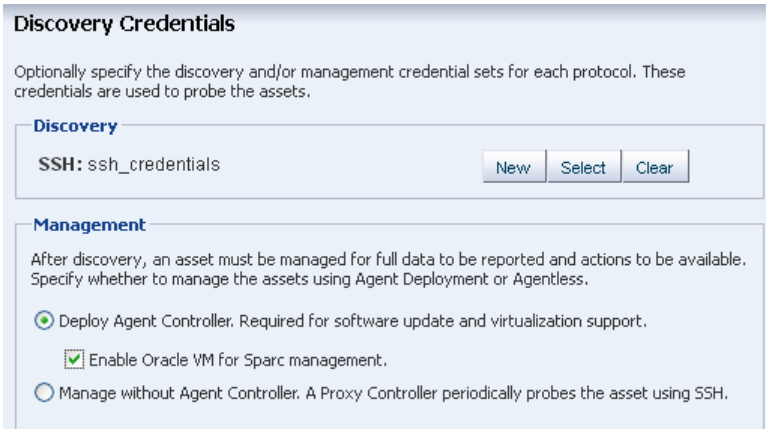

   h. Review the Summary page, then click **Finish**.
3. Add the assets using the OS discovery profile. See Adding Assets Using a Discovery Profile.
   a. Expand **Assets** in the Navigation pane, then click **Add Assets** in the Actions pane.
   b. Select **Add and manage various types of assets via discovery probes**, then click **Next**.
   c. Select the Discovery profile you created in the previous step. You are prompted to complete the host names or IP addresses and network.

d. Enter a comma-separated list of host names or IP addresses. Select a managed network with which the host is associated, or select Automatic to route the job to the most appropriate Proxy Controller.

   The IP address of a target must resolve to only one known network for automatic routing to succeed.

e. (Optional) The Discovery credentials and management option are completed based on the profile. You can edit these fields.

f. Click **Add Now**.

The control domain appears in the Asset tree under its server hardware. The control domain operating system appears under its control domain. The non-global zones, if any, appears under the control domain operating system with a limited set of capabilities.

When an Agent Controller is installed on the control domain, all the networks that are connected to the control domain and configured with an IP address are discovered in Oracle Enterprise Manager Ops Center.

Also, when the Agent Controller is installed on the logical domains of the control domain, the extra networks that are connected to the logical domains and configured with an IP address are discovered in Oracle Enterprise Manager Ops Center.

To Agentlessly Manage an Existing Oracle VM Server for SPARC System

Agentlessly managed discovery locates the logical domains through the proxy management point.

1. (Optional) Create SSH credentials. If you have not created SSH credentials, you can create them before creating an OS Discovery profile, or you can create them as part of the OS Discovery profile. See Using Management Credentials to create SSH credentials.
2. Create an agentless OS Discovery profile. See Creating a Discovery Profile for more information.
   a. Expand **Plan Management** in the Navigation pane.
   b. Click **Discovery** in Profiles and Policies.
   c. Click **Create Profile** in the Actions pane.
   d. Enter a name and description for the discovery profile. Expand Operating Systems and select **Solaris, Linux OS**. Click **Next**.
   e. (Optional) Click **Next** to skip tags. See Using Tags for information about using tags in a Discovery profile.
   f. (Optional) Click **Next** to skip adding IP Ranges. You are prompted to provide IP addresses and host names during discovery.
   g. Select **Manage without Agent Controller**. Click **Replace** and choose the SSH management credentials. Click **Next**.

Figure 19–3 Discovery Credentials for Oracle VM for SPARC Management

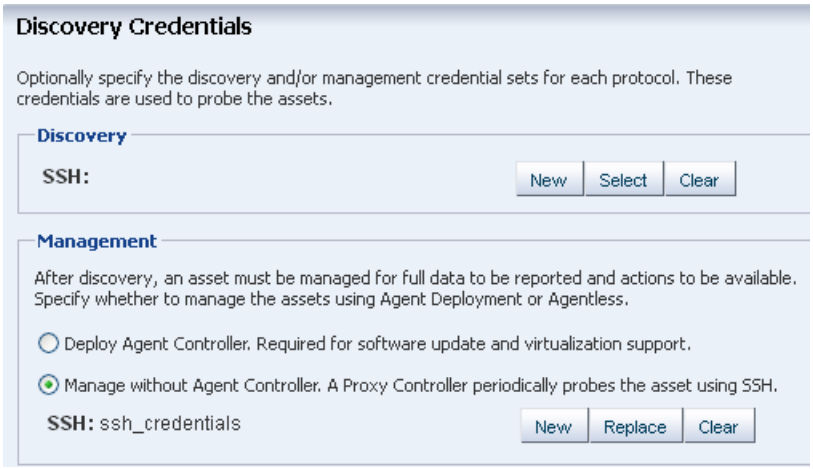

h. Review the Summary page, then click **Finish**.

3. Use the OS Discovery profile to add the Oracle VM Server for SPARC:

   a. Expand **Assets** in the Navigation pane, then click **Add Assets** in the Actions pane.

   b. Select **Add and manage various types of assets via discovery probes**, then click **Next**.

   c. Select the Discovery profile you created in the previous step. You are prompted to complete the host names or IP addresses and network.

   d. Enter a comma-separated list of host names or IP addresses. Select a managed network with which the host is associated, or select Automatic to route the job to the most appropriate Proxy Controller.

      The IP address of a target must resolve to only one known network for automatic routing to succeed.

   e. (Optional) The Discovery credentials and management option are completed based on the profile. You can edit these fields.

   f. Click **Add Now**.

The control domain appears in the Asset tree under its server hardware. The control domain operating system appears under the control domain. The logical domains appear under the control domain operating system.

Because it is agentlessly-managed, the control domain has a limited level of monitoring and management actions.

Note: To change the server to agent managed, use **Switch Management Access**. To monitor and manage the Oracle VM Server for SPARC server, control domain, and operating system, select the LDom Virtualization Controller agent. To monitor and manage the global zone, select the Zone Virtualization Controller Agent.

Provisioning Oracle VM Server for SPARC

The following information is covered in this section:

- Prerequisites for Provisioning Oracle VM Server for SPARC

- Recommended Minimum Configuration
- Profiles for Provisioning and Configuring Oracle VM Server for SPARC
- Creating an OS Provisioning Profile for Oracle VM Server for SPARC
- Creating an OS Configuration Profile for Oracle VM Server for SPARC
- Creating a Link Aggregation in Oracle VM Server for SPARC
- Creating IPMP Group in Oracle VM Server for SPARC
- Deployment Plans for Oracle VM Server for SPARC
- Overview of Oracle VM Server for SPARC Installation

You can provision the following versions of Oracle VM Server for SPARC using Oracle Enterprise Manager Ops Center:

- 1.2
- 1.3
- 2.0
- 2.1
- 2.2
- 3.0
- 3.1

Prerequisites for Provisioning Oracle VM Server for SPARC

The specific hardware, OS, and firmware requirements must be met for provisioning Oracle VM Server for SPARC.

Operating System

Configure your Oracle Solaris 11 Software Update Library in the Enterprise Controller. The library is a local copy of the parent repository of Oracle Solaris 11 Image Packaging System (IPS).

When you create OS provisioning profile, the supported Oracle VM Server for SPARC version, Oracle Solaris 11 OS and the Support Repository Update (SRU) are populated and listed.

When you want to use Oracle Solaris 10 OS, import or upload the required update versions of Oracle Solaris 10 OS in the Initial EC Library or in a software library created using a local file system on the Enterprise Controller or on a shared file system on a NFS server. Refer to the Table 19–3 for the required update versions of Oracle Solaris 10 OS.

Table 19–3 Supported Oracle VM Server for SPARC version

Oracle VM Server for SPARC Version	Oracle Solaris 10 OS
1.2, 1.3, 2.0, and 2.1	Oracle Solaris 10 9/10 OS or later release
2.2	Oracle Solaris 10 8/11 OS or later release
3.0	Oracle Solaris 10 8/11 OS or later release
3.1	Oracle Solaris 10 1/13 OS or later release

Upload the required Oracle Solaris OS image to the software library in Oracle Enterprise Manager Ops Center. See Chapter 5, "Software Libraries" for more information about uploading or importing OS images.

Refer to http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html for more information about supported Oracle Solaris OS versions for different Oracle VM Server for SPARC versions.

Hardware and Firmware

Refer to the Oracle VM Server for SPARC Release Notes in http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html for the supported platforms to install different versions of Oracle VM Server for SPARC.

Ensure that you install the correct system firmware. You can update the firmware using Oracle Enterprise Manager Ops Center. See Chapter 11, "Hardware" for downloading and updating the required firmware version.

Refer to **PCIe SR-IOV Hardware and Software Requirements** and **Non-primary Root Domain Hardware and Software Requirements** sections in the Release Notes of Oracle VM Server for SPARC documentation at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html.

When you use Oracle Enterprise Manager Ops Center to provision the Oracle VM Server for SPARC software, it installs the Oracle Solaris OS, the control domain, and an agent on the target system. The provisioning action removes existing virtualization software, including any previous logical domains installed on the service processor.

To install Oracle VM Server for SPARC software, apply a profile that specifies the values for resources such as CPU threads, Crypto units, and memory as described in Profiles for Provisioning and Configuring Oracle VM Server for SPARC.

When you have existing Oracle VM Server for SPARC environment that you want to discover and manage in Oracle Enterprise Manager, refer to the section Discovering Existing Oracle VM Server for SPARC Environments for more information.

Recommended Minimum Configuration

The recommended minimum configurations for the control domain are described in the following sections:

- CPU Resource Allocation
- Crypto Units
- RAM

CPU Resource Allocation

The number of system CPUs determines the number of control domain CPU threads:

- For less than 16 system CPUs, set the control domain CPU Threads to 2.
- For between 16 and 64 system CPUs, set the control domain CPU Threads to 4.
- For more than 64 system CPUs, set the control domain CPU Threads to 8.

Whole-Core

You can select to allocate CPU resources either as CPU Threads or Whole-core, Whole-core is the default value in Oracle Enterprise Manager Ops Center. When you allocate as Whole-core, all the CPU Threads in the core are allocated to the control

domain. For example, when you allocate two cores in UltraSPARC T2 servers, the control domain is allocated with all the 16 CPU Threads in the core. You can also set the maximum cores constraint when you select Whole-core allocation type. The maximum number of cores constraint specifies the number of cores that must be assigned to the domain.

Crypto Units

Crypto units are the resources on the supported platforms that provide high-performance, dedicated cryptographic engines. These can be used for tasks such as encrypting and decrypting network traffic between a Secure Socket Layer (SSL) web server and an application server.

Each CPU core has one Crypto unit and four or eight CPU threads. Because the Crypto unit is part of a core, the Crypto unit is bound only to domains that contain at least one thread from the parent core. Crypto units cannot be split as CPU threads are split. For example, you have assigned the Crypto unit for the first CPU core to the control domain. When a new logical domain is assigned a thread from the first CPU core and the Crypto unit for that core is already assigned, the control domain cannot assign that Crypto unit to the new logical domain. Allocation of Crypto units might not succeed, especially when a core is split between domains. An Oracle VM Server might allocate fewer Crypto units or none at all.

You must assign at least one Crypto unit to the control domain because the Crypto unit enables domain migration.

The use of Crypto Units is not mandatory although it might speed the logical domain migration. Allocation of Crypto Units might not be available in all the hardware supported. See the corresponding hardware data sheet and documentation for more information before planning for Crypto Units.

Example 19–1 Example of Crypto Unit Assignments

In UltraSPARC T1 based servers, one core is four CPU threads. Therefore, assign one Crypto unit and four CPU threads to the control domain. These values are set in the OS profile for Oracle VM Server for SPARC.

In UltraSPARC T2 and T2 Plus based servers, one core is eight CPU threads. Therefore, assign one Crypto unit and eight CPU threads to the control domain.

RAM

The amount of RAM for the control domain depends on the size of the system RAM and the load of the system.

- For system RAM less than 8 GB, set the control domain's RAM to 1 GB.
- For system RAM between 8 GB to 16 GB, set the control domain's RAM to 2 GB.
- For system RAM greater than 64 GB, set the control domain's RAM to 8 GB.

In Oracle Enterprise Manager Ops Center the default value is 4 GB as a recommended starting point for logical domains, and the minimum value is 1GB.

Note: See Oracle VM Server for SPARC Administration Guide at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html for more information about allocating memory size.

NVRAMRC Value

Automatic booting on a SPARC system uses the default boot device that is defined in the non-volatile RAM (NVRAM). User-defined commands that are executed during start-up are stored in the NVRAMRC file in the NVRAM. When you run an OS provisioning job on a SPARC machine, Oracle Enterprise Manager Ops Center resets the configuration to the factory default configuration and removes the information that is stored in the NVRAMRC file. The control domain OS configuration profile gives you the option to preserve the information in the NVRAMRC file before resetting the server to the factory defaults, and then restore the information after the reset.

Profiles for Provisioning and Configuring Oracle VM Server for SPARC

Oracle Enterprise Manager Ops Center provides two profiles for defining the parameters required for provisioning Oracle VM Server for SPARC:

- **OS Provisioning Profile**: The profile collects information required for provisioning the OS.
- **OS Configuration Profile**: The profile collects information about control domain parameters and network configuration.

Create an OS provisioning plan that encapsulates the OS provisioning and configuration profile for installing Oracle VM Server for SPARC software, the OS, and to set the parameters of the control domain.

The OS provisioning profile collects the following details:

- OS provisioning parameters such as OS image, and Oracle VM Server for SPARC version.
- OS setup details such as time zone, language, root and console serial port and baud rate, root password, and options to enable manual net boot and save NVRAMRC values.
- File system layout.
- Naming services that includes, DNS, NIS, and LDAP.

The OS configuration profile collects the following information:

- Control domain configuration such as CPU Threads or Whole-cores, memory, and Crypto Units.
- Options such as enabling SR-IOV, and detaching the unused buses.
- OS management option to install the Agent Controller.
- Enabling multiplexed I/O to connect to block storage libraries.
- Networking options to use IPMP or Link Aggregation.

Creating an OS Provisioning Profile for Oracle VM Server for SPARC

1. Select the **Plan Management** section and expand **Profiles and Policies**.
2. Select **OS Provisioning** and click **Create Profile** in the Actions pane.
3. Specify the profile details in the **Create Profile - OS Provisioning** wizard:
   a. Enter the name and description of the profile.
   b. Select Oracle VM Server for SPARC as the Subtype.

Click **Next** to select the OS image and distribution.

4. Select the following parameters in the Specify OSP Parameters step:
   - Oracle VM Server for SPARC Version
   - Software Group

   Refer to Table 19–3 for selecting Oracle Solaris 10 OS. You have option to include custom scripts when you use Oracle Solaris 10. For Oracle Solaris 11 OS, the SRU and the supported Oracle VM Server for SPARC version are listed.

   Click **Next** to specify the OS setup.

5. Specify the following OS setup parameters:
   - Select a language from the list.
   - Specify the time zone.
   - Specify a terminal type.
   - To monitor the installation using a serial connection, select the console serial port device and the baud rate.
   - Enter the NFS4 domain name for the target system to use. A dynamic NFSv4 domain name enables the name to be derived at run time, based on the naming service configuration. You can also enter a static domain name.
   - Enter the root password for the root user on systems provisioned using this profile. Re-enter the password for confirmation.
   - Select the **Manual Net Boot** option to enable manual control of network boot operations for the target system. You must select this option for a target system that does not have a service processor because Oracle Enterprise Manager Ops Center cannot control the network boot process remotely on these systems.
   - Select the option **Save NVRAMRC** values to preserve your user-defined commands that were executed during start-up.

   Click **Next**.

6. If it is Oracle Solaris 11 OS, then you are directed to Step 7 to create a user account to log in.

   Otherwise, to Step 8 for Oracle Solaris 10 OS.

7. For Oracle Solaris 11 OS, root login to the system is not enabled. Create a user account through which you can SSH to the OS after provisioning. Enter the user name and password for the user account.

   Click **Next** to specify whether you want to install the OS on an iSCSI disk.

8. Select the option **Use iSCSI Disk** to use iSCSI disks for provisioning Oracle VM Server for SPARC. The storage server is identified in the network and select the volume group that must be used to create a new iSCSI disk. A new iSCSI disk is created in the storage server for each target when necessary.

   A Dynamic Block Storage Library is automatically created for each storage server.

   If you want to enter the details of the storage server, select **Manually Specify iSCSI Disk** to enter the storage server IP address and the LU Number while deploying the plan.

   Click **Next** to specify the file system layout.

9. Specify the disk partitions and file systems that you want to create on the target system. The root (/) and a swap file system are defined by default. You can modify the size of the file system. Click the **Add** icon to define a new partition. For each partition that you define, provide the following information:
   - **File System Type**: Select a file system type: ufs, unnamed, zfs, or swap.
   - **Mount Point**: Enter a directory to use as a mount point for partitions.
   - **Device**: Enter the `rootdisk` keyword and a slice value to describe a partition on the target system's boot disk, for example, `rootdisk.s0`, or enter the logical device name, for example, `c1t0d0s0`, of the partition that you want to create.
   - **Size**: Enter the size that you want to assign to the partition, expressed in MB. Do not enter any value for the size when you want to allocate the remaining unused disk space to a file system.

   Click **Next** to specify the name service.
10. Specify the name service, domain name and the corresponding name server. Select one of the following name services:
    - **DNS**: Enter the domain name of the DNS server and enter the IP address of the DNS server in the **Name Server** field. You can enter up to three IP addresses as the value for the Name Server. Provide the additional domains to search for name service information in the Domain Name Search List. You can specify up to six domain names to search. The maximum length of each search entry is 250 characters.
    - **NIS or NIS+**: Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, select the option **Specify an NIS Server** and enter the NIS server host name and the IP address.
    - **LDAP**: Enter the domain name of the LDAP server. Specify the name of the LDAP Profile you want to use to configure the system. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.
    - **None**: Select **None** when there is no naming service configured.

    Click **Next** to review the properties of the profile.
11. Review the summary of your selections. Click **Finish** to create the profile.

Creating an OS Configuration Profile for Oracle VM Server for SPARC

1. Select the **Plan Management** section and expand **Profiles and Policies**.
2. Select **OS Configuration** and click **Create Profile** in the Actions pane.
3. Specify the profile details:
   a. Enter the name and description of the profile.
   b. Select Oracle VM Server for SPARC as the Subtype.

   Click **Next** to specify the setup for configuring control domain.
4. Specify the resources that you want to assign to the control domain. The remaining resources are available for the logical domains.
   - **Oracle VM Server Version:** Select the Oracle VM Server for SPARC version to be configured.

- **CPU Model**: Select how you want to allocate the CPU resources, Virtual CPU or Whole-Core. The default value is Whole-Core.
- **CPU Threads**: This field is displayed when you select the CPU Model as Virtual CPU. Specify the number of CPU threads that you want to assign to the control domain.
- **CPU Cores**: This field is displayed when you select the CPU Model as Whole-Core. Specify the number of cores to be allocated to the control domain.
- **Max CPU Cores**: This field is displayed when you select the CPU Model as Whole-core. Specify the maximum number of cores that must be assigned to the domain.
- **Memory**: Specify the amount of memory that you want to assign to the control domain. The default value is 4 GB and the minimum value is 1 GB.
- **Requested Crypto Units**: Specify the number of Crypto units that you want to assign to the control domain.
- **Virtual Console Port Range**: Specify the minimum port and maximum port of the virtual console of the control domain. The default port range for virtual console is 5000 to 6000.
- **Virtual Disk Server**: A virtual disk server (vds), `primary-vds0`, is added to the control domain by default. The virtual disk services allow you to export virtual disks to logical domains. You can modify the default name of the virtual disk server.
- **Enable Single Root I/O Virtualization (SR-IOV)**: The supported PCIe buses are SR-IOV enabled and maximum number of virtual functions (VFs) are created on them. The type of network card decides the number of VFs. The option to enable SR-IOV is available from Oracle VM Server for SPARC 2.2 version.
- **Detach unused buses**: The PCIe buses that are not in use by the control domain are released. All the unused PCIe buses are detached so that it is available for creating root domains.

Note: For Oracle VM Server for SPARC 1.2 version, you can select **Enable JASS** to install SUNWJass package to harden the system.

Click **Next** to specify the OS management for installing Agent Controller.

5. The option to install the Agent Controller is by default and cannot be changed.

   Select **Enable Multiplexed I/O** option to enable multiple path for SAN library connectivity to the control domain. The action enables LUNs to be accessed through multiple host controller interfaces from a single instance of the I/O device.

   Click **Next** to specify the networking details for provisioning.

6. Select one of the network options for the target system:
   - When you select **Use Link Aggregation**, refer to the section Creating a Link Aggregation in Oracle VM Server for SPARC to define the parameters for Link Aggregation.
   - When you select **Use IPMP**, refer to the section Creating IPMP Group in Oracle VM Server for SPARC for IP Multipathing.

- When you select **None**, you are forwarded to the next step to define the network interfaces and IP address allocation for the selected networks.

Select an option and click **Next** to continue.

7. Select the following details for network configuration:
   - **Controller**: The network interfaces built in with the server hardware are displayed as **default** in the **Controller** column. For bare-metal OS provisioning, only the in-built network interfaces of the server are displayed.
   - **Interface**: Select the network interface which you want to configure on the OS. The network interfaces are represented as *net_x*. x ranges from 0 to 31. The number of interfaces for each controller available is 32. For bare-metal ILOM servers, the device map are not available. For such instances, you can specify the MAC address instead of interface while deploying the provisioning plan.
   - **Address Allocation Method**: The address allocation for the network is to use static IP. You have to provide the static IP address while deploying the OS provisioning plan.
   - **Virtual Switch:** Select **Auto** to create virtual switch automatically for the network connection. Or select **None** if you do not want to create virtual switch. You can also enter a name for the virtual switch. Enter a switch name directly in the virtual switch list.

     Click the **Add** icon to add more networks and define the interfaces.

   Click **Next** to go to the Summary step.
8. Review the Summary of your selections. Click **Finish** to create the profile.

Creating a Link Aggregation in Oracle VM Server for SPARC

While creating the profile for installing Oracle VM Server for SPARC, follow the procedure to create link aggregation:

1. Select **Use Link Aggregation** in the Specify Networks step of the profile.
2. Specify and configure the IEE 802.3ad Link Aggregation details:
   - **Link Aggregation Name**: Select the name of the Link Aggregation. The names are set to *aggr<x>.*
   - **Load Balancing Policy**: Define the policy for outgoing traffic.
   - **Aggregation Mode and Switches**: When the aggregation topology connects through a switch, determine whether the switch supports the Link Aggregation Control Protocol (LACP). When the switch supports LACP, you must configure LACP for the switch and the aggregation. Define one of the modes in which LACP must operate.
   - **LACP Timer**: Indicates the LACP timer value, either short or long.
   - **MAC Address Policy**: Define whether the MAC address of the NICs are fixed.
   - **Virtual Switch:** Select **Auto** to create virtual switch automatically for the network connection. Or select **None** if you do not want to create virtual switch. You can also enter a name for the virtual switch. Enter a switch name directly in the virtual switch list.
   - **Number of Interfaces**: By default, two interfaces are defined in the aggregation. You can modify the number of interfaces for the aggregation.

You can add multiple link aggregation and define the configuration parameters for each aggregation.

Click **Next** to specify the link aggregation interfaces.

3. Select the interfaces for each link aggregation defined in the previous step. Choose the **Controller** from which the interfaces are used in the link aggregation.

   You must define the number of interfaces in the previous step. You cannot add or delete the required interfaces in this step.

   Click **Next** to continue.

Creating IPMP Group in Oracle VM Server for SPARC

While creating the profile for installing Oracle VM Server for SPARC, follow the procedure to create IPMP groups.

1. Select **Use IPMP** in the Specify Networks step of the profile.
2. Specify and configure the IPMP groups:
   - **IPMP Group Name**: The name of the IPMP group is automatically displayed in the format of *ipmp<x>*. You can also modify the default name of the IPMP group.
   - **Failure Detection**: The detection can be Link-Based or Link-Based and Probe-Based.
   - **Number of Interfaces**: By default, two interfaces are defined in the group. You can modify the number of interfaces for the IPMP group.

   You can add multiple IPMP groups and select the failure detection for each group.

   Click **Next** to specify the IPMP interfaces.
3. Select the **Controller** from which the interfaces are used in the IPMP group. Select the interfaces for each IPMP group and choose whether they are **Failover** or **Standby** interface. You can add multiple interfaces to each IPMP group.

   Select whether to assign IP address to the specified NIC during configuration. The data and test addresses are assigned during profile execution.

   Select **Auto** to create virtual switch automatically for the network connection. Or select **None** if you do not want to create virtual switch. You can also enter a name for the virtual switch. Enter a switch name directly in the virtual switch list.

   Click **Next** to continue the profile.

Deployment Plans for Oracle VM Server for SPARC

The Provision OS deployment plan provides the steps to provision and configure Oracle VM Server for SPARC. The Provision OS plan consists of the following profiles:

- OS Provisioning Profile
- OS Configuration Profile

Create a Provision OS deployment plan with the two profiles defined in the steps of the plan and apply the plan on the target server.

See Chapter 8, "Plans and Profiles" for more information about creating and managing the profiles and plans.

Applying a Deployment Plan for Oracle VM Server for SPARC

When you apply a deployment plan to provision Oracle VM Server for SPARC, you must have the following information to complete the installation:

- In the server hardware, obtain the network interface that is physically connected to the network that is managed by Oracle Enterprise Manager Ops Center.
- Both tagged and untagged networks are listed for network configuration. When OpenBoot PROM (OBP) is used, only untagged networks can be used for OS provisioning as you cannot boot from a tagged network.
- The IP address for the boot interface.
- If you want to identify the network interface using the MAC address, you can select to enter the MAC address instead of the boot interface.
- The details of network connection that you want to use to configure the OS after booting. For multiple network configuration, the first interface is always overwritten by the boot interface that is defined during the OS provisioning step. For single network configuration, it is same as the boot interface network.
- Select to place the Oracle VM Server for SPARC in a server pool or not.

To Apply a Deployment Plan for Oracle VM Server for SPARC

1. Select the **Provision OS** plan in the **Deployment Plans** list.
2. Select **Apply Deployment Plan** in the Actions pane.
3. Select one or more assets and add to the target list.
4. Select the plan to be applied with minimal interaction. If required to change the profile parameters, then select **Allow me to override any profile values**.
5. Select not to review the steps that are not included in the plan and click **Next**.
6. The wizard collects information for provisioning Oracle VM Server for SPARC. Click **Next**.
7. In the Boot Interface Resource Assignments step, provide the following information:
   - **Network**: The network for the boot interface.
   - **Controller**: Select the controller that provides the network interface for OS provisioning. It is always default for the Oracle VM Server for SPARC provisioning.
   - **Interface**: Select the interface from the list. The network interface that is physically connected to the selected network.
   - **IP Address**: Enter the IP address for the boot interface.
   - (Optional) **Primary Hostname**: Enter the host name for Oracle VM Server for SPARC.

   If you want to identify the network interface by its MAC address, then select the option **Identify Network Interface by MAC Address** and enter the MAC address instead of selecting the Controller and the Interface.

   Click **Next** to view the OS provisioning summary.
8. Review the parameter of OS provisioning and click **Next**.
9. The following steps in the wizard collects information about OS configuration.

   Click **Next**.

10. Specify the network resources that were defined in the profile. Select the network and for each network, select the network interface and enter the IP address.

    The first network interface listed is the boot interface. For multiple network configuration, the first network interface is always overwritten by the boot interface network. You can select which is the primary network interface after the provisioning of the OS.

    For single network, the boot interface network will be the defined as the primary network during OS configuration.

    Click **Next**.

11. Select whether you want to add the Oracle VM Server for SPARC to a server pool. You also have option to create a new server pool with default server pool settings. For creating a new server pool, you must provide a NAS storage library for the server pool.

    Click **Next**.

12. Review the summary of the OS configuration parameters and click **Next** to schedule the job.

13. Schedule the provisioning job to run immediately, or at a later time.

    Click **Apply** to apply the deployment plan on the selected targets.

Overview of Oracle VM Server for SPARC Installation

You can install the software on bare-metal or on systems already configured with logical domains. When you provision Oracle VM Server through Oracle Enterprise Manager Ops Center, any previous configuration is removed and the service processor is reset to its factory defaults.

When you provision Oracle VM Server for SPARC, the instance of the Oracle Solaris OS becomes the control domain which is the first domain to be created. The virtualization host on which this instance of Oracle Solaris OS runs is called the Oracle VM Server Host or simply, Oracle VM Server. The Logical Domains Manager runs in the control domain and provides the functions to create and manage logical domains. You can have only one control domain per server. You cannot change the name or destroy the control domain.

This section describes how to install the Oracle VM Server for SPARC software on a service processor:

1. The Oracle VM Server for SPARC software depends on particular Oracle Solaris OS versions, required software patches, and particular versions of system firmware. See Selection of Oracle VM Server Version.

2. Discover the target system on which you want to install Oracle VM Server for SPARC software.

3. Refer to Table 13–2 for information about configuring DHCP services on Proxy Controllers, if required.

4. Create an OS provisioning and OS Configuration profiles for provisioning Oracle VM Server for SPARC. See Profiles for Provisioning and Configuring Oracle VM Server for SPARC.

5. Create a deployment plan with the profiles created for provisioning and configuring Oracle VM Server for SPARC.

6. (Optional) Adjust the amount of time allowed for the provisioning job. The default time is three hours or 180 minutes. You can edit the `default_timeout` value for OS provisioning in the Enterprise Controller Configuration in the Administration section, as shown in Figure 19–4.

Figure 19–4 Default Timeout Settings

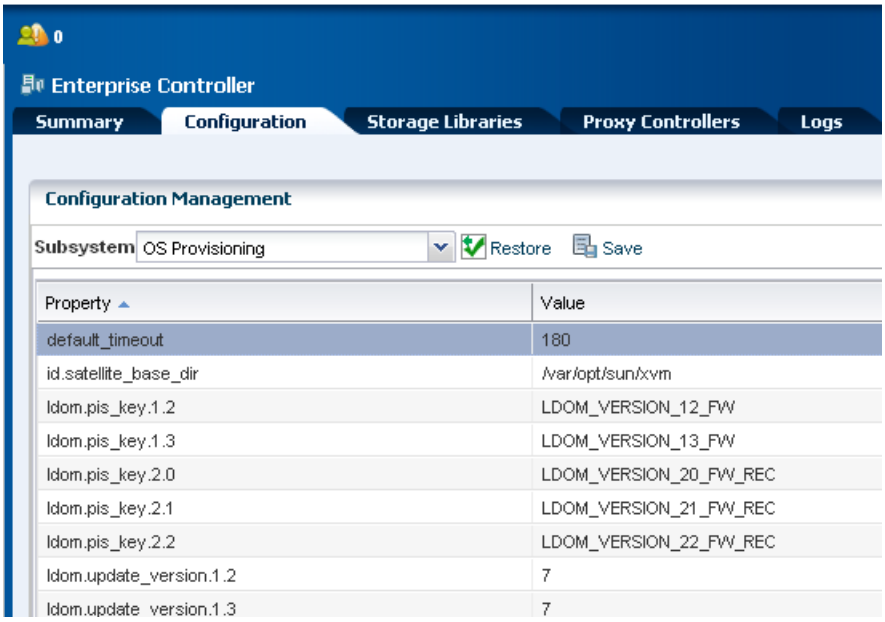

The Oracle VM Server for SPARC provisioning job performs the following major tasks:

- Downloads the appropriate OS image.
- Initiates a net boot action on the service processor.
- Installs the Oracle VM Server Host.
- Configures Oracle VM Server Host according to the values set in the profile such as memory, CPU threads, Crypto units, and virtual console port range.
- When selected, installs the SUNWJass package to harden the system. For example, you can install the SUNWJass package for version 1.2.
- Enables the Fibre Channel ports on the storage system.
- Installs and configures the Agent Controller. Do not install the Agent Controller manually on the Oracle VM Server host.
- Detaches unused buses and enables SR-IOV feature on PCIe buses as defined in the profile.

Selection of Oracle VM Server Version

Verify that the target system has the correct hardware and firmware configuration to support the installation of selected version of Oracle VM Server for SPARC.

When the supported hardware and firmware versions are not there, the provisioning job might fail. You cannot force set the version of the software to be installed. For the complete list of supported hardware, firmware version, and the supported Oracle VM Server for SPARC version, refer to the Release Notes in the library `http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html` for more information.

Important Notes for Installation

Some important notes for installing Oracle VM Server for SPARC in Oracle Enterprise Manager Ops Center:

- The OS in the control domain must have the default locale set to C. The control domain provisioning is supported on Oracle Solaris SPARC or x86 Proxy Controllers, but not on Linux Proxy Controllers.
- To provision the Oracle VM Server for SPARC software, the proper system model name must be populated on the service processor. The model name is not available on the service processor of the Sun Blade T6320 hardware.
- Oracle Enterprise Manager Ops Center does not support use of the LDoms configuration manager in LDoms 1.2.

Manual Net Boot Initiation

A target system requires a manual net boot operation when Oracle Enterprise Manager Ops Center cannot perform a remote network boot process. Profiles that provision this type of system contain the Manual Net Boot option enabled. You must initiate the net boot operation when the Oracle VM Server provisioning job completes on the UI.

1. Log in to the service processor of the target system.
2. When the target system is running, enter `halt`.
3. Enter the following command

```
boot net:dhcp - install
```

Additional Configurations

A virtual console concentrator named *primary-vcc0* is created during the installation of Oracle VM Server for SPARC and is not required to be included or defined in the profiles.

The virtual disk server is created by the default name *primary-vds0*. You can modify the name of the virtual disk server in the OS configuration profile.

Manage Oracle VM Server for SPARC

The following actions and information are available for managing the Oracle VM Server:

- Virtual Services in Control Domain
- View I/O Resources
- Modify Configuration and Tags
- Reboot Oracle VM Server
- Performance Management

After you have provisioned or discovered the Oracle VM Server for SPARC, you can dynamically manage the Oracle VM Server control domain resource configurations such as CPU Threads, Crypto Units, and memory. You can associate storage libraries and attach networks with Oracle VM Server for SPARC, create logical domains, and monitor the performance of logical domains.

You can also create server pools for Oracle VM Server for SPARC for an efficient management of your resources. For creating server pools, see Chapter 21, "Server Pools".

Virtual Services in Control Domain

In Oracle Enterprise Manager, the control domain is also a service domain and the following virtual device services are created:

- **Virtual disk server**: A virtual disk server (vds), `primary-vds0`, is added to the control domain by default. The virtual disk services allow you to export virtual disks to logical domains. The virtual storage infrastructure enables logical domains to access the storage disks that are not directly assigned to them. The virtual disk service processes the requests from the logical domains and submits them to the back end storage.
- **Virtual network switches**: The virtual network switch enables networking between virtual network devices in logical domains. When you connect a network to control domain, a virtual switch is created for each network connection. You can also define the switch name. You use virtual switch to connect a logical domain to the external network for communication. A virtual network (vnet) device is defined in the logical domain when connected to a virtual switch. There is no virtual switch creation for SR-IOV enabled networks and you do not require virtual switch for connecting to SR-IOV enabled network. You can also edit the virtual switch attributes such as Oracle Solaris Cluster Mode, Physical Link State, and Inter-vnet Link.
- **Virtual console concentrator**: This functions as a concentrator for all logical domains console and for use by the virtual network terminal server daemon. The console I/O from all other domains is redirected to the control domain that is running the virtual console concentrator (vcc).

View I/O Resources

The I/O Resources tab in the Center pane displays the PCIe root complexes, PCIe Endpoint devices, Network Interface Units (NIU), and Single Root I/O virtualization (SR-IOV) functions.

Buses and End Point Devices

The PCIe root complexes and the NIUs that are available in the servers are listed here. The number of PCIe root complexes and NIUs depend on the type of server hardware. Initially, when you provision Oracle VM Server for SPARC, all the PCIe root complexes and NIUs are assigned to the control domain. When you use the option **Detach Unused Buses** option in the OS configuration profile, the PCIe buses are released from the control domain.

You allocate PCIe buses to create root domains. You allocate PCIe Endpoint devices to create Physical I/O domains.

To assign PCIe Endpoint devices to Physical I/O domains, you must first release the devices from the corresponding PCIe bus that is allocated to the control domain or root domain. You must manually release the PCIe Endpoint devices using the CLI.

The PCIe Endpoint devices displays whether the PCIe card slot is occupied or empty. If occupied, the domain to which it is allocated is displayed. The devices to which the PCIe bus is attached to are also grouped and displayed.

SR-IOV Services

Single Root I/O Virtualization is a PCI-SIG standard specification that enables efficient use of PCIe devices. In this, a single PCIe card owned by a PCIe root complex is made to physically appear in multiple domains simultaneously. A single I/O resource also know as physical function is shared by many virtual machines. A physical function is

a PCIe device that is SR-IOV enabled with appropriate hardware and OS support can appear as multiple, separate physical devices, each with its own configuration space.

A virtual function is a lightweight PCIe function that shares one or more physical resources with the physical function and with virtual functions that are associated with that physical function. The number of virtual functions that are supported on a physical function depends on the hardware.

The physical functions and their corresponding virtual functions are listed in the SR-IOV services tab. The number of virtual functions supported for the physical function and the logical domains to which the virtual functions are assigned are displayed.

Currently, the SR-IOV feature is enabled only for the SR-IOV cards that are installed on the control domain. SR-IOV feature are not enabled when you assign the SR-IOV card to a physical I/O domain.

With the release of Oracle VM Server for SPARC 3.1 version, the SR-IOV feature is enabled when the SR-IOV cards are assigned to a root domain.

Modify Configuration and Tags

You can modify the control domain configuration that includes CPU Model, CPU Threads or cores, and memory. You can select to switch between the CPU Models. When you select Whole-core option, you can edit the number of whole-cores and the maximum number of cores to be allocated to the Oracle VM Server.

You can also modify the name and description of the control domain.

Use the **Edit Attributes** option to modify the configuration. The Oracle VM Server must have Oracle Solaris 10 10/09 OS or higher version to edit the attributes.

When you modify the memory size, Oracle VM Server is rebooted unless it is running Oracle VM Server for SPARC 2.0 or higher version. You can choose to cancel the reboot option after editing the changes.

If you have any physical bindings constraint for the memory, CPU cores, or both, then you cannot edit CPU and memory configuration of the control domain in the UI. You can edit the CPU Model even if you have any physical bindings constraint, but the change of CPU Model will undo the physical binding as well.

Oracle Enterprise Manager Ops Center also does not provide options to explicitly assign physical resources such as CPU and memory. Whereas, you can use the CLI to assign physical resources to the logical domains. These are displayed in the Summary tab of the control domain as Physical Bindings.

Use **Edit Tags** option to modify the existing tags and add new tags.

Delayed Reconfiguration Mode

Some resources of a domain can be configured dynamically on a running logical domain, while others must be configured on a stopped domain. If a resource cannot be dynamically configured, the logical domain can be set in delayed reconfiguration state to postpone the configuration activities until after rebooting the domain. Delayed reconfiguration is restricted to the control domain or a PCIe root domain (if supported). For all other logical domains, you must stop the domain to modify the configuration unless the resource can be dynamically reconfigured.

You cannot initiate a delayed reconfiguration using Oracle Enterprise Manager Ops Center UI, but Oracle Enterprise Manager Ops Center detects when a logical domain is in delayed reconfiguration state. When a domain is in delayed reconfiguration state, a

message in the center pane is visible to all logical domains attached to the affected control domain. When a delayed reconfiguration is in progress, trying to perform actions such as start, shutdown, reboot, or change attributes in any domain attached to the affected control domain will fail.

You can either reboot the domain in delayed reconfiguration to apply changes and exit the delayed configuration mode, or use the **Cancel Delayed Reconfiguration** action to cancel the delayed reconfiguration in the logical domain.

Reboot Oracle VM Server

You can reboot the control domain regardless of the state of the logical domains in it. During the reboot, the logical domains might suspend temporarily but keep their current states.

Performance Management

You can view the following information about Oracle VM Server performance:

- The **Summary** tab displays the Oracle VM Server for SPARC configuration information, and its health status.
- The **Dashboard** tab displays the general information about the selected Oracle VM Server for SPARC. A graphical representation of the group members and its relationship are displayed. The unassigned incidents and the list of recent incidents are also displayed.
- The **Analytics** tab displays the CPU, memory and network utilization by the logical domains in the Oracle VM Server. The consumption is represented in the form of charts.
- The **Incidents** tab lists the incidents that are reported for the Oracle VM Server. The tab provides the list of unresolved incidents and a graphical representation of the problem composition.

  To see the corresponding alerts for each incident, see Chapter 9, "Incidents" for more detailed information.
- The **Monitoring** tab, as shown in Figure 19–5, shows the monitoring values and boundaries that are set for Oracle VM Server activity. You can review the rules, monitoring attributes, specific time period for monitoring, and the level of severity for monitoring the activity of Oracle VM Server. To change these values, see Chapter 4, "Monitoring Rules and Policies".

Figure 19–5 Oracle VM Server for SPARC Monitoring Options

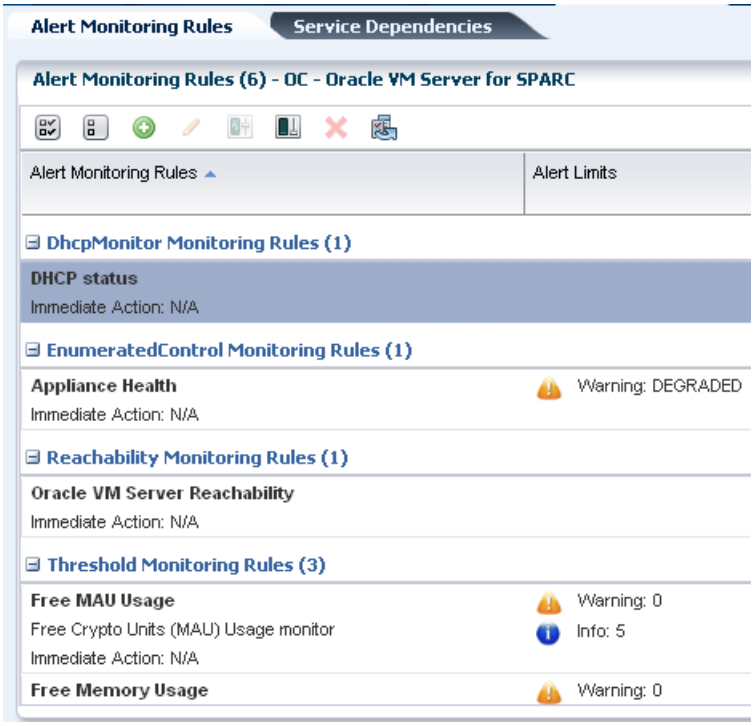

- The **Configuration** tab displays information such as remote logging, routing, NFS domain name, and name service information for the Oracle VM Server.
- The **Charts** tab displays the CPU, memory and network utilization of the Oracle VM Server. The data are collected in five-minute intervals and displayed graphically. The utilization data are provided for different time intervals. See Chapter 12, "Operating System Management" for more information about reading charts.

Managing Storage Resources in Oracle VM Server for SPARC

The following information is covered in this section:

- Associating Storage Library with the Domains
- Disassociating Storage Libraries

Storage libraries are required for storing logical domain metadata, ISO images and to provide virtual disks for logical domain storage. You must first associate the storage library to the Oracle VM Server and then it is made available for all the logical domains under it.

You can associate the following libraries for Oracle VM Server:

- **Filesystem storage**: This includes the NAS storage libraries.
- **Block storage**: This includes the Static and Dynamic Block Storage libraries.
  - **Static Block Storage**: The Static Storage library includes the LUNs of Fibre Channel (FC) and iSCSI disks of storage servers. Oracle Enterprise Manager Ops Center manages the storage device on which the LUNs are already created. The existing LUNs are assigned to the library. You cannot create, modify or delete the LUNs in Static Storage Library.

- **Dynamic Block Storage**: The Dynamic Storage Library is automatically created when the storage server is discovered and managed in Oracle Enterprise Manager Ops Center. It includes the LUNs of iSCSI disks of storage servers. You can create LUNs as you attach the library to an Oracle VM Server. You can create, modify and delete the LUNs as required and it is defined as Dynamic Storage Library.

When you are associating the libraries with the control domain, I/O domain and root domains, the storage resources are applicable at the domain level. You can associate the libraries with the control domain, I/O domain, or root domain. You must select at least one domain per server to be associated with the library. You must not attempt to associate the storage libraries from the OS of the domain.

The storage resource that is to be used for a guest domain might not be exclusively provided by the control domain but also by the I/O domains and the root domains configured on the Oracle VM Server. The same storage resource can be accessed through multiple I/O domains which results in redundancy for the access to the storage resource.

When you associate the storage library with Oracle VM Server that are configured with I/O domains and root domains, you can select the domain to which you want to associate the storage library.

Select the option **Associate Library** to associate the libraries with the selected Oracle VM Server.

Associating Storage Library with the Domains

An Oracle VM Server can be configured with I/O domains, root domains, and guest domains. When you associate storage library with Oracle VM Server, you can select to associate the library with the preferred domains.

Associate Libraries wizard provides with options to select the appropriate domains to which you want to associate the storage library.

To Associate Libraries with Oracle VM Server for SPARC

1. Select the Oracle VM Server or the control domain in the Assets section.
2. Click **Associate Libraries** in the Actions pane.

   The Associate Library Wizard is displayed.
3. Select the libraries that you want to associate with the domains.

   Click **Next**.
4. Select the **Associate** action for the control domain, I/O domain, or both. You must select at least one domain per server.

   When the storage library is already associated with any of the domains, the **Associate** action is selected. You can deselect the option to disassociate the storage library from the domain.

   This action is repeated for each selected storage library. Click **Next**.
5. Review the summary of the association and click **Finish** to associate the library with the selected domains.

The domains for which the association was removed are disassociated from the storage library.

Disassociating Storage Libraries

1. Select the Oracle VM Server or the control domain in the Assets section.
2. Select the **Libraries** tab in the Center pane.

   All the storage libraries that are associated with the system are displayed.
3. Select the storage library that you want to disassociate from the domain, then click the Disassociate Library icon.

   The Disassociate Library window is displayed.
4. The domains that are associated with the storage library are displayed with the **Disassociate** option selected. You can deselect the **Disassociate** action if you do not want to remove the association.
5. Click **Finish** to disassociate the storage library from the selected domains.

Managing Network Resources in Oracle VM Server for SPARC

The following information is covered in this section:

Attach networks to Oracle VM Server to provide networking facilities for the logical domains.

With the support for advanced I/O domain configurations in Oracle Enterprise Manager Ops Center UI, the network configuration for Oracle VM Server is also enhanced to support the access to the available network interfaces from the I/O domains and root domains.

You can attach networks to the Oracle VM Server that involves the physical interfaces or `etherstub` device belonging to the control domain, root domain or the I/O domains. When you attach networks to the Oracle VM Server, you can select the domain and the physical network interfaces or the `etherstub` devices that belong to the domain.

`etherstub` is an Oracle Solaris 11 network virtualization feature. `etherstub` is a pseudo network device which provides functionality similar to physical network devices but only for private communications with its clients. This pseudo device can be used as a network back-end device for a virtual switch that provides the private communications between virtual networks.

When you discover and manage existing Oracle VM Server for SPARC environments, you must have the Agent Controller installed on the control domain to discover all the networks that are attached and configured with an IP address in the control domain.

Also, when the Agent Controller is installed on the logical domains of the control domain, the extra networks that are connected to the logical domains and configured with an IP address are discovered in Oracle Enterprise Manager Ops Center.

Virtual Switches

Virtual switches are created or re-used when you attach networks to Oracle VM Server. This is not applicable when you select SR-IOV enabled network interface. See SR-IOV Enabled Networks for more information about attaching SR-IOV enabled networks.

You can make multiple connections to a network in the Oracle VM Server. For each network connection, a virtual switch is required. If there is an already existing virtual switch for the physical interface or `etherstub` device, you can re-use the virtual switch. If there is no virtual switch for the physical interface or `etherstub` device, you can either provide a user-friendly name for the virtual switch or a virtual switch is automatically created with a default naming pattern. For example, the network 1.1.1.0/24, the virtual switches take the name as 1.1.1.0_24, 1.1.1.0_24_1, and 1.1.1.0_24_2. The virtual switches have a unique name. When you create and start a logical domain, you define the virtual switch that connects the logical domain to the network. Each virtual switch is connected to a NIC.

Service Domains

When you connect to the network interfaces from physical I/O domains and root domains, the virtual switch is created in the control domain. You cannot define the IP address allocation for the network connection. Instead, you can later define the IP address in the OS of the logical domain as required. You can define the IP address only when the network interfaces are used from the control domain.

Tagged and Untagged Mode

You can attach the network to Oracle VM Server for SPARC in tagged or untagged mode. This option is available for networks configured with VLAN IDs. You can set the tagging mode while attaching the network to Oracle VM Server. When you make multiple connections to the same network, the connection must be either in tagged or untagged mode. You cannot mix the tagging mode for the multiple connection networks.

Refer to the section Mixed Network Tagging Mode Configurations in Server Pool for the limitations of having Oracle VM Servers in different tagging mode in a server pool.

IPMP and Link Aggregation

You can create IPMP groups and aggregate links in the control domain. Navigate to the Oracle Solaris OS of the control domain and create IPMP groups or Link Aggregation. Refer to Chapter 17, "Networks for Virtualization" for more information about creating IPMP groups and Link Aggregations in an Oracle Solaris OS.

Network Tagging Mode Conditions

There are certain scenarios in which the network configuration must be applied to avoid any networking issues. The scenarios that are explained here are for Oracle VM Server for SPARC in stand-alone mode:

- You can select networks without VLAN ID. The UI does not provide the option to select Tagged or Untagged mode.
- You can select to associate and configure the networks with VLAN ID in Tagged mode.
- You can select to associate and configure the networks with VLAN ID in Untagged mode.
- You can select to configure the networks in mixed tagging mode in the server pool. For example, you can attach the network N1 with VLAN ID = 100 in tagged mode

with the server S1 and in untagged mode for server S2. Refer to Mixed Network Tagging Mode Configurations in Server Pool for more detailed information.

- You can attach networks whose VLAN ID is similar to another network already connected to the servers. For example, a server S1 is already connected to network N1 with VLAN ID = 100, then while creating the server pool with S1 as the member of the pool, you can also attach another network N2 with VLAN ID =100.
- You can edit the VLAN ID of a network when you are attaching the network in Tagged mode for the first time.
- When you can edit the VLAN ID of the network, you cannot enter -1 as the value for the VLAN ID.
- If the selected network with a VLAN ID is already connected to the selected assets in Tagged mode, then you cannot edit the VLAN ID and make another connection.
- You cannot make multiple network connections to the Oracle VM Server over the same network in both tagged and untagged modes. The mode can be either in tagged or untagged mode only. For example, if you attach network N1 with VLAN ID =100 for the first time to server S1 in Tagged mode, then you cannot make another connection to the same network N1 in Untagged mode. Every other connection with network N1 must always be in Tagged mode for server S1.
- If the selected members of the pool are already connected to network N1 with VLAN ID =100, then you cannot select the same network with different VLAN ID to be connected for the server pool.

Attaching Networks to Oracle VM Server for SPARC

The procedure to attach networks is applicable for Oracle VM Server for SPARC in stand-alone mode. The procedure varies if the Oracle VM Server is placed in a server pool. Refer to Chapter 21, "Server Pools" for more information.

To Attach Networks to Oracle VM Server

1. Select the Oracle VM Server in the Assets section.
2. Click **Attach Networks** in the Actions pane.

   The list of available networks in Oracle Enterprise Manager Ops Center are displayed. The list also displays the existing number of connections with the server.
3. Select one or more networks from the list. You can make multiple connections to a network.

   Click **Next**.
4. Specify the Mode for networks configured with VLAN ID and the number of connections for each selected networks. Increase the total number of connections.

   If required, the VLAN ID of the network can be modified based on the condition that the network is being attached to the server for the first time in tagged mode. The VLAN ID can be already in use by the network attached to the server.

   Note: The number of connections does not limit the number of logical domains that can be connected to this network.

   Click **Next**.

5. Select the domain in the Service Domain list from which the network interfaces are selected. The service domain can be the control domain, I/O domains, or root domains. Root domains and I/O domains OS must be managed by Oracle Enterprise Manager Ops Center.

   For control domain, it is listed as `primary` in the Service Domain list. Whereas for I/O domains and root domains, the names of the domains are listed.

   Note: When you select an network interface from an I/O domain or root domain, the virtual switch is created in the control domain.

6. Specify the NIC and IP address for each network connection.

   For each network connection, you must provide a NIC. The virtual switch requires a physical interface to allow communication between the logical domains and external network. Follow these rules while assigning the NIC and IP address:

   - Specify the network interface or NIC. Specify the same NIC to different network when the network has different VLAN ID and every connection is in tagged mode. Otherwise, you cannot assign the same NIC to different networks.
   - If required, you can modify the network tagging mode specified in the previous step.
   - If there is a virtual switch available for the network, the virtual switch is displayed. Otherwise, enter a name for the virtual switch or leave it blank for the software to automatically create a name for the virtual switch using the default naming pattern.
   - Select an IP address allocation method according to your requirements:
     - When supported by the network, select **Assign by DHCP** to enable the system to automatically allocate the IP address.
     - Select **Use Static IP** to provide an IP address for the network connection.
     - Select **Do Not Allocate IP** for the IP address when you do not want to assign IP address to the selected network interface.
     - Select **Do not Plumb Interface** for the IP address when you do not want to assign IP address and plumb the selected network interface. This option is available for non SR-IOV enabled networks.

     Note: For network interfaces from I/O domains or root domains, the Address Allocation Method is automatically set to Do Not Allocate IP.

   Click **Next**.

7. Review the summary and click **Finish** to attach the selected network to the Oracle VM Server.

SR-IOV Enabled Networks

When you connect to network using SR-IOV enabled network interface, there is no virtual switch creation. An SR-IOV enabled network interface means that there are virtual functions created on the physical function of PCIe Endpoint device and you can assign the virtual functions to the logical domains. When you select SR-IOV option

while attaching networks to the Oracle VM Server, only the interfaces on which the virtual functions are created are available for network configuration.

When you assign SR-IOV enabled networks to logical domains, you cannot migrate the domains. SR-IOV enabled networks are available only from control domain and root domain.

SR-IOV enabled networks on root domain are available only in the following conditions:

- Oracle Solaris 11.1.4.5.0 or higher version is necessary for dynamic attach of networks.
- Oracle VM Server for SPARC 3.1 version to support SR-IOV networks on root domain.
- Refer to Oracle VM Server for SPARC Release Notes at `http://docs.oracle.com/cd/E38405_01/html/E38409/index.html` for hardware and firmware requirements for SR-IOV feature.

To Attach SR-IOV Enabled Networks to Oracle VM Server

1. Select the Oracle VM Server in the Assets section.
2. Click **Attach Networks** in the Actions pane.

   The list of available networks in Oracle Enterprise Manager Ops Center are displayed. The list also displays the existing number of connections with the server.
3. Select one or more networks from the list. You can make multiple connections to a network.

   Click **Next**.
4. Specify the number of connections for the selected networks. Increase the total number of connections.

   Note: The number of connections does not limit the number of logical domains that can be connected to this network.
5. Select the domain in the Domain list from which the network interfaces are selected. If the network interface is from the control domain, then it displays as `primary`. For network interfaces from the root domain is listed with the name of the domain.
6. Select the option SR-IOV and the NIC list is populated with the NICs that are SR-IOV enabled in the selected domain. The NICs that are defined as a physical function are listed.
7. Skip the option to define the IP address for the network interface. You can define the IP address in the logical domain OS later.

   Click **Next**.
8. Review the summary and click **Finish** to attach the selected network to the Oracle VM Server.

Network Options

To manage the attached networks, use the icons available in the **Network** tab of the Oracle VM Server for SPARC, as shown in Figure 19–6.

Figure 19–6 Network Options

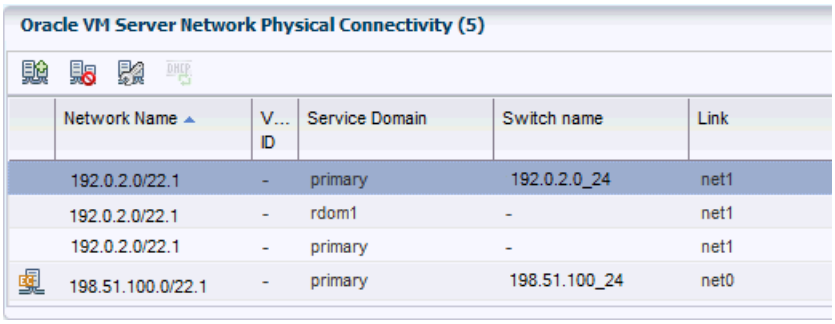

- **Refresh DHCP connectivity**: Reconnects to the DHCP server.
- **Unbind networks**: Disconnects the selected network from Oracle VM Server.
- **Modify physical connectivity**: Changes the connectivity attributes of the selected network, as shown in Figure 19–7.

Figure 19–7 Modify Physical connectivity

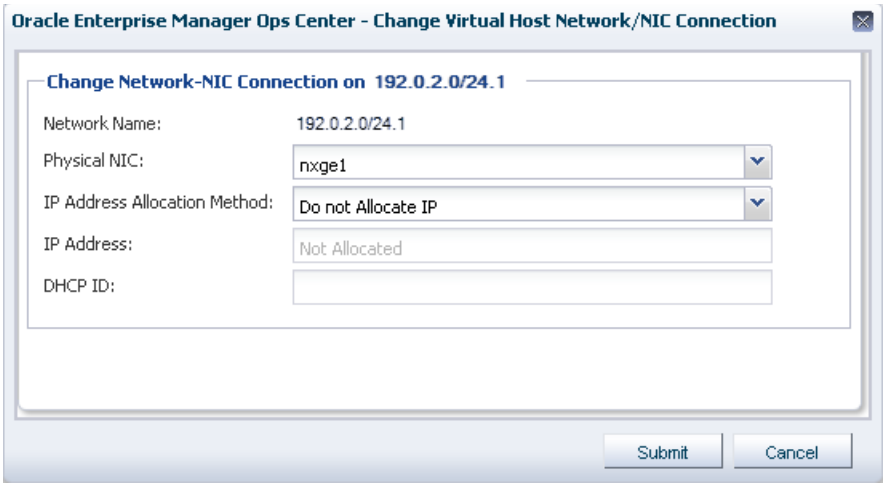

Unbinding Networks from Oracle VM Server for SPARC

You can remove the network connection from Oracle VM Server for SPARC. The removal of a network connection does not remove the virtual switch. You must select the appropriate virtual switch or the physical function to remove the network connection.

To Unbind Networks from Oracle VM Server for SPARC

1. Select the Oracle VM Server or control domain in the **Assets** section.
2. Select the **Networks** tab in the center pane.

   The list of networks that are connected to Oracle VM Server are displayed.
3. Select the network that you want to unbind from Oracle VM Server. You can unbind only one network at a time.
4. Click the **Unbind Network** icon.

   The Unbind Network Wizard is displayed.

5. In the Remove Network Connection step, the number of **New Connections Count** is automatically reduced by one and displayed. If you want to remove more than one network connection, enter the number of New Connections Count after the removal of the required network connections.

   Click **Next**.

6. Select the virtual switch or the physical function that connects the network to the physical interface.

   Ensure to select the correct virtual switch on the interface and click **Next**.

7. Review the summary of the network connection to be removed and click **Unbind** to remove the network connection.

Note: You cannot detach a network from Oracle VM Server if the network is connected to a virtual switch that is already in use.

Maximum Transmission Unit (MTU) Size

The default size of the network's MTU is 1500 bytes. A network can be created with MTU size that varies between 576 and 9216. Oracle Enterprise Manager Ops Center configures the IP configuration of Oracle VM Server for SPARC with the MTU value set at the network level.

When a network's MTU is modified and the network is attached a control domain and its logical domains, you must reboot the control domain for changes to take effect. You must also shut down and start the running logical domains attached to the network.

For a logical domain, the network to be attached must have minimum MTU size as 1500 bytes. The network interface cards driver configuration must be updated to support MTU size greater than or equal to 9216 in the following scenario:

- You have VLAN tagged networks created over the virtual switches with Link Aggregation as the uplink.
- You are using Oracle Solaris 11 OS.

Otherwise, the tagged networks become unreachable, and Oracle VM Server and logical domains cannot connect to these networks.

About Logical Domains

A logical domain is a virtual machine that has its own operating system and identity within a single SPARC server. Each logical domain can be created, destroyed, reconfigured, and rebooted independently, without requiring the server to be powered off. You can run a variety of applications in different logical domains to keep them independent for performance and security purposes.

Using Oracle Enterprise Manager Ops Center, you can create logical domains and provision Oracle Solaris OS on them. Using profiles and deployment plans, you can create more than one logical domain simultaneously and then save the configuration for future use.

You can use one of the following methods to create logical domains:

- **Create Logical Domain Profile and Plan**

  1. Create a profile which defines the configuration of the logical domain.
  2. Using the profile, create a deployment plan.

3. Apply the plan on an Oracle VM Server to create logical domains.
4. The logical domains do not have the OS installed. You must select each logical domain and apply an OS provisioning plan to install the OS. See Provisioning OS on Logical Domains for provisioning OS on the logical domains.

- **Configure and Install Logical Domains**

  This is a complex plan which contains deployment plans that create logical domains and install Oracle Solaris OS on them. You must have the required profiles and deployment plans available to create the complex plan. The outline to create a plan is as follows:

  1. Create a logical domain profile.
  2. Create an OS provisioning profile to install the OS.
  3. Create a deployment plan to install the OS and other updates if any.
  4. Create the Configure and Install Logical Domains plan.

You can now manage logical domains that were created manually using the native CLI in Oracle Enterprise Manager Ops Center. See Discovering Existing Oracle VM Server for SPARC Environments for more information.

Selection of Domains Types

Oracle Enterprise Manager Ops Center provides options to select the type of logical domain in the logical domain profile creation. You can select the following subtypes:

- Physical I/O domain
- Root domain
- Guest domain
- HA Guest domain

Select the appropriate subtype in the profile for creating logical domains.

Create Logical Domains

You can create logical domains with the following roles in Oracle Enterprise Manager Ops Center:

- Root domain: See Creating a Root Domain Profile
- Physical I/O domain: See Creating a Physical I/O Domain Profile
- Guest domain: See Creating a Guest Domain Profile
- HA Guest domain: See Creating HA Guest Domain Profile

Plan Your Domain Configuration

A logical domain profile captures the requirements and configuration of a logical domain, that includes the CPU Threads or whole-cores, memory, storage, and network details. You must provision Oracle Solaris OS on each domain separately. As an alternative, you can also use the complex plan Configuring and Installing Logical Domains that includes the OS provisioning profile.

Ensure that you have the following information before you create a profile:

- **System Requirements**: For creating I/O domains and root domains, the hardware and firmware requirements must be met. Refer to the Release Notes of Oracle VM Server for SPARC documentation at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html for hardware, firmware and software requirements for different versions of Oracle VM Server for SPARC and hardware.

- **CPU Resource Allocation – Virtual or Whole-Core**

  Each SPARC processor has multiple cores and each core has multiple CPU threads (virtual CPUs). The CPU threads are grouped into cores. For example, there are 4 Threads per core in the Oracle Sun Fire and SPARC Enterprise T1000 CMT processors.

  To a logical domain, you can either select to allocate the CPU resource as virtual CPU or as Whole-core, the default value is Whole-core. When you select to allocate as whole core, the specified number of cores and all their CPU threads are allocated to the logical domain.

- **Virtual CPU Allocation**

  The physical CPUs of the Oracle VM Server for SPARC are shared among the CPU threads of all the logical domains. Each logical domain requires:

  – At least one GB of memory, the default value is 4 GB.

  – At least one CPU thread.

    Each CPU thread can be allocated independently to logical domain. Some hardware resources are provided on a per-core basis and therefore, shared between the threads in the core. When the threads in a core are allocated to two or more separate logical domains, it can lead to suboptimal performance of those threads. To get the best performance, it is best to avoid allocating the threads of a core to many logical domains. The best approach would be create large logical domains first, the logical domains with complete cores, and then the smaller logical domains.

- **Whole-Core CPU Allocation**

  The CPU cores are allocated to a logical domain rather than virtual CPUs. You can also limit the maximum number of CPU cores that must be allocated to a logical domain. By default all logical domains are created for maximum throughput by tuning CPU cores to use a maximum number of CPU Threads.

- **CPU Architecture**

  The CPU architecture is of importance when you want to migrate the guest domains between systems that have different CPU processor type. The following are the supported types for CPU architecture:

  – `native`: This is the default value which enables the logical domains to be migrated only between systems that have the same CPU architecture type.

  – `generic`: Use this type when you want to enable the logical domain migration to take place independent of the CPU type of the systems. Though this might result in reduced performance, it provides an increased flexibility to migrate the domains between systems that have different CPU types.

- **Crypto Units**

  Crypto units are assigned based on CPU thread assignments. You can request the number of Crypto units to be assigned to the logical domain. However, the number of Crypto units assigned might be different than the amount requested

because we can only allocate a Crypto unit for every given number of CPU thread allocation, depending on the server hardware. After the creation of logical domain, view the job notification to see the actual number of Crypto Units assigned to the logical domain. You can also edit the logical domain configuration, storage, and network resource configuration later.

- **PCIe Buses**

  You must release the PCIe buses from the control domain. Only, then you can assign the PCIe buses to the root domain. When you configure and deploy Oracle VM Server for SPARC, you can select the option **Detach Unused Buses** from the control domain. Refer to Creating an OS Configuration Profile for Oracle VM Server for SPARC for more information. Otherwise, you must manually detach the buses to be used to create root domains.

- **Storage resources**

  The libraries for virtual disk can be local, local device, NAS, SAN, or Dynamic Storage libraries. The default value for storage size of a logical domain is 20 GB. If Oracle Solaris 11.2 is later installed, the OS provisioning job might fail if the storage size is less than 20GB.

  Oracle Enterprise Manager Ops Center provides multipathing configuration for logical domains. For each virtual disks that have more than one path to access the back-end storage, a multipathing group is created.

  Oracle Enterprise Manager Ops Center provides storage redundant access for logical domains. You can enable virtual disk multipathing to access the back-end storage of the domain by more than one path. Provide the name of the multipathing group and select the alternate path to access the virtual disk.

  The virtual disk of a logical domain remains accessible even if one of the service domains go down. For each virtual disks of the logical domain that has more than one alternate path, a new multipathing group is created. You must always specify the virtual disk of another service domain for the alternate path to access the back-end storage. You can select to enter the name of the multipathing group or a group is created automatically in the format of *logical domain name_mpGroup_devID*. The devID is the disk index. For example, if the name of the logical domain is *ldom*, then the multipathing group name is in the format of *ldom_mpGroup_1*.

  When you select more than one alternate path to the virtual disk, you must select which is the active path during the failure of the main storage path.

  When the virtual disk has only one alternate path to access the back-end storage, multipathing group is not created automatically unless you enter the name of the group.

- **Network resources and number of connections for a network**

  You can connect a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain. A vnet is created for each connection of the guest to the network. Vnets are re-used when the underlying virtual switch is used for networks with different VLAN IDs.

  When you start a logical domain, you must define the virtual switches or the virtual functions through which you connect the logical domain to the external network. If the root domain or I/O domain has connection to the physical network interfaces, you can use the direct physical network connection.

- **Automatic Recovery**

Select the automatic recovery option to recover the guest domains when the server hardware fails. This option is helpful when the Oracle VM Server is placed in a server pool and the guest domains can be recovered on the other Oracle VM Servers in the pool. If selected, enter the value for Priority of Recovery. The value must be between 0 and 100. A guest domain with a higher value is recovered first.

The guest domains are created with virtual disk multipathing configuration by providing the alternate path to access the back-end storage. During server hardware failure in a server pool, Oracle Enterprise Manager Ops Center tries to recover the guest domains on server that provides redundant access to the storage. You can select the option **Authorize Recovery without Redundant I/O** to authorize the recovery of the logical domains on other servers without redundant storage access in the server pool. The configuration of the guest domain is changed to lose the redundant storage access. If you do not select this option and there are no servers with redundant storage access, the automatic recovery is not performed.

About Root Domains

Root domain is the domain which has the entire PCIe bus also known as root complex assigned to it. The entire PCIe bus consists of the PCIe bus and all its switches and devices. When you assign a PCIe bus to the root domain, all the devices on that bus are owned by the domain. You cannot directly assign the PCIe Endpoint devices on that bus to any other domains. You must release the PCIe Endpoint devices so that it can be allocated to other I/O domains.

The root domain is non-migratable domain and the domain metadata is stored in the local file system.

The root domain can provide virtual I/O services to guest domains. When you create a root domain, you can define the name of the virtual disk storage server that provide virtual disk storage to guest domains. The names are checked for uniqueness during creation.

Creating a Root Domain Profile

To Create a Root Domain Profile

1. Select the **Plan Management** section in the Navigation pane.
2. Expand **Profiles and Policies** and select **Logical Domain** from the list.
3. Click **Create Profile** from the Actions pane.

   The Create Logical Domain Profile Wizard is displayed.
4. Provide a name and description to identify the profile and select Root Domain as the role for the logical domain.

   The option **Create a deployment plan for this profile** is already selected and creates a plan using this profile. If required, you can deselect this option.

   Click **Next** to specify the domain identity.
5. Provide a name and the starting number for the logical domain.

   Using this profile, you can create more than one logical domain. To identify each new logical domain, provide a prefix start name and starting number. For example, for a Start Name defined as TestDomain and a starting number of 10, three new logical domains have the name TestDomain10, TestDomain11, and TestDomain12.

To avoid number in the suffix, you can add the number in between the name using the wild card '%?d', where '?' is the number of digit displayed.

Click **Next** to configure the CPU and memory allocation for the domain.

6. Select one of the following options for CPU resource allocation:
   - **Virtual CPU**: You allocate CPU Threads to the root domain. Enter the number of **CPU Threads** that must be allocated to the root domain.
   - **Whole-core**: You allocate whole cores to the root domain. Provide the following data for whole-core allocation:
     - **CPU Cores**: Enter the number of CPU cores that you want to allocate to the root domain.
     - **Max CPU Cores**: The maximum number of cores that you can assign to an active domain.
7. (Optional) Enter the number of **Crypto Units** to be assigned to the root domain.

   The number of allocated Crypto Units might be different depending on the CPU Threads allocated and the server hardware. Check the job details for viewing the actual number of Crypto Units assigned to the logical domain.

   Note: The Crypto Units are not applicable from SPARC T4 server onwards and the values are ignored.

8. Enter the size of the memory to be allocated to the logical domain. The root domain requires at least four GB of memory per I/O device.

   Click **Next** to specify the I/O buses.

9. Enter the number of PCIe buses to be assigned to this domain. The number depends on the number of PCIe root complexes available on the hardware platform.

   Click **Next** to specify the library for virtual storage disks of the root domain.

10. The root domains are non-migratable as it is an I/O domain that has direct access to physical I/O devices. Therefore, the root domain metadata is stored in the local filesystem.
11. (Optional) Select one or more libraries that form the logical domain virtual disks. The libraries for virtual disk can be local, local device, NAS, SAN, or Dynamic Storage libraries.

    Note: It is recommended not to allocate virtual devices to root domain.

    Any physical storage or network devices attached to the root domain are made available and presented during OS provisioning.

    Click **Next** to specify the networks for the domains.

12. (Optional) Select the network from the list of networks identified by Oracle Enterprise Manager Ops Center. You can associate a root domain with more than one network.

You can use this step to attach networks using vnets. Skip this step if you want to connect to the networks available through the physical I/O devices. The physical connectivity is presented during OS provisioning.

13. (Optional) Enter the number of connections for each network.

    You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices (vnet) created for the logical domain. Oracle Enterprise Manager Ops Center re-uses the vnets and reduces the necessity to create vnets for each network connection.

    When you start a logical domain, you must define the virtual switches or the virtual function of SR-IOV enabled network interface through which you connect the logical domain to the external network.

    Click **Next** to view the summary of the details selected for creating a logical domain.

14. Review the information and click **Finish** to save the profile.

    When you want to create a deployment plan with this profile in the first step, then a corresponding logical domain plan is also created.

Applying Deployment Plan for Creating a Root Domain

You can create root domains on Oracle VM Servers that are stand-alone or placed in a server pool. When you want to select Oracle VM Servers placed in a server pool, remove the filter option in the Select Target Assets window and then select the Oracle VM Server which can be placed in a server pool.

Virtual Disk Server Creation

When the root domain has access to storage devices, then it can provide storage services to other guest domains. For providing storage services to other domains, a virtual disk server (vds) must be configured on the root domain. vds is required to export virtual disk devices to other domains. During root domain creation, a single vds is created by default with the default name *<domain name>-vds0*. You can change the name of the vds or retain the default name. You can also configure as many vds on the root domain. When you migrate guest domains that are utilizing the storage services from the root domain, the vds must be available to the target server.

To Apply Deployment Plan for Creating a Root Domain

1. Select the deployment plan created for root domain creation.
2. Click **Apply Deployment Plan** in the Actions pane.

   The Select Target Assets window is displayed.
3. The Oracle VM Server for SPARC servers are listed. To list the Oracle VM Servers placed in server pools, deselect the filter option.

   Select the targets and click **Add Targets to List**.
4. Select **Apply with minimum interaction** to apply the plan with the selected profile values. Otherwise, select **Allow me to override any profile values** to change the profile values.

   Click **Next** to specify the resource assignments for root domain.
5. If required, modify the name of the logical domains starting name and the number appended to it.

Click **Next** to configure the CPU Threads and memory.

6. Edit the CPU model, and memory as required.

   Click **Next** to specify the I/O buses to the domain.

7. Select the PCIe buses that have been released from the control domain.

   Click **Next** to specify the storage resource assignments and virtual disk server name.

8. (Optional) A default virtual disk server name is automatically set in the Name of the Virtual Disk Server to be created in the format of *<domain name>-vds0.* You can select to modify the name. The vds is to provide storage services to other domains.

9. (Optional) If required, you can add virtual disks to the root domain.

   The root domain is not migratable and therefore, the root domain metadata is automatically stored in the local library of the control domain.

   When the storage resources from the profile are not available for the selected target, it is flagged in red color. Modify the storage resources accordingly. Provide the name of the multipathing group and select the virtual disk server of the service domains and the active path. When a group name is not provided, a multipathing group is automatically created with the default naming procedure.

   Click **Next** to specify the network resource assignments.

10. (Optional) If you want to attach networks using vnets, use this step to provide the network connection. Select the networks and define the network resource assignment.

    When the network resource assignments from the profile are not available for the selected target, then it is flagged in red color. Modify the network resources accordingly in the profile or apply the plan in override any profile values mode.

    Click **Next** to schedule the job.

11. Schedule the job to run now or at a later time. Click **Next**.

12. Review the summary and click **Apply** to execute the deployment plan on the selected targets.

When you provision OS on the root domain, virtual functions are automatically created on the SR-IOV enabled PCIe buses that are allocated to the root domain.

Creating a Physical I/O Domain Profile

An I/O domain has direct access to and ownership of physical I/O devices. In Oracle Enterprise Manager Ops Center, you can create I/O domains by assigning PCIe Endpoint devices.

In the Logical Domain profile, select Physical I/O Domain subtype. In this guide, the I/O domains refer to the domains assigned with PCIe Endpoint devices only.

A domain that is assigned with a virtual function from SR-IOV enabled network interfaces is also defined as an I/O domain in Oracle VM Server for SPARC.

In Oracle Enterprise Manager Ops Center, to create domains assigned with virtual functions, you must use the Guest Domain subtype in the Logical Domain profile.

The number of I/O domains that you can create is limited by the number of PCIe Endpoint devices. For more information about assigning PCIe Endpoint devices and hardware and software requirements, see Oracle VM Server for SPARC documentation at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html.

Note: You cannot migrate I/O domains as they are assigned with PCIe Endpoint devices.

This section describes about creating I/O domains that are assigned with PCIe Endpoint devices.

To Create an I/O Domain Profile

1. Select the **Plan Management** section in the Navigation pane.
2. Expand **Profiles and Polices** and select **Logical Domain** from the list.
3. Click **Create Profile** from the Actions pane.

   The Create Logical Domain Profile Wizard is displayed.
4. Provide a name and description to identify the profile. The option **Create a deployment plan** is already selected for this profile and creates a plan using this profile. If required, you can deselect this option.
5. Select Physical I/O Domain from the Subtype.

   Click **Next** to specify the domain identity.
6. Provide a name and the starting number for the logical domain.

   Using this profile, you can create more than one logical domain. To identify each new logical domain, provide a prefix start name and starting number. For example, for a Start Name defined as TestDomain and a starting number of 10, three new logical domains have the name TestDomain10, TestDomain11, and TestDomain12.

   To avoid number in the suffix, you can add the number in between the name using the wild card '%?d', where '?' is the number of digit displayed.

   Click **Next** to configure the CPU and memory allocation for the I/O domain.
7. Select one of the following options for CPU resource allocation:
   - **Virtual CPU**: You define the CPU resource allocation in CPU Threads to the I/O domain. Enter the number of **CPU Threads** that must be allocated to the I/O domain.
   - **Whole-core**: You allocate whole cores to the I/O domain. Provide the following data for whole-core allocation:
     - **CPU Cores**: Enter the number of CPU cores that you want to allocate to the I/O domain.
     - **Max CPU Cores**: The maximum number of cores that can be assigned to an active domain.
8. Provide the values for Crypto units and memory allocation:
   - **Memory**: Enter the size of the memory to be allocated to the I/O domain.The size of the memory must be set according to the system size and load. The default is 4 GB and the minimum value is 1 GB.
   - **Requested Crypto Units**: Though you can enter the number of Crypto Units to be assigned to the logical domain, the number might be different depending on the CPU threads allocated and the server hardware. Check the job details for viewing the actual number of Crypto Units assigned to the logical domain.

   Click **Next** to specify the PCIe Endpoint devices for the I/O domain.

9. Select the type of PCIe Endpoint device and the number of devices. Select the type of device to filter only those I/O resources attached to it during resource assignment:
   - Ethernet Device
   - InfiniBand Device
   - Fibre Channel Device
   - SCSI Device

   Select Any type when you do not want to filter the I/O resources attached to the PCIe Endpoint.

   Click **Next** to specify the library for the storage disks for the logical domain.

10. (Optional) Select one or more libraries that form the I/O domain storage disks.

    The libraries for virtual disk can be local, local device, NAS, SAN, or Dynamic Storage libraries. You can skip this step when you want to provide the storage disks from the directly attached storage resources from the device.

    Click **Next** to specify the networks for the domains.

11. (Optional) Select at least one network from the list of networks identified by Oracle Enterprise Manager Ops Center. You can associate a logical domain with more than one network.

    Enter the number of connections for each network. You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain. When you start a logical domain, you must define the virtual switches or the virtual functions of SR-IOV enabled networks through which you connect the logical domain to the external network.

    Skip this step when you want to use the networks available from the physical devices attached to the I/O domain. You can define the physical connectivity during OS provisioning.

    Click **Next** to view the summary of the details selected for creating an I/O domain.

12. Review the information and click **Finish** to save the profile.

    When you want to create a deployment plan with this profile in the first step, then a corresponding logical domain plan is also created.

Applying Deployment Plan for Creating an I/O Domain

Virtual Disk Server Creation

When the I/O domain has access to storage devices, then it can provide storage services to other guest domains. For providing storage services to other domains, a virtual disk server (vds) must be configured on the I/O domain. vds is required to export virtual disk devices to other domains. During I/O domain creation, a single vds is created by default with the default name *<domain name>-vds0*. You can change the name of the vds or retain the default name. You can also configure as many vds on the I/O domain. When you migrate guest domains that are utilizing the storage services from the I/O domain, the vds must be available to the target server.

To Apply Deployment Plan for Creating an I/O Domain

1. Select the deployment plan created for I/O domain creation.
2. Click **Apply Deployment Plan** in the Actions pane.

   The Select Target Assets window is displayed.
3. The Oracle VM Server for SPARC servers that are available and eligible are listed. Select the targets and click **Add Targets to List**.
4. Select **Apply with minimum interaction** to apply the plan with the selected profile values. Otherwise, select **Allow me to override any profile values** to change the profile values.

   Click **Next** to specify the resource assignments for root domain.
5. If required, modify the name of the logical domains starting name and the number appended to it.

   Click **Next** to configure the CPU Threads and memory.
6. Edit the CPU model, and memory as required.

   Click **Next** to specify the PCIe Endpoint devices to be assigned to the domain.
7. The selected PCIe Endpoint Type and the corresponding PCIe Endpoint devices that have been released from the control domain or the root domain are listed.

   Select the PCIe Endpoint devices that you want to assign to the I/O domain.

   Click **Next** to specify the storage resource assignments and virtual disk server name.
8. (Optional) A default virtual disk server name is set in the format of <*domain name*>*-vds0* for the Name of the Virtual Disk Server to be created. You can select to modify the name. The vds is for providing virtual disk services to other domains.

   You can also edit the assignment of your storage resources.
9. (Optional) If required, you can add virtual disks to the I/O domain.

   The I/O domains are not migratable as they are attached to physical I/O devices and therefore, the domain metadata is automatically stored in the local library of the control domain.

   When the storage resources from the profile are not available for the selected target, it is flagged in red color. Modify the storage resources accordingly. Provide the name of the multipathing group and select the virtual disk server of the service domains and the active path. When you do not provide the group name and there is more than one alternate path to the virtual disk, a multipathing group is automatically created with the default naming procedure.

   Click **Next** to specify the network connection settings.
10. (Optional) All network connection require a virtual switch or virtual function to connect the network to the logical domain. Select the following options for each network connection:

    - **SR-IOV**: Select this option if you want to connect to the networks using the virtual function.
    - **Mode**: This option is applicable for networks that are configured with a VLAN ID. Select **Untagged** or **Tagged** mode from the list.

    Click **Next** to specify the network resource assignments.

11. In the network resource assignment, specify the domain that provides the network interface, and the virtual switch or the virtual function for the network connection.

    Note: You cannot edit the SR-IOV option in this step. Click **Back** to edit the SR-IOV option.

    Click **Next** to schedule the job.

12. Schedule the job to run now or at a later time. Click **Next**.
13. Review the summary and click **Apply** to execute the deployment plan on the selected targets.

Creating a Guest Domain Profile

To Create a Guest Domain Profile

1. Select the **Plan Management** section in the Navigation pane.
2. Expand **Profiles and Policies** and select **Logical Domain** from the list.
3. Click **Create Profile** from the Actions pane.

   The Create Logical Domain Profile Wizard is displayed.
4. Provide a name and description to identify the profile. The option **Create a deployment plan** is already selected and creates a plan using this profile. If required, you can deselect this option.
5. Select Guest Domain from the Subtype.

   Click **Next** to specify the domain identity.
6. Provide a name and the starting number for the logical domain.

   Using this profile, you can create more than one logical domain. To identify each new logical domain, provide a prefix start name and starting number. For example, for a Start Name defined as TestDomain and a starting number of 10, three new logical domains have the name TestDomain10, TestDomain11, and TestDomain12.

   To avoid number in the suffix, you can add the number in between the name using the wild card '%?d', where '?' is the number of digit displayed.

   Click **Next** to configure the CPU and memory allocation for the domain.
7. Select one of the following options for CPU resource allocation:
   - **Virtual CPU**: You allocate CPU Threads to the logical domain. Enter the number of **CPU Threads** that must be allocated to the logical domain.
   - **Whole-core**: You allocate whole cores to the logical domain. Provide the following data for whole-core allocation:
     - **CPU Cores**: Enter the number of CPU cores that you want to allocate to the logical domain.
     - **Max CPU Cores**: The maximum number of cores that can be assigned to an active domain.
8. Select the CPU architecture, Crypto units, memory allocation, and automatic recovery options:

- **CPU architecture**: Select the architecture as `native` or `generic` according to your requirement to migrate the logical domains.
- **Memory**: Enter the size of the memory allocated to the logical domain. The size of the memory must be set according to the system size and load. The default is 4 GB and the minimum value is one GB.
- **Requested Crypto Units**: Though you can enter the number of Crypto Units to be assigned to the logical domain, the number might be different depending on the CPU threads allocated and the server hardware. Check the job details for viewing the actual number of Crypto Units assigned to the logical domain.
- **Automatic Recovery**: Select whether you want the logical domain to be recovered automatically when the server hardware fails. If selected, enter the value for **Priority of Recovery**. The value must be between 0 and 100. A logical domain with a higher value is recovered first.

Click **Next** to specify the library for storing the domain metadata and storage disks for the logical domain.

9. Select a library to store the logical domain metadata.

   Select local storage library or NAS library to store the domain metadata.

10. Select one or more libraries that form the logical domain storage disks.

    The libraries for virtual disk can be local storage, local device, NAS, SAN, or Dynamic Storage libraries. Click **Next** to specify the networks for the domains.

11. Select at least one network from the list of networks identified by Oracle Enterprise Manager Ops Center. You can associate a logical domain with more than one network.

12. Enter the number of connections for each network.

    You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain.

    When you start a logical domain, you must define the virtual switches or the virtual functions through which you connect the logical domain to the external network.

13. Click **Next** to view the summary of the details selected for creating a logical domain.

14. Review the information and click **Finish** to save the profile.

    When you want to create a deployment plan with this profile in the first step, then a corresponding logical domain plan is also created.

Applying Deployment Plan for Creating a Guest Domain

Apply the created logical domain deployment plan on the target server and create the logical domains.

1. Expand **Deployment Plans** in the **Plan Management** section.

   and select the created logical domain deployment plan.

2. Select **Create Logical Domain** in the deployment plans.

3. Click **Apply Deployment Plan** in the Actions pane.

   The Select Target Assets window is displayed.

4. Select the target server from the list.

   Select one or more targets to apply the plan.

5. Click **Add to Target List** option to add the selected targets on which you want to apply the plan.

6. Select whether you want to run the plan with minimal interaction or you want to override the profile values.

7. Click **Next** to specify the resource assignments in the profile.

8. If required, modify the name of the logical domains starting name and the number appended to it.

   Click **Next** to configure the CPU Threads and memory.

9. Edit the CPU Threads, memory, and automatic recovery as required.

   Click **Next** to specify the storage resource assignments.

10. When the storage resources from the profile are not available for the selected target, it is flagged in red color. Modify the storage resources accordingly. Provide the name of the multipathing group and select the virtual disk server of the service domains and the active path. When you do not provide the group name and the virtual disk has more than one alternate path, a multipathing group is automatically created with the default naming procedure.

    Click **Next** to specify the network connection settings.

11. You can assign a virtual function or a vnet to the guest domain for the network connection.

    Select the following options for each network connection:

    - **SR-IOV**: Select this option when you want to assign a virtual function to connect to the network.
    - **Mode**: If the network is assigned a VLAN ID, you can select the network to be connected in tagged or untagged mode.

    Click **Next** to specify the network resource assignments.

12. You can select the domain which provides the network interface for the network connection. Select the following details for each network connection:

    - **Service domain**: Select the service domain that provides the network interface for the network connection. The domain can be primary, I/O domains or the root domains.
    - **Map Connection**: Select the virtual switch through which you want to connect the guest domain to the network. For SR-IOV enabled network connection, select the physical function through which you want to connect to the network.

    Click **Next** to schedule the job.

13. Schedule the job to run now or at a later time. Click **Next**.

14. Review the summary and click **Apply** to execute the deployment plan on the selected targets.

Creating HA Guest Domain Profile

The guest domain has the following high available features defined in the profile:

- Two network connections for each networks
- Redundant storage access to the virtual disks
- Automatic recovery of the guests during server failure in a server pool

The HA guest domain profile requires that you provide at least two connections for each selected network.

To Create a HA Guest Domain Profile

1. Select the **Plan Management** section in the Navigation pane.
2. Expand **Profiles and Policies** and select **Logical Domain** from the list.
3. Click **Create Profile** from the Actions pane.

   The Create Logical Domain Profile Wizard is displayed.
4. Provide a name and description to identify the profile. The option **Create a deployment plan for this profile** automatically creates a plan using this profile. If required, you can deselect this option.
5. Select Guest Domain from the Subtype.

   Click **Next** to specify the domain identity.
6. Provide a name and the starting number for the logical domain.

   Using this profile, you can create more than one logical domain. To identify each new logical domain, provide a prefix start name and starting number. For example, for a Start Name defined as TestDomain and a starting number of 10, three new logical domains have the name TestDomain10, TestDomain11, and TestDomain12.

   To avoid number in the suffix, you can add the number in between the name using the wild card '%?d', where '?' is the number of digit displayed.

   Click **Next** to configure the CPU and memory allocation for the domain.
7. Select one of the following options for CPU resource allocation:
   - **Virtual CPU**: You allocate CPU Threads to the logical domain. Enter the number of **CPU Threads** that must be allocated to the logical domain.
   - **Whole-core**: You allocate whole cores to the logical domain. Provide the following data for whole-core allocation:
     - **CPU Cores**: Enter the number of CPU cores that you want to allocate to the logical domain.
     - **Max CPU Cores**: The maximum number of cores that can be assigned to an active domain.
8. Select the CPU architecture, Crypto units, memory allocation, and recovery options:
   - **CPU architecture**: Select the architecture as `native` or `generic` according to your requirement to migrate the logical domains.
   - **Memory**: Enter the size of the memory allocated to the logical domain. The size of the memory must be set according to the system size and load. The default is 4 GB and the minimum value is one GB.
   - **Requested Crypto Units**: Though you can enter the number of Crypto Units to be assigned to the logical domain, the number might be different depending

on the CPU threads allocated and the server hardware. Check the job details for viewing the actual number of Crypto Units assigned to the logical domain.

- **Automatic Recovery**: Select whether you want the logical domain to be recovered automatically when the server hardware fails. When you select this option, provide the following details: If selected, enter the value for **Priority of Recovery**. The value must be between 0 and 100. A logical domain with a higher value is recovered first.
  - **Priority of Recovery**: The value for priority of recovery must be between 0 and 100. A guest domain with a higher value is recovered first.
  - **Authorize recovery without redundant I/O**: Select this option when you want to recover the guest domain in a server pool even when there is no redundant storage access in the other servers in the pool.

Click **Next** to specify the library for storing the domain metadata and storage disks for the logical domain.

9. Select a library to store the logical domain metadata.

   Select only the local storage or NAS library to store the domain metadata.

10. Select one or more libraries that form the logical domain storage disks.

    The libraries for virtual disk can be local storage, local device, NAS, SAN, or Dynamic Storage libraries. Click **Next** to specify the networks for the domains.

11. Select at least one network from the list of networks identified by Oracle Enterprise Manager Ops Center. You can associate a logical domain with more than one network.

12. Enter the number of connections for each network. You must enter at least two connections for each network.

    You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain.

    When you start a logical domain, you must define the virtual switches or the virtual functions through which you connect the logical domain to the external network.

13. Click **Next** to view the summary of the details selected for creating a logical domain.

14. Review the information and click **Finish** to save the profile.

    When you want to create a deployment plan with this profile in the first step, then a corresponding logical domain plan is also created.

Applying Deployment Plan for Creating a HA Guest Domain

Apply the created logical domain deployment plan on the target server and create the logical domains.

1. Expand **Deployment Plans** in the **Plan Management** section and select the created logical domain deployment plan.

2. Click **Apply Deployment Plan** in the Actions pane.

   The Select Target Assets window is displayed.

3. Select the target server from the list.

Select one or more targets to apply the plan.

4. Click **Add to Target List** option to add the selected targets on which you want to apply the plan.
5. Select whether you want to run the plan with minimal interaction or you want to override the profile values.
6. Click **Next** to specify the resource assignments in the profile.
7. If required, modify the name of the logical domains starting name and the number appended to it.

   Click **Next** to configure the CPU Threads and memory.
8. Edit the CPU Threads, memory, and automatic recovery as required.

   Click **Next** to specify the storage resource assignments.
9. When the storage resources from the profile are not available for the selected target, it is flagged in red color. Modify the storage resources accordingly. Provide the name of the multipathing group and select the virtual disk server of the service domains and the active path. When you do not provide the group name and the virtual disk has more than one alternate paths, a multipathing group is automatically created with the default naming procedure.

   Click **Next** to specify the network connection settings.
10. You can assign a virtual function or a vnet to the guest domain for the network connection.

    Select the following options for each network connection:

    - **SR-IOV**: Select this option when you want to assign a virtual function to connect to the network.
    - **Mode**: If the network is assigned a VLAN ID, you can select the network to be connected in tagged or untagged mode.

    Click **Next** to specify the network resource assignments.
11. You can select the domain which provides the network interface for the network connection. Select the following details for each network connection:

    - **Service domain**: Select the service domain that provides the network interface for the network connection. The domain can be primary, I/O domains or the root domains.
    - **Map Connection**: Select the virtual switch through which you want to connect the guest domain to the network. For SR-IOV enabled network connection, select the physical function through which you want to connect to the network.

    Click **Next** to schedule the job.
12. Schedule the job to run now or at a later time. Click **Next**.
13. Review the summary and click **Apply** to execute the deployment plan on the selected targets.

Selection of CPU Architecture

Newer platforms have a new class for the CPU architecture, named migration-class1 and sparc64-class1. The class1 architecture enables you to migrate the guest domains across systems while maintaining the full capabilities of the domains.

When you have a guest domain with the CPU architecture as `generic` in Oracle Enterprise Manager Ops Center and you start the guest in the following servers:

- Oracle SPARC T3 and T2 Servers: Oracle Enterprise Manager Ops Center sets the CPU architecture to generic.
- Oracle SPARC T4, T5, M5, and M6 servers: Oracle Enterprise Manager Ops Center sets the CPU architecture to `generic (migration-class1).migration-class1` is a cross-CPU migration family for SPARC platforms starting with the SPARC T4 servers.
- Fujitsu M10 Systems: Oracle Enterprise Manager Ops Center sets the CPU architecture to `generic (sparc64-class1).sparc64-class1` is a cross-CPU migration family for SPARC 64 platforms. sparc64-class1 value has more instructions than generic value. This value is compatible only with Fujitsu M10 systems.

Creating a Deployment Plan for Installing Logical Domain

When you create a logical domain profile, you can select to create a deployment plan with that profile. You can also create logical domain plans. You can either create a simple plan to install only the logical domains or a complex plan to create the logical domains and provision OS on it. The following procedure describes how to create a simple plan to install the logical domains only. You require to have created a logical domain profile to select in the plan.

To Create a Logical Domain Plan

1. Expand **Deployment Plans** in the **Plan Management** section and select **Create Logical Domains** plan.
2. Click **Create Plan from Template** in the Actions pane.

   The Create a Deployment Plan window is displayed.
3. Enter a name and description for the plan.
4. Select the failure policy.
5. Select the logical domain profile in the Create Logical Domain step.

   This is a simple single step plan to select the logical domain profile. The logical domain profiles that are available are listed in the list of **Associated Profile/Deployment Plan**.
6. Enter the number of logical domains to be created.
7. Click **Save** to create the deployment plan.

The deployment plan is created with the associated profile. Now, apply the plan on a suitable target to create the logical domains.

Logical Domains Created Using CLI

You can also create logical domains using the native CLI. The new logical domain created is automatically discovered and displayed under the corresponding Oracle VM Server for SPARC in the UI. The discovered logical domain metadata is stored in the local library of the control domain and the virtual disk storage is defined as Opaque.

The logical domains cannot be migrated as such and it requires the following actions to enable the migrate option:

- Move the metadata of logical domain to shared storage like NFS storage. See Moving Metadata to Another Library for complete procedure.
- Enable sharing of the virtual disk storage of the logical domain.

Provisioning OS on Logical Domains

The following information is covered in this section:

- Plan Your Network and Storage Resources
- Profiles and Deployment Plans for OS Provisioning
- Creating an OS Provisioning Profile
- Creating an OS Configuration Profile
- Creating Link Aggregation in Logical Domains
- Creating IPMP Groups in Logical Domains
- Apply the OS Deployment Plan for Logical Domains

The Oracle OS version that is provisioned on the logical domains is independent of the OS version on the Oracle VM Server for SPARC control domain. You can provision Oracle Solaris 10 or Oracle Solaris 11 OS on the logical domains irrespective of the Oracle Solaris OS version of the control domain.

After you create a logical domain, you must provision an OS on the domain.

- Ensure that you have imported the appropriate ISO image into the library. If it is Oracle Solaris 10, then it must be at least Oracle Solaris 10 8/07.
- Create or identify an OS provisioning profile for SPARC. See Chapter 13, "Operating System Provisioning" for more information.
- Create or identify a deployment plan with the OS provisioning and configuration profile and apply the deployment plan.

> **Note:** When you want to monitor the OS provisioning on the logical domains, you must first enable the logical domain console. See Connect to Logical Domain Console for enabling the console.

Oracle Enterprise Manager Ops Center provides OS provisioning and OS configuration profiles that capture the required parameters for provisioning OS. The profiles provide Logical Domain subtype that collects the parameters required for logical domain OS provisioning. You cannot use the Oracle Solaris OS provisioning profiles created for bare-metal provisioning.

The procedures to create the OS provisioning and configuration profiles especially for logical domains are described in this section.

Plan Your Network and Storage Resources

When you provision OS on the root domains and I/O domains, you can plan to provide the physical resources to the domains that are available from the PCIe Endpoint devices assigned to them. The physical network or the storage resources are presented while deploying the OS provisioning plan on the domains.

The storage resource can be either used by the domain or provide virtual disk server services to other guest domains. Whether the I/O domains are attached to a physical

storage device or not, a virtual disk server service is created for the virtual disks added to it. You can modify the name of the Virtual Disk Server name during the deployment of OS provisioning plan.

For network resources, the OS provisioning deployment displays the network controllers that provides the Ethernet connection and the available network ports in the controller. The PCIe Endpoint devices that are provided with Ethernet devices and the network ports that are connected must be known before deploying the OS provisioning plan on the I/O domain or root domain. Check with your administrator for the network connection details in the server hardware.

Profiles and Deployment Plans for OS Provisioning

Oracle Enterprise Manager Ops Center provides the following profiles and deployment plans for OS provisioning on logical domains:

- Profiles
  - OS Provisioning
  - OS Configuration
- Deployment Plans
  - Provision OS
  - Install Server
  - Configure and Install Logical Domains

The Provision OS deployment plan is a simple deployment plan that consists of two steps to provision and configure OS. You use this plan to do OS provisioning on the already created logical domains.

Install Server deployment plan is a multi-step deployment plan that includes step to update the OS, and install any software apart from OS provisioning. You use this plan on the already created logical domains.

Configure and Install Logical Domains is a complex plan that includes the step to create logical domains and provision OS on them.

Depending on the requirement in your environment, choose the appropriate method to create logical domains and provision OS on them. Using the complex plan Configure and Install Logical Domains might show some differences in the boot interface assignment and network resource assignment steps. For example, when you apply a provisioning plan on already created root domain, the OS provisioning deployment plan filters the network controllers that provides the network device and the network ports that are available from that device. Whereas, when you use a complex plan such as Configure and Install Logical Domains, the network resources are not filtered for the network devices. Instead, all the PCIe Endpoint devices assigned to the domain are listed. You must enter the correct network port for the OS provisioning job.

Oracle Enterprise Manager Ops Center also provides the option to identify the network interface by providing the MAC address.

Creating an OS Provisioning Profile

The OS provisioning profile collects the following details:

- OS image and version.

- OS setup parameters such as time zone, language, console, terminal type, and root password for the OS.
- File system layout.
- Name services.
- User account for Oracle Solaris 11

To Create an OS Provisioning Profile for Provisioning Logical Domain

1. Select the **Plan Management** section in the Navigation pane.
2. Expand **Profiles and Policies** and select **OS Provisioning** profile.
3. Click **Create Profile** in the Actions pane.

   The Create Profile - OS Provisioning Wizard is displayed.
4. Provide the following details for the profile identification:
   - Provide a name and description for the profile identification.
   - Select **Logical Domain** as the Subtype.

   Click **Next** to specify the provisioning parameters.
5. Select the OSP parameters:
   - Select the Oracle Solaris OS image from the list of available images. For Oracle Solaris 11 OS, select the SRU from the list.
   - Select the Software Group from the list.

   Click **Next** to specify the OS Setup.
6. Specify the OS setup parameters:
   - Enter the time zone, language, terminal type, console serial port, and console baud rate.
   - Enter the root password.
   - The NFS4 domain is set to dynamic in this example. If a naming service is configured in your environment, enter the NFS4 domain value.

   Click **Next** to provide a user account for Oracle Solaris 11. Otherwise, skip to Step 8.
7. Root login is not enabled in Oracle Solaris 11 OS. Create a user account to SSH to the OS after provisioning. Provide a user name and password for the account.

   Click **Next** to specify whether you want to use iSCSI disks for OS provisioning.
8. Select the option **Use iSCSI Disks** if you want to use iSCSI for OS provisioning.

   You must provide the storage server IP address and the volume group during deployment.

   Click **Next** to specify the file system.
9. The root (/) and a swap file system are defined by default. Click the **Add** icon to add more file systems. You can also change the file system type to UFS, or ZFS.

   Click **Next** to specify the name service.
10. If you have a naming service in place, select the appropriate one and provide the setup details. Otherwise, select **None** and proceed further.

- **DNS**: Enter the domain name of the DNS server and enter the IP address of the DNS server in the Name Server field. You can enter up to three IP addresses as the value for the Name Server. Provide the additional domains to search for name service information in the Domain Name Search List. You can specify up to six domain names to search. The maximum length of each search entry is 250 characters.
- **NIS or NIS+**: Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the option Specify an NIS Server and enter the NIS server host name and the IP address.
- **LDAP**: Enter the domain name of the LDAP server. Specify the name of the LDAP Profile you want to use to configure the system. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.

Click **Next** to view the summary of the parameters selected for the profile.

11. Review the parameters selected for the profile and click **Finish** to create the OS provisioning profile.

Creating an OS Configuration Profile

The OS configuration profile collects the information related to networking details for the OS. You can specify the number of interfaces that you want to specify for the OS.

Depending on the number of interfaces that you want to configure for the OS, the network connections are constructed for the OS during the deployment. The

1. Select the **Plan Management** section and expand **Profiles and Policies**.
2. Select **OS Configuration** and click **Create Profile** in the Actions pane.
3. Enter the following details to identify the profile:
   - Name and description of the profile.

     Note: Ensure that the name of the profile is not same as the OS provisioning profile.

   - Select Logical Domain as the Subtype and Virtual Machine as the Target Type.

   Click **Next** to set the OS Management properties
4. Select to manage the OS automatically and deploy the Agent Controller to manage the asset. Select the option **Enable Multiplexed I/O** so that you can associate block storage libraries such as SAN and iSCSI libraries for storage with the OS. If required, deselect the option **Enable Single Root I/O Virtualization (SR-IOV)** to disable SR-IOV on root domains.

   Click **Next** to specify the networking details.
5. Select one of the network options for the target system:
   - When you select **Use Link Aggregation**, refer to the section Creating Link Aggregation in Logical Domains to define the parameters for Link Aggregation.
   - When you select **Use IPMP**, refer to the section Creating IPMP Groups in Logical Domains for IP Multipathing.

- When you select **None**, you are forwarded to the next step to define the network interfaces and IP address allocation for the selected networks.

Click **Next** to specify the number of network interfaces that must be used on the OS.

6. Enter the number of network interfaces that must be used on the OS. The details of the interfaces are collected while deploying the plan.

   Click **Next** to view the summary of the parameters selected for OS configuration.

7. Review the parameters and click **Finish** to create the OS configuration profile.

Creating Link Aggregation in Logical Domains

While creating the profile for OS configuration, follow the procedure to create link aggregation:

1. Select **Use Link Aggregation** in the Specify Networks step of the profile.
2. Specify and configure the IEE 802.3ad Link Aggregation details:
   - **Link Aggregation Name**: Select the name of the Link Aggregation. The names are set to *aggr<x>*.
   - **Load Balancing Policy**: Define the policy for outgoing traffic.
   - **Aggregation Mode and Switches**: When the aggregation topology connects through a switch, determine whether the switch supports the Link Aggregation Control Protocol (LACP). When the switch supports LACP, you must configure LACP for the switch and the aggregation. Define one of the modes in which LACP must operate.
   - **LACP Timer**: Indicates the LACP timer value, either short or long.
   - **MAC Address Policy**: The MAC address policy is always defined to Auto.
   - **Number of Interfaces**: By default, two interfaces are defined in the aggregation. You can modify the number of interfaces for the aggregation.

   You can add multiple link aggregation and define the configuration parameters for each aggregation.

   You can define the interfaces during deployment. Click **Next** to continue the profile.

Creating IPMP Groups in Logical Domains

While creating the profile for OS configuration, follow the procedure to create IPMP groups.

1. Select **USe IPMP** in the Specify Networks step of the profile.
2. Specify and configure the IPMP groups:
   - **IPMP Name**: The name of the IPMP group is automatically displayed in the format of *ipmp<x>*. You can also modify the default name of the IPMP group.
   - **Failure Detection**: The detection can be Link-Based or Link-Based and Probe-Based.
   - **Number of Interfaces**: By default, two interfaces are defined in the group. You can modify the number of interfaces for the IPMP group.

   You can add multiple IPMP groups and select the failure detection for each group.

Click **Next** to specify the IPMP interfaces.

3. The interfaces are defined during the deployment of the profile. Specify the following information for each interface:
   - Select whether the interface is Failover or Standby Interface.
   - Select to Assign IP address or not during deployment to the specified NICs. The data and test addresses are defined during deployment.

   Click **Next** to continue the profile.

Apply the OS Deployment Plan for Logical Domains

When you apply the plan for provisioning the OS on logical domains, you must provide resource information in the following steps:

- Boot Interface Resource Assignments
- Storage Resource Assignments
- Specify Network Resource Assignments

When assigning a boot interface, if OpenBoot PROM (OBP) is used, only untagged networks are listed for selection. Only untagged networks can be used for OS provisioning as you cannot boot from a tagged network.

Other than these steps, the deployment plan in minimal interaction mode takes you through the confirmation of the OS provisioning and configuration profile information.

For logical domains that use vnets for network resource, the network resource step is displayed as shown in Figure 19–8.

Figure 19–8 Boot Interface Step for Virtual Controller

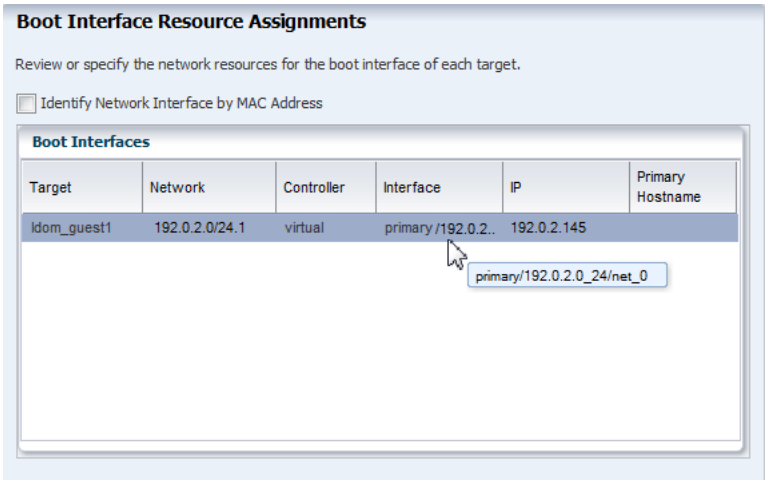

The Controller is displayed as virtual and the interface is the network interface that provides the virtual switch for the network connection of the logical domain.

For the other vnet network connections that must be configured with the OS are displayed as shown in the Figure 19–9.

Figure 19–9 Network Resource Assignments

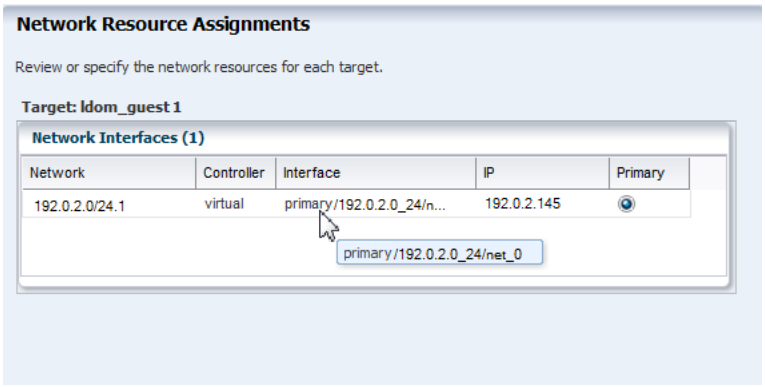

For every network connection that must be configured with the OS, requires a virtual switch.

Root Domains and I/O Domains

The physical network resources for root domains and I/O domains are displayed as shown in the Figure 19–10

Figure 19–10 Boot Interface Resource for Root Domain

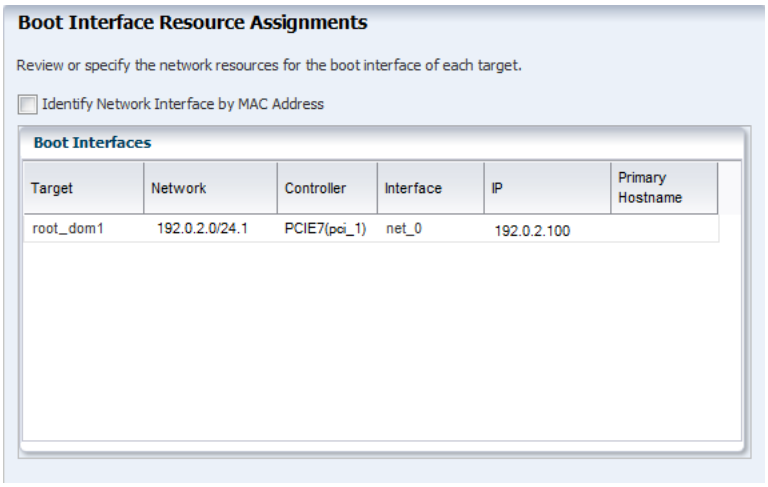

The network devices are filtered in the Controller list and displayed for selection. The Interface lists the network interfaces or the ports that are available in the selected network device.

If you know the MAC address of the network interface that must be used for the booting, you can use the option **Identify Network Interface by MAC Address**.

Figure 19–11 Boot Interface by Using MAC Address

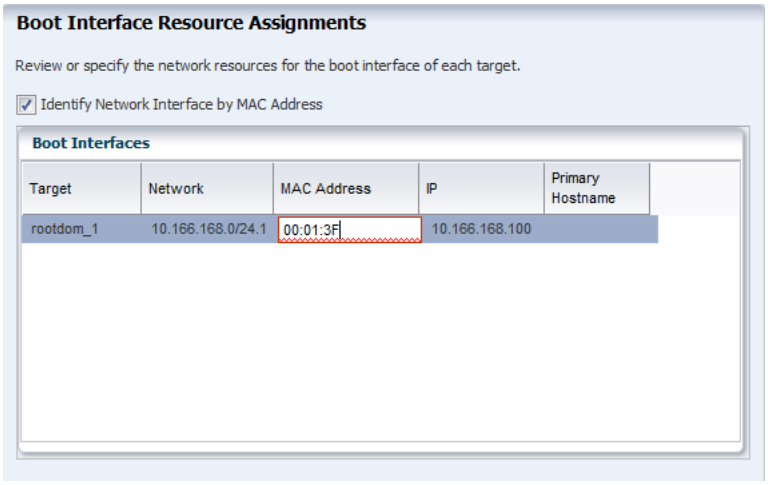

While configuring network on the OS, the first interface is always overwritten by the boot interface resource assignment. You can select the network resource that will be the primary network resource for the OS.

When you apply the combined plan of configuring and installing guest domain on a server pool, then in the OS deployment steps, the Interface column does not display the service domain, and the virtual switch. Only the network interface is listed as *net_x*. You must wait for the job to complete to see in which Oracle VM Server for SPARC system, the guest domain is created. The service domain and the virtual switch are auto assigned. See for the actual assignment after the job is completed.

For detailed information about creating plans and applying the deployment plan for OS provisioning, refer to Chapter 13, "Operating System Provisioning".

Creation of Virtual Functions

When you provision OS on the root domain, virtual functions are automatically created on SR-IOV enabled PCIe buses that are allocated to the domain. When you want to delete the root domain, it is required to delete all the virtual functions before proceeding to delete the root domain.

Manage Logical Domains

Oracle Enterprise Manager Ops Center provides options from the UI to manage the logical domain state, add resources, migrate and recover them. You can interchangeably use the CLI or the UI to manage the logical domains. The UI reflects the current state of the logical domains.

The logical domains that are discovered and managed in the UI does not have any difference in the management options in the UI except for the logical domain metadata.

You can manage logical domains with Oracle Enterprise Manager Ops Center or from the control domain. See Discovering Existing Oracle VM Server for SPARC Environments for information about managing from the control domain.

The following operations are available for managing the logical domains that are created using Oracle Enterprise Manager Ops Center or that are discovered and managed with an LDom VC Agent in Oracle Enterprise Manager Ops Center:

- View I/O Resources

- View Virtual Services
- Edit Logical Domain Configuration
- Shut Down a Logical Domain
- Starting a Logical Domain
- Add Storage to Logical Domain
- Moving Metadata to Another Library
- Enabling Shared Access for Opaque Storage Disks
- Connect to Logical Domain Console
- Delete a Logical Domain
- Cancel Delayed Reconfiguration

View I/O Resources

The logical domain that has I/O resources assigned to it displays the PCIe buses and the corresponding endpoints in the **I/O Resources** tab in the center pane.

The I/O resources tab also displays the network interfaces that are SR-IOV enabled. The physical functions and their corresponding virtual functions are listed in the **SR-IOV Services** sub tab. The guest domains to which the virtual functions are assigned are displayed.

View Virtual Services

Depending on the type of PCIe bus or the PCIe Endpoint assigned to the logical domain, it can provide services to other domains. If the PCIe Endpoint is attached to storage device, then the logical domain can provide virtual disk services to other domains. The type of device attached to the PCIe Endpoint or the bus, the I/O domains provide virtual services to other domains. These details are listed in the **Virtual Services** tab in the center pane of the selected logical domain.

- **Virtual disk server**: The virtual disk services allow you to export virtual disks to other logical domains. The virtual storage infrastructure enables other logical domains to access the storage disks that are not directly assigned to them. The virtual disk service processes the requests from the logical domains and submits them to the back end storage.
- **Virtual network switches**: The virtual network switch enables networking between virtual network devices in logical domains. The virtual switches that are created for the network interfaces in the I/O domain are listed. You can use the virtual switches to start or create logical domains. There is no virtual switch creation for SR-IOV enabled networks. See View I/O Resources for the SR-IOV enabled network interfaces and virtual functions.
- **Virtual console concentrator**: The virtual console concentrator that has been created in the logical domain is listed. You can use this console to redirect the console I/O from other domains.

Edit Logical Domain Configuration

Use the option **Edit Attributes** in the Actions pane of a selected logical domain. You can edit the name, description, CPU Model, Crypto Units, memory, and automatic recovery priority value.

You can modify the CPU Model of a logical domain between Virtual CPU and Whole-Core. The logical domain must be in shutdown or shutdown and detached state to modify the CPU Model.

When the logical domain is shutdown and detached, then you can modify the CPU architecture between `generic` and `native`. To modify the CPU architecture, the Oracle VM Server for SPARC must be at least running 2.2 version.

If you have any physical bindings constraint for the memory, CPU cores, or both, then you cannot edit CPU and memory configuration of the logical domain in the UI. Oracle Enterprise Manager Ops Center also does not provide options to explicitly assign physical resources such as CPU and memory. Whereas, you can use the CLI to assign physical resources to the logical domains. These are displayed in the Summary tab of the logical domains as Physical Bindings.

Also, you cannot switch the CPU Model for logical domains that are assigned with physical CPU and memory resources.

Use the option **Edit Tags** to modify the tag values or to add new tags.

Shut Down a Logical Domain

You can either shutdown or shutdown and detach the domain from its running state. Shutting down a logical domain means that the domain is still associated with the control domain, and connected to its network and storage resources. The domain is said to be in the state of shutdown. When you start a shutdown domain, you are not required to define the server and the resources.

When you shutdown and detach the domain from its running state, the domain is detached from the control domain, and disconnected from its network and storage resources. To start the domain, you have to define and select the server, network and storage resources. The option to detach the domain is not available in the following conditions of the domain:

- The domain is root or I/O domain.
- The domain is managed agentless.
- The domain metadata is on unmanaged storage. The scenario can occur when you discover and manage logical domains that are created using the CLI and storage of the logical domain is not moved to a managed storage.

A logical domain might be attached to virtual I/O devices on multiple I/O domains. You can still shut down the logical domain unless it does not provide any network, storage and console services to other domains.

Table 19–4 lists the logical domain configuration and shows the state, either shutdown or shutdown and detached.

Table 19–4 Shutdown or Shutdown and Detached State

Logical Domain	Shutdown	Shutdown and Detached
Created using Oracle Enterprise Manager Ops Center and Agent Managed	Yes	Yes
Discovered and LDom VC Agent Managed	Yes	Yes
Metadata on unmanaged storage	Yes	No
Agentlessly Managed	Yes	No

Table 19–4 (Cont.) Shutdown or Shutdown and Detached State

Logical Domain	Shutdown	Shutdown and Detached
I/O Domain	Yes	No
Root Domain	Yes	No

Starting a Logical Domain

You can start a logical domain from a shutdown or shutdown and detached state.

The logical domain in shutdown state is not detached from the server, and network resources and therefore, starting the logical domain immediately starts the logical domain in the Oracle VM Server. The Start action does not initiate the Start Logical Domain Wizard.

When the logical domain is in shutdown and detached state, it is disconnected from its network resources. To start a logical domain that is in a shutdown and detached state, you must define the network resources.

For logical domains placed in a server pool, in shutdown and detached state, the logical domain is detached from its control domain and disconnected from its network and storage resources. Starting this logical domain requires you to select the Oracle VM Server in which you want to run the logical domain, storage and network resources.

You must select the virtual switch or the virtual functions of SR-IOV enabled networks that connects to the logical domain MAC address. A virtual switch can be connected to only one MAC address. You can select not to connect to network for a virtual switch.

When a shutdown and detached guest is started, then Oracle Enterprise Manager Ops Center selects the appropriate and best option for the guest domain with generic CPU architecture.

When you have a guest with the CPU architecture as `generic` in Oracle Enterprise Manager Ops Center and you start the guest in the following servers:

- Oracle SPARC T3 and T2 Servers: Oracle Enterprise Manager Ops Center sets the CPU architecture to generic.
- Oracle SPARC T4, T5, M5, and M6 servers: Oracle Enterprise Manager Ops Center sets the CPU architecture to `generic (migration-class1).migration-class1` is a cross-CPU migration family for SPARC platforms starting with the SPARC T4 servers.
- Fujitsu M10 Systems: Oracle Enterprise Manager Ops Center sets the CPU architecture to `generic (sparc64-class1).sparc64-class1` is a cross-CPU migration family for SPARC 64 platforms. sparc64-class1 value has more instructions than generic value. This value is compatible only with Fujitsu M10 systems.

To Start a Logical Domain

1. Select the logical domain in shutdown and detached state.
2. Click **Start** in the Actions pane.

   The Start Logical Domain Wizard is displayed.
3. Select the server pool in which you want to start the logical domain.

   When you shutdown a logical domain, the logical domain is disassociated from the control domain and it can be started on a server.

When you have networks attached to a logical domain and is not available in the selected server pool, then that network is not available when the logical domain starts.

4. Select the Oracle VM Server for SPARC in which you want to start the logical domain. You can select the same Oracle VM Server in which the logical domain is running already. Otherwise, select an Oracle VM Server for SPARC from the list.

   The current load for all the Oracle VM Server for SPARC are displayed. Use this information to place your logical domain.

   Click **Next** to specify the network interfaces.

5. Select the virtual switch or the virtual function that must be associated with the logical domain MAC address.

   The number of virtual switches displayed depends on the number of network connections for the server pool. For each network connection, you have a virtual switch created. You can associate a virtual switch with only one MAC address.

6. Select **Do Not Connect** when you do not want the virtual switch to connect to MAC address.

   Click **Next** to schedule the job.

7. Schedule the job to run now or at a later time.

8. Review the properties and click **Start** to run the logical domain.

Add Storage to Logical Domain

You can add virtual disks to a logical domain. The associated libraries to the Oracle VM Server for SPARC are available to be added as additional storage to logical domains. Use the option **Add Storage** in the Actions pane to add virtual disks to the logical domain. The logical domain can be in running state for adding storage.

You can add virtual disks from the following type of storage libraries:

- **File system storage**: These are NAS libraries. Specify a virtual disk name and size of the disk.
- **Static block storage**: Select LUNs from SAN or iSCSI storage servers. The LUNs sizes are fixed.
- **Dynamic block storage**: Add LUNs from SAN or iSCSI storage servers discovered and managed in Oracle Enterprise Manager Ops Center. You can also create LUNs by selecting a volume group and specifying the size of the LUN.

Starting in 12.2.2.0.0 release, you can add multiple SAN LUNs to a logical domain from different libraries using the **Add Multiple SAN LUNs** option of the **Add Storage** wizard. Refer to *Oracle Enterprise Manager Ops Center Adding Volumes to SAN Storage Libraries* for more details about using this option.

Virtual Disk Multipathing

Virtual disk multipathing enables you to configure a virtual disk on a logical domain to access its back-end storage by more than one path. The paths lead through different service domains that provide access to the same back-end storage. The virtual disk of a logical domain remains accessible even if one of the service domains go down. For each virtual disks of the logical domain that have more than one alternate path to the backend storage, a new multipathing group is created automatically. You must always specify the virtual disk of another service domain for the alternate path to access the back-end storage. You can select to enter the name of the multipathing group or a

group is created automatically in the format of *logical domain name_mpGroup_devID*. The devID is the disk index. For example, if the name of the logical domain is *ldom*, then the multipathing group name is in the format of *ldom_mpGroup_1*.

When you have more than one alternate path for the virtual disks, you can select the active path which serves during a failure in one of the service domains.

Unless you enter the name of the multipathing group, the group is not created when there is only one alternate path to access the back-end storage.

Moving Metadata to Another Library

Metadata for a logical domain is saved in either a local library or a shared library. When you install Oracle VM Server for SPARC on a system without using Oracle Enterprise Manager Ops Center, the metadata is saved in a default local library. When you install Oracle VM Server for SPARC with Oracle Enterprise Manager Ops Center, you define a shared storage location for metadata.

To migrate a logical domain to another system, both systems must use the same virtual disks for shared storage and the same shared library for the logical domain metadata.

The following conditions must be met before you can move the logical domain metadata to another library:

- The control domain must be managed with an LDom VC Agent.
- The Logical Domain must be known to Oracle Enterprise Manager Ops Center.
  - The Logical Domain and its associated resources are created with Oracle Enterprise Manager Ops Center
  - A manually created logical domain is agent managed with Oracle Enterprise Manager Ops Center

To Move Metadata to Another Library From the Control Domain View

1. Expand **Libraries**, expand **Storage Libraries**, then select the source library in the Navigation pane.
2. Click the **Usage** tab. Select the guest in the Logical Domains table, then click the **Move Metadata** icon.

Figure 19–12 Move Metadata Icon

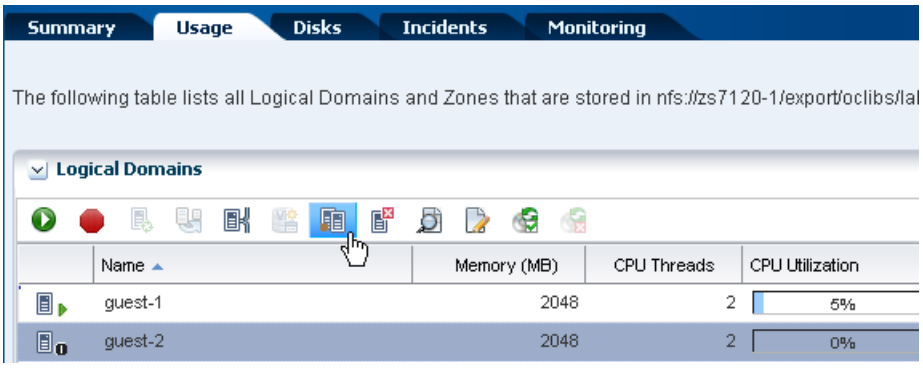

3. Select the new library from the list of available libraries to store the logical domain metadata. Optionally, add a description and add tags. Click **Next**.
4. Review the Summary, then click **Submit**.

When the job completes, the logical domain's metadata is located on the new library.

To Move Metadata to Another Library From the Logical Domain View

1. Expand **Assets**, then select the shutdown guest in the Assets tree.
2. Click **Move Metadata** in the Actions pane.
3. Select the new library from the list of available libraries to store the logical domain metadata. Optionally, add a description and add tags. Click **Next**.
4. Review the Summary, then click **Submit**.

When the job completes, the logical domain's metadata is located on the new library.

Enabling Shared Access for Opaque Storage Disks

When you create logical domains using the native CLI, it is automatically discovered and displayed under the corresponding Oracle VM Server for SPARC in the UI. The metadata of the logical domain is stored in the local library of the control domain. If the logical domain storage disks are on storage servers that are not managed by Oracle Enterprise Manager Ops Center, then they are tagged as Opaque and displayed in the UI. Oracle Enterprise Manager Ops Center does not identify whether the storage is local or shared. The migration action is disabled for the logical domain that has the metadata on local library and the storage disks tagged as opaque. Use **Move Metadata** option to move the logical domain metadata to a shared storage.

Mark the logical domain storage disks as shared to indicate that the storage is available to other managed Oracle VM Servers using the same backend name.

The storage that is marked as shared must be available for all the members in a server pool.

To Enable Sharing for Opaque Disks

1. Select the logical domain that has been created using CLI.
2. Select the **Storage** tab in the center pane.
3. Select the disks tagged as Opaque in the Disk Type.
4. Click the **Enable Sharing** icon.

The selected disks are marked as sharing in the UI.

Connect to Logical Domain Console

You can attach to the logical domain console within the Oracle Enterprise Manager Ops Center UI. You enable the console connection and then you connect to the console.

Note: When the logical domain is not in view in the Assets tree, the console is logged out automatically but the connection exists until the connection time expires. You must log in again.

Delete a Logical Domain

When you delete a logical domain, it is disconnected from the associated networks and is disassociated from the Oracle VM Server. All of the associated resources are released and the domain configuration is removed from the library. All references to the logical domain, including its disk image and snapshots are removed from the system. However, the logical domain profile and plan remain unaffected.

When you delete a logical domain manually, that is, using the native CLI, the UI reflects the removal of the logical domain.

Starting in the 12.2.2.0.0 release, you can use the **Delete** action to delete root domains that have virtual functions on SR-IOV enabled PCIe devices if other domains are not using a virtual function from that root domain.

In the 12.2.1.0.0 release and in previous releases, to delete root domains that have virtual functions on SR-IOV enabled PCIe devices, you must first remove the virtual functions manually by using he following command from the respective control domain:

```
ldm destroy-vf VF-DEVICE-NAME
```

> **Note:** You must delete the virtual functions in the reverse order. You must start deleting from the last virtual function on the PCIe device

If you want to delete guest domains assigned with virtual functions, you must first remove the virtual functions from the domain and then delete the guest domains.

Refer to Oracle VM Server for SPARC documentation at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html for more information.

Cancel Delayed Reconfiguration

The action **Cancel Delayed Reconfiguration** is restricted to the control domain or a PCIe root domain, if supported.This action is disabled for all other logical domains. See Delayed Reconfiguration Mode, for more information about the delayed reconfiguration.

Managing Logical Domain Networks

The following information is covered in this section:

- VLAN Tagging Support
- Alternate MAC Addresses
- Connecting to Network
- Port Connectivity

You can connect or disconnect a network from a logical domain in running state. You can connect to a network multiple times. Use the **Connect Network** and **Disconnect Network** icons, shown in Figure 19–6, in the Networks tab.

You can also connect to SR-IOV enabled networks and thereby assign virtual functions to the logical domain. The SR-IOV enabled networks are available only on control domain as the SR-IOV feature is enabled only on the SR-IOV cards on the primary domain. From Oracle VM Server for SPARC 3.1 version, SR-IOV enabled networks are also available from root domains.

Each connection requires a virtual switch or the virtual function of SR-IOV enabled networks to connect the logical domain to the network.

Virtual network device (vnet) is a virtual device that is defined in the logical domain when it is connected to a virtual switch. You can create several vnets over a virtual switch. Creating too many vnets also decrease the performance. Therefore, the creation of vnets are reduced and Oracle Enterprise Manager Ops Center ensures to re-use

when possible. For example, when there is an existing vnet that has the connectivity to the required network and there is no VLAN ID incompatibility then the vnet is re-used.

You can re-use a virtual switch to make multiple connections to a network from the logical domain or use the same virtual switch to connect to a network for different logical domains.

For SR-IOV enabled networks, you can use a virtual function to connect to a network multiple times for a logical domain or connect to a network for multiple logical domains.

Migration is disabled for domains that has one or more virtual functions assigned to it. For virtual switches, the virtual switch name must be identical in the source and target machine to migrate the guest domain.

Example 19–2 Re-using Vnets

Consider a virtual switch *vsw0* connected to networks N1 and N2. Network N1 is a regular LAN without a VLAN ID and Network N2 is a VLAN ID with VID = 11. When you connect a guest G1 to network N1 through *vsw0*, a virtual network device *vnet0* is created. The *vnet0* is created with PVID = 1 and VID = <null>.

When you connect the guest G1 to network N2 through *vsw0*, then the vnet, *vnet0*, is re-used by setting the VID = 11 which is the VLAN ID of the network N2.

VLAN Tagging Support

Oracle VM Server for SPARC software supports 802.1Q VLAN-Tagging in the network infrastructure. The option to define the VLAN tagging mode is available from the Oracle Enterprise Manager Ops Center UI when you connect networks configured with a VLAN ID to the guest domains.

Configure the VLAN interfaces over the physical network devices in Oracle VM Server for SPARC. When you connect these networks to the guest domains, define the tagging mode as tagged or untagged.

In untagged mode, the Port VLAN ID (PVID) is set to the VLAN ID of the network. The outbound frames from the virtual network that are untagged are tagged with this PVID by the virtual switch. You can have only one PVID for a vnet or a virtual function (VF). If the PVID option is selected for a virtual function, you cannot use the VID option; if the VID option is selected, you cannot use the PVID option.

In tagged mode, the network's VLAN ID is added to the VID of the vnet or virtual function. The vnet sends and receives tagged frames over the VLANs specified by its VLAN IDs (VIDs).

You must have at least one untagged network connection to be used for provisioning OS on the logical domain. If the network to be used for OS provisioning is already VLAN tagged, then select the untagged mode while connecting the network. The network is assigned the PVID with the VID of the network.

When the tagging mode for a network configured with VLAN ID is set to tagged at the Oracle VM Server, then you have the option to select tagged or untagged option when you connect the logical domain to the network. If the tagging mode is set to untagged, then you cannot the change the tagging mode and it is set to untagged.

Alternate MAC Addresses

When you connect network to a logical domains, alternate MAC addresses are created automatically for the vnets or the virtual functions assigned to the logical domain. These alternate MAC addresses can be used to create zones in the logical domains.

The number of MAC addresses to create is to set to 10 by default. If required, you can change the number of MAC addresses to create in the `/opt/sun/n1gc/lib/XVM_SATELLITE.properties` file. In the file, set the following property

```
ldom.auto.alt.mac.addrs.count=20
```

If you do not want any MAC address to be created, you can change to:

```
ldom.auto.alt.mac.addrs.count=0
```

The property is set globally for each Oracle VM Server for SPARC managed in Oracle Enterprise Manager Ops Center. Restart the Enterprise Controller after you change the property. Only after the restart, the number of alternate MAC address change is applied when you connect a network to the logical domain. The number is not altered for any existing connection.

To create alternate MAC addresses, the domain must be in shutdown or shutdown/detached state.

The alternate MAC address are not created in the following scenarios or conditions:

- When you connect to a network dynamically, that is when the logical domain is in running state.
- The network connection is not done using Oracle Enterprise Manager Ops Center.

> **Note:** You need alternate MAC addresses to create zones on logical domains. Refer to Oracle VM Server for SPARC documentation at http://docs.oracle.com/cd/E38405_01/html/E38406/assignmacaddressesautomaticallyormanually.html#scrolltoc to create alternate MAC address manually.

- When you use InfiniBand virtual functions.
- The Oracle VM Server for SPARC is lower than 3.0.0.2 version.

Connecting to Network

Figure 19–13 Logical Domain Network Options

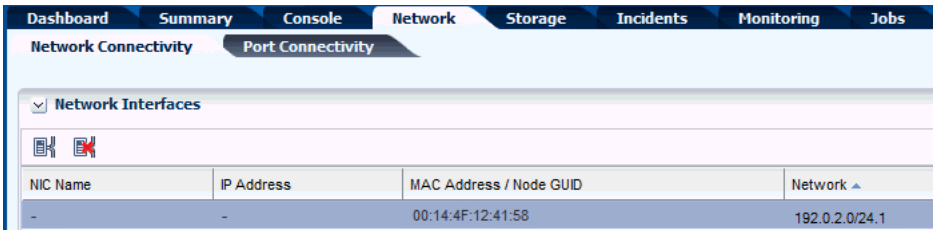

To Connect Networks to Logical Domains

1. Click the **Connect Network** icon to display the Connect Guests to Network window. Select the network.

2. Select the networks from the list. For each selected network, choose the following details:
   - Select the domain that provides the network services or the SR-IOV network device. The domain can be either control domain or root domain.
   - Select the Mode as Tagged or Untagged for networks configured with a VLAN ID.
   - If you want to assign virtual function of a SR-IOV enabled network, then select the SR-IOV option. The Map Connection lists the physical function that are available from the selected domain.
   - From the Map Connection list, select the virtual switch or the physical function through which you want to connect to the network.

     An existing vnet is re-used to connect to the same virtual switch or virtual function rather than creating a new vnet for each connection.
3. Click **Connect to Network** to connect the networks to the logical domains.

When you connect networks to a guest domain that is in shutdown/detached state, then the service domain and Map Connection options are not available in the Connect Networks window.

You can create IPMP groups or aggregate the NICs allocated to the logical domain. See Chapter 17, "Networks for Virtualization" for more information about creating IPMP groups and link aggregation.

When you disconnect networks from logical domains, the vnets are destroyed. If the vnets are re-used with another network connection, then it cannot be removed. To remove the vnet completely, you must remove the network connections in the following order:

- You must detach the network in the order of highest to lowest for net*x*. For example, a network connection with net5 as the highest must be removed first.
- The network connection for a network with a VLAN ID must be the last network connection to be removed.

To Disconnect Networks from Logical Domains

1. Select the networks that you want to disconnect from the logical domain in the Network tab of the selected logical domain.
2. Click **Disconnect Network** icon.

   Unbind a Network window is displayed.
3. Click **Disconnect From Network** to confirm the delete action.

Port Connectivity

The Port Connectivity sub tab in the Network tab provides the details of the logical domain network connection. The port connectivity tab displays the following information:

- Port Type
- Media Type
- VLAN ID and Port VLAN ID
- Alternate MAC Address

- MAC Address
- Port Name

Figure 19–14 Port Connectivity

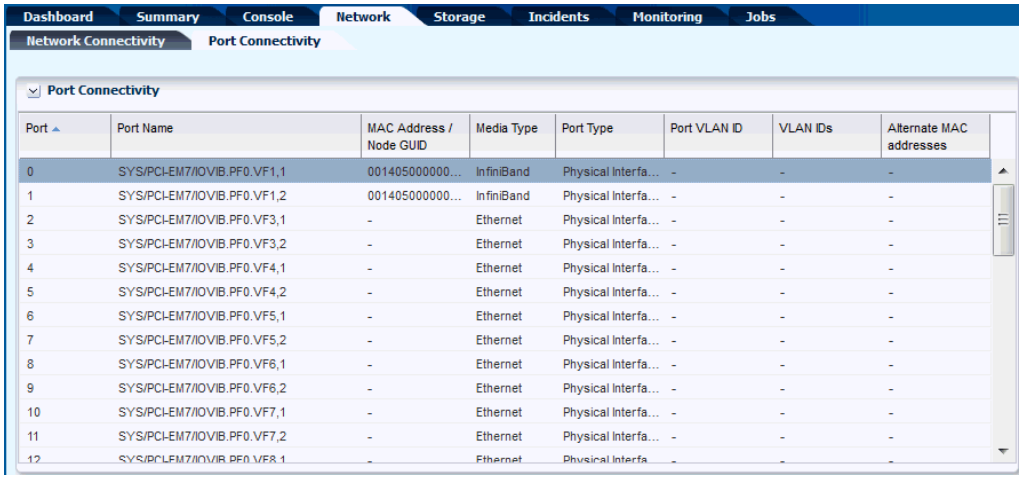

The Port Name displays the name of the vnet created. When the name of the vnet is of the format *vnicxxxxx*, it indicates that the vnet is created by Oracle Enterprise Manager Ops Center. Whereas, all other formats indicate that the network was not connected to the logical domain using Oracle Enterprise Manager Ops Center.

Also, when the network is connected to logical domain using Oracle Enterprise Manager Ops Center, then the alternate MAC address shows the list of created address provided the properties is set to create MAC address.

The Port Type indicates whether the type of underlying network connection for the logical domain. The type can be Physical Interface, SR-IOV Virtual Function, or Ldom vnet. The Media Type indicates whether the network media is Ethernet or InfiniBand.

Migrate Logical Domains

The following information is covered in this section:

- Setting User Accounts for Migration
- Migration Requirements
- Migrating a Logical Domain
- Migrating Multiple Logical Domains

Migrating a logical domain means moving the logical domain from one server (source) to another server (target). When a domain is migrated from the source to the target server in its running state without any impact to its availability is called live migration.

When the logical domain is stopped and then migrated to the target server, it is called as cold migration. You can migrate the logical domains when the requirements are met for successful migration.

Apart from these, Oracle Enterprise Manager Ops Center also migrate domains to other servers in the following scenarios:

- Balancing the load of the servers in a server pool

- Recovery of logical domains when the server on which it is running fails

Live migration, logical domain recovery and load balancing are available and supported only in the context of a server pool.

> **Note:** From Oracle VM Server for SPARC 2.1 and higher versions, you can migrate the logical domains only when the source and target machines are running at least Oracle VM Server for SPARC 2.1 version.

Oracle Enterprise Manager Ops Center UI provides the Migrate option to migrate the logical domains. The Migrate Logical Domain Wizard provides a list of available compatible target machines to migrate the logical domain. Select a target and migrate the logical domain.

Setting User Accounts for Migration

Temporary user accounts are created dynamically when a logical domain migration is initiated. The user account is deleted after the migration. If your datacenter environment does not support dynamic creation of user accounts, then you can set the logical domain migration to use an existing user account instead of creating a temporary one. You must have the Security Admin role to create the credentials. If not, contact the Ops Center Administrator for creating a user account for logical domain migration.

To create a user account, follow the procedure:

1. Select the **Administration** section in the Navigation pane.
2. Select **Credentials** in the Administration section.
3. Click **Create Credentials** in the Actions pane.

   The Create Credentials window is displayed.
4. Select the protocol as OVM For SPARC Migration and enter the following details:
   - Name of the credential
   - User name for the account
   - Password for the user account
5. Click **Create** to create the user account.

Set the following system property in the configuration management of the Enterprise Controller:

1. Select the **Administration** section in the Navigation pane.
2. Select the **Configuration** tab in the center pane.
3. Select **Virtualization** in the Subsystem.

   The property Virtualization.Ldoms.Migration.Username is displayed.
4. Enter the user name of OVM For SPARC Migration user account created previously to be used for all logical domain migration.
5. Click **Save**.

The migration job proceeds to use the user account as defined in the system property, that user account must exist on the source and target systems. If the system property is

set default to null, then the migration job dynamically creates a temporary user account. Ensure to set this property before initiating any logical domain migration, if required.

Migration Requirements

When you migrate logical domains, there are certain requirements and restrictions that must be noted for successful migration. Otherwise, there can be loss in the configuration of the logical domains or unsuccessful migration job. The requirements and restrictions that affect the migration of the logical domains are:

- **Shared Storage**: The source and the target servers must have access to common storage resource. The logical domain metadata must be stored on NFS storage that must be associated with the server pool or the stand-alone Oracle VM Servers. The virtual disks must be shared. If the virtual disks are on shared storage that are not managed by Oracle Enterprise Manager Ops Center (other than FC or iSCSI LUNs, or NFS shares), enable shared access for the storage disks. Otherwise, use shared FC or iSCSI LUNs, NFS shares for domain's storage.
- **CPU Architecture**: The CPU architecture is of importance when you want to migrate the logical domains between systems that have different CPU processor type.

  The following are the supported types for CPU architecture:

  - `native`: The logical domains that are created with this CPU type enable them to be migrated only between systems that have the same CPU architecture type.
  - `generic`: The logical domains that are created with this CPU type enable them to migrate to systems independent of the CPU type. Though this might result in reduced performance, it provides an increased flexibility to migrate the domains between systems that have different CPU types.

    From Oracle VM Server for SPARC 3.1 and higher versions, the value of generic varies according to the platform. For example, if generic is set for a guest CPU architecture on an Oracle SPARC T2 server, then the CPU architecture is set to generic. Whereas for servers starting from Oracle SPARC T4, when the CPU architecture is set to generic in the UI, Oracle Enterprise Manager Ops Center sets the CPU architecture to migration-class1. In the UI, the CPU architecture is displayed as `generic (migration-class1)`. For Oracle M5 and M6 servers also, `generic (migration-class1)` is displayed. For Fujitsu M10 systems, the guests that are set to generic CPU architecture, are automatically set to sparc64-class1 CPU architecture. In the UI, the CPU architecture is displayed as `generic (sparc64-class1)`.

    Depending on the server platform onto which the guest is started after shutdown-detach, the best option for the CPU architecture is selected when the guest has generic as the CPU architecture.

  You can edit the CPU architecture of the logical domain as required to facilitate the cross CPU live migration. You must shut down and detach the guest to edit the CPU architecture.
- **Logical Domains Created Using CLI**: For logical domains that are not created using the Oracle Enterprise Manager Ops Center UI, you must move the metadata of the discovered logical domains to a shared NFS storage. The virtual disks of the discovered logical domains are marked as opaque and it must be shared to enable migration.

- **Whole-core**: When you migrate the logical domains from latest versions of Oracle VM Server for SPARC to earlier versions results in losing the whole core configuration. For example, the logical domains configured with whole cores on Oracle VM Server for SPARC 3.0 version loses it whole-core configuration when migrated to Oracle VM Server for SPARC 2.2 version. Depending on the original number of cores allocated, it is translated to the number of CPU Threads and allocated to the logical domain.
- **Power Management Policy**: The power management policy for Oracle Integrated Lights Out Manager (ILOM) firmware can be set to Disabled, Elastic or Performance. For Oracle VM Server for SPARC 2.2 and earlier versions, the migration of logical domains are not supported when the source and target machines are set with elastic power management policy. You must change to Performance policy to migrate the logical domains.

  From Oracle VM Server for SPARC 3.0 version onwards, the Elastic policy does not affect the migration of the logical domains. The power management policy Performance and Elastic are fully compatible with the whole-core constraint.
- **Physical I/O devices**: The logical domains that are attached to physical I/O devices, namely the I/O domains and root domains, cannot be migrated.
- **Virtual Switch**: When you migrate logical domains, the virtual switch name must be identical in the source and the target Oracle VM Server for SPARC.
- **Virtual Functions**: You cannot migrate logical domains that are connected to virtual functions.
- **Virtual I/O Devices**: The logical domains can be attached to virtual I/O devices from multiple I/O domains. You can still migrate the logical domains provided all the virtual I/O services are available on the target machine.
- **Virtual Disk Server**: The target and the source Oracle VM Server must have the same names for the virtual disk server to migrate the guest domain. When the guest domains uses I/O resources from I/O domains or root domains, then the target Oracle VM Server must have domains that use the same name for the virtual disk server as the source Oracle VM Server.
- **MAU**: For Oracle VM Server for SPARC 1.2 version, you can migrate with only one MAU. For Oracle VM Server for SPARC 1.3 and later versions, you can migrate with any number of MAUs.
- **Tagged and Untagged Networks**: When you migrate logical domains between servers that have different tagging mode, then the network configuration of the logical domain OS can be lost. You can avoid this issue by migrating domains between servers that have similar network configuration mode. When a guest domain is connected to a network configured with a VLAN ID, migration is not allowed between servers with the same VLAN network connected using different tagging mode.

Migrating a Logical Domain

1. Select the logical domain to be migrated from the Assets section.
2. Click **Migrate** in the Actions pane.

   The Migrate Logical Domain Wizard is displayed.
3. The eligible Oracle VM Servers appear in decreasing order of preference. Select an Oracle VM Server from the list.

   Click **Next**.

4. Review the information and click **Finish** to migrate the logical domain.

This migration is initiated from the logical domain and it results in migration of single logical domain. When the Oracle VM Server is placed in a server pool and there are more than one logical domain in the host, you can use Migrate Logical Domains to migrate more than one logical domain from the host.

Migrating Multiple Logical Domains

1. Select an Oracle VM Server that is placed in a server pool.
2. Click **Migrate Logical Domains** in the Actions pane.

   The Migrate Logical Domains Wizard is displayed.
3. The list of logical domains running in the Oracle VM Server are listed. Select one or more logical domains from the list.

   Click **Next**.
4. The eligible Oracle VM Servers in the same server pool that have the required resources to host all the logical domains are listed in the decreasing order of preference. Select an Oracle VM Server from the list.

   Click **Next**.
5. Review the summary and click **Finish** to migrate logical domains.

Automatic Recovery of Logical Domains

The following information is covered in this section:

- Virtual Disk Multipathing and Automatic Recovery Process
- Recovering Logical Domains
- Re-introducing the Failed Server

You can enable automatic recovery of logical domains. This ensures that when the Oracle VM Server for SPARC placed in a server pool fails, the logical domains are recovered and restarted on another server in the server pool.

You can set the priority of recovery value for the logical domain that decides the order of recovery of the logical domains in the Oracle VM Server for SPARC. The priority of recovery value can be set between 0 to 100. For example, if the priority of recovery value set for logical domain A is 12 and logical domain B is 15, then the logical domain B is recovered first.

Use the **Enable Automatic Recovery** and **Disable Automatic Recovery** actions, shown in Figure 19–15, to set the automatic recovery for a logical domain. If you have disabled the automatic recovery of a logical domain, then it is not migrated to other servers in the pool when the underlying virtualization host fails. Instead, it is listed under Shutdown Guests in the Server Pool. You can restart them on other servers as required.

Figure 19–15 Enable and Disable Automatic Recovery

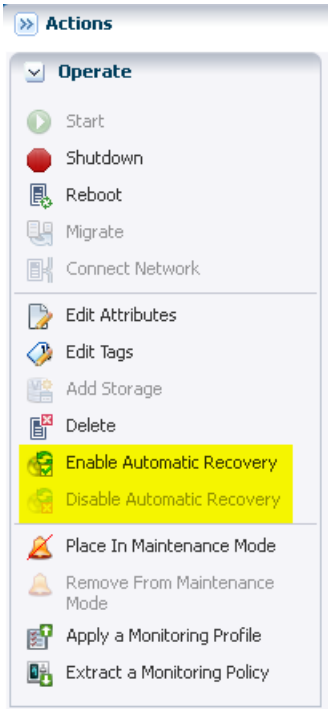

When you edit the attributes of a logical domain, you can edit the priority of recovery value.

If there are no resources available to recover the logical domains in the server pool, Oracle Enterprise Manager Ops Center checks periodically for every minute for free resources to retry the automatic recovery mechanism.

You can also follow the procedure described in Recovering Logical Domains to recover the attached logical domains.

When a logical domain is configured for automatic recovery, the auto boot value for the logical domain is set to false. The auto boot value is controlled by Oracle Enterprise Manager Ops Center. This is because, when the failed server is repaired and restarted, the logical domains are not started automatically. Oracle Enterprise Manager Ops Center checks whether the logical domains are recovered and running on other servers. If the logical domains are not recovered, then the logical domains are started. If the logical domains are recovered on other servers, then Oracle Enterprise Manager Ops Center cleans up those logical domains on the server.

When you perform shut down and start operations using the UI for a logical domain that is configured to automatically recover, the operating system is also booted automatically by Oracle Enterprise Manager Ops Center.

When you perform the shut down and start operations using the CLI for a logical domain that is configured to automatically recover, you must boot the OS separately. This is not automatically done by Oracle Enterprise Manager Ops Center.

See Chapter 21, "Server Pools" and Automatic Recovery of Oracle VM Server for SPARC in a server pool for more information.

Virtual Disk Multipathing and Automatic Recovery Process

When you have I/O domains and root domains, it is possible to have several paths to access the LUNs allocated for virtual disk storage. The redundant storage access

requires the automatic recovery options of the logical domains to be set or modified accordingly to recover the logical domains with or without the redundant storage.

For example, if an Oracle VM Server for SPARC has two HBA cards that have access to the Fiber Channel LUNs. One HBA card is assigned to the control domain or the primary domain. Other HBA card is made available to an I/O domain by assigning the PCIe bus to it. Now, the LUN from the storage array is accessible from the primary domain and the I/O domain. When you create a guest domain, select that the virtual disk is created from the virtual disk server of the primary domain and the alternate path to the storage is the virtual disk server on the I/O domain. A multipathing group is created for the virtual disk of the guest domain with both the paths to the storage. Also, while creating the guest domain select the option to automatic recovery and authorize recovery without Redundant I/O.

If the PCIe card on the primary domain fails, then the guest domain is still up and running as it has access to its virtual disk through the alternate path on the I/O domain.

The Oracle VM Server for SPARC system is placed in a server pool with other servers in the pool that do not have an I/O domain or root domain. The primary domain is the only access to the LUN.

If the server fails, then Oracle Enterprise Manager Ops Center tries to recover the guest domain on other servers in the server pool only when the option Authorize Recovery without Redundant I/O is selected. Otherwise, the recovery is not done as there is no target server that provides redundant access to the storage.

Recovering Logical Domains

The following procedure describes the actions that must be performed to recover the attached logical domains when an Oracle VM Server for SPARC server in a server pool has failed.

To Recover Logical Domains

1. Isolate the failed server.

   Log in to the ALOM or ILOM of the physical server and shut down the server.

2. Power-off the failed server.

3. In the Oracle Enterprise Manager Ops Center UI, select the server pool in which the failed server is the member. Check whether the server is flagged as unavailable. This status is updated within 5 minutes approximately.

4. Select the control domain of the server in the UI and click **Delete Asset** in the Actions pane.

5. Select the server pool in which the failed server was a member and expand the **Shutdown Guests** list.

6. The attached logical domains of the failed server are listed. Select the logical domains and click **Start** to start the logical domains in the desired server in the server pool.

7. To bring back the failed server to the server pool, the server must be repaired and provisioned again with Oracle VM Server for SPARC using Oracle Enterprise Manager Ops Center.

   To avoid provisioning the server again with Oracle VM Server for SPARC, follow the steps as defined in the section Re-introducing the Failed Server.

Re-introducing the Failed Server

You can bring in the repaired server back into Oracle Enterprise Manager Ops Center without the need of provisioning Oracle VM Server for SPARC again.

You must do the following when you want to re-introduce the repaired Oracle VM Server back into Oracle Enterprise Manager Ops Center:

1. Start the server and login to the console of the server and check whether the logical domains that have been recovered on other servers are still present on the server without the OS being booted. The logical domain OS not being booted prevents data corruption.
2. During the startup of the Oracle Enterprise Manager Ops Center Agent Controller on the server, the logical domain that were recovered on other servers without their OS being booted are removed.
3. Wait until the agent starts up on the control domain. When the agent is starting up, the following message is displayed on executing the command **`/usr/bin/svcs -xv`**:

   ```
   svc:/application/management/common-agent-container-1:scn-agent is
   starting
   ```

   This message stops when the agent is started.
4. Discover the control domain from the Oracle Enterprise Manager Ops Center UI using the discovery procedure as explained in the section Discovering Existing Oracle VM Server for SPARC Environments.

Layered Virtualization

Oracle Enterprise Manager Ops Center now supports non-global zones within the logical domains. The consolidated view of the control domain, and its service domains helps Oracle Enterprise Manager Ops Center to understand the network and storage resources that can be used to create zones within a logical domain OS. When you provision OS on logical domains, it will automatically install the Zone Virtualization Controller Agent. The Zone Virtualization Controller Agent enables you to create non-global zones on the logical domain OS.

Connecting networks to a logical domain results in creation of alternate MAC address which can be used to allocate to the zones.

The virtual disks storage that are free and available for use in the logical domains are available for zone storage. The storage that are in use by other zones, and those exported by the logical domains and used as a virtual disk to other domains, are not available for zone creation.

When you need to migrate the zones on the logical domains placed in a server pool, then the zone must not be using the logical domains virtual disk for storage. To enable migration of zones created on logical domains, the following conditions must be satisfied:

- The zone must be directly attached to NFS, FC LUNs or iSCSI LUNs.
- For zones on I/O domains or root domains and to use FC LUNs, the domains must have the FC card to enable creation and migration of zones.

Refer to the Chapter 18, "Oracle Solaris Zones," for more information about creating zones.

Server Pools

You can create server pools of your Oracle VM Server for SPARC systems in Oracle Enterprise Manager Ops Center. The Oracle VM Server for SPARC can have logical domains running in it.

When you want to create logical domains for the Oracle VM Servers placed in the server pool, the designation of the domains and the virtual switches or virtual functions to provide network connection are auto assigned. You cannot specify the network connection resource assignments.

The detailed description and procedure for creating server pools are described in Chapter 21, "Server Pools".

Related Resources for Oracle VM Server for SPARC

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources.

- For Oracle VM Server for SPARC documentation, see `http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html`
- Chapter 2, "Asset Management"
- Chapter 21, "Server Pools"
- Chapter 17, "Networks for Virtualization"
- Chapter 16, "Storage Libraries for Virtualization"
- *Oracle Enterprise Manager Ops Center Administration Guide*
- *Oracle Enterprise Manager Ops Center Configuring and Deploying Oracle VM Server for SPARC*

20

Oracle VM Server for x86

The following information includes:

Introduction to Oracle VM Server for x86

Oracle VM Server for x86 is a platform that provides a fully equipped environment to leverage the benefits of x86 virtualization technology. Oracle VM Server enables you to deploy operating systems and application software within a supported virtualization environment.

Oracle VM Server for x86 is a Xen-based server virtualization technology which supports Linux, Oracle Solaris, and Windows guests, and provide features to manage guests lifecycle operations, allocate and monitor guest resource consumption. The components of Oracle VM are:

- **Oracle VM Manager**: Provides the user interface to manage Oracle VM Servers, virtual machines, and resources.
- **Oracle VM Server**: A self-contained environment that is designed to produce a lightweight, secure server-based platform for running virtual machines. The Oracle VM Server can perform one or more of the following functions:

- **Server Pool Master**: The master is the core of the server pool operations and it acts as the contact point for the server pool to Oracle VM Manager, and also as the dispatcher to other Oracle VM Servers in the server pool.
- **Utility Server**: Its function focuses on the creation and removal operations of virtual machines, Oracle VM Servers, and server pools.
- **Virtual Machine Server**: The primary function of virtual machine server is to run virtual machines.

Roles for Oracle VM Server for x86

Table 20–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 20–1 Oracle VM Server for x86 Tasks and Roles

Task	Role
Discover Oracle VM Manager and Oracle VM Server	Virtualization admin
Provision Oracle VM Manager	Ops Center admin
Manage Oracle VM Manager	Virtualization admin
Create Link Aggregation	Network admin
Create, manage, update, and delete virtual machines	Virtualization admin
Set monitor threshold	Asset admin
Manage storage resources	Storage admin
Create and manage profiles and plans	Profile and plan admin

Integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center

The integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center provides the platform to manage Oracle VM Manager, Oracle VM Servers, server pools, and the virtual machines through Oracle Enterprise Manager Ops Center UI.

The following operations can be done through Oracle Enterprise Manager Ops Center to manage Oracle VM Server for x86 deployments:

- Discover deployed Oracle VM Managers
- Provision Oracle VM Servers
- Discover existing Oracle VM Servers
- Launch Oracle VM Manager UI
- Create virtual machines
- Provision OS on virtual machines
- Create server pools
- Connect to Oracle VM Manager console

- Manage storage repositories of Oracle VM Server for x86
- Perform management operations on Oracle VM Servers and virtual machines

Location of Oracle VM Server for x86 Information in the User Interface

Table 20–2 lists where to find different information for Oracle VM Server for x86 in the UI.

Table 20–2 Location of Oracle VM Server for x86 Information

To See	Location
Discovered Oracle VM Manager	Expand Administration in the Navigation pane. All the discovered Oracle VM Managers are displayed under Oracle VM Manager.
Discovered Oracle VM Server	Expand Assets in the Navigation pane. All the discovered Oracle VM Servers are displayed in the All Assets tree.
Server Pools	Expand Assets in the Navigation pane and select Server Pools in the Resource Management Views
Virtual Machine	Expand Assets in the Navigation pane and select Server Pools in the Resource Management Views. The virtual machines are listed under corresponding Oracle VM Servers.
Options for managing Oracle VM Servers from Oracle VM Manager	Expand Administration in the Navigation pane and select the Oracle VM Manager. All options to manage Oracle VM Servers are listed in the Actions pane.

Installing Oracle VM Manager

Installation of Oracle VM Manager is outside the scope of this document. For detailed instructions about how to install the Oracle VM Manager, see *Oracle VM Installation and Upgrade Guide for x86* at http://www.oracle.com/technetwork/documentation/vm-096300.html.

By default, the Enterprise Controller, Proxy Controller, and Agent Controller use the most recent version of JDK available on the system. However, Oracle VM Server for x86 code will not work when the JDK version on the Proxy Controller is higher than JDK version 6. If the Proxy Controller has a higher version, you must manually set the JDK to version 6 on the Proxy Controller by editing the JAVA_HOME variable to force the use of JDK 6.

Discover the Oracle VM Manager in Oracle Enterprise Manager Ops Center and launch the console to access it.

Discovering Oracle VM Manager

When you discover the Oracle VM Manager in Oracle Enterprise Manager Ops Center, the following Oracle VM resources are automatically discovered and populated in the UI:

- Oracle VM Servers
- Server pools
- Virtual machines
- Storage servers
- Storage repositories

- Networks

To discover an existing Oracle VM Manager, you must have the following information:

- Discovery profile for Oracle VM Manager
- Credentials for Oracle VM Manager access

You must create a discovery profile for discovering the installed Oracle VM Manager. The discovery profile provides options to create the discovery profiles for all the components related to Oracle VM Server for x86.

Figure 20–1 shows the different assets that are available under Oracle VM in the Discovery Profile Wizard.

Figure 20–1 Discovery Asset Type

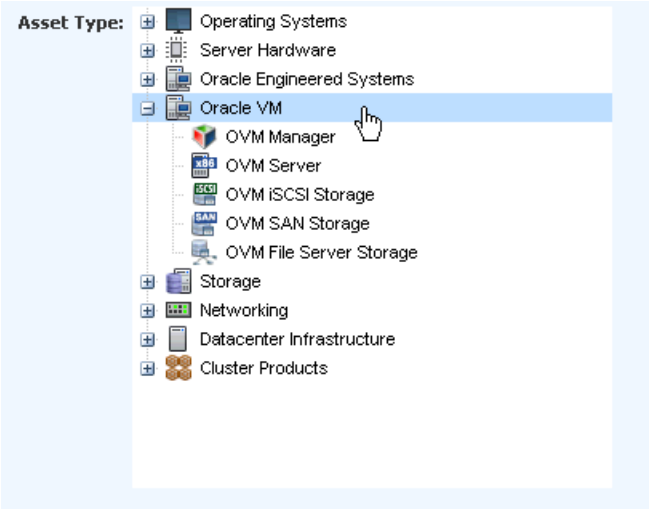

To Create a Discovery Profile for Oracle VM Manager

The following procedure outlines the steps to create a discovery profile for Oracle VM Manager:

1. Select Discovery from the Plan Management section.
2. Click **Create Profile** in the Actions pane.

   The Create Profile-Discovery Wizard is displayed.
3. Enter a name and description for the discovery profile.
4. Select the Asset Type as Oracle VM Manager from the list.

   Click **Next** to define the tags.
5. (Optional) Define the tags for the asset to be discovered. Tags help to group the assets.

   Click **Next** to enter the IP range.
6. (Optional) Enter the IP address of the Oracle VM Manager to be discovered. Enter the host name or IP Address when you execute the profile.

   Click **Next** to select the credentials of the Oracle VM Manager.
7. When you have the credentials of the Oracle VM Manager, click **Select** and the list of saved credentials are displayed. Select the appropriate credential from the list.

Click **New** to display the Create Credentials window. Enter the following information:

- Name and description for the credential.
- Enter the username and password of the Oracle VM Manager. Re-enter the password to confirm.
- Select the protocol and enter the port number.

Click **Ok** to create the credential and click **Next** to review the summary of the Oracle VM Manager discovery parameters.

8. Click **Finish** to create the profile.

After creating the profile, execute the profile to discover the Oracle VM Manager. The following procedure describes how to execute the profile:

1. Select the discovery profile created for Oracle VM Manager.
2. Click the **Add Assets Using Profile** icon to launch the Add Assets Wizard.
3. When you have not provided the details of the IP address or host name in the profile, you must enter the details of the IP address or host name of the Oracle VM Manager.
4. Click **Add Now** to initiate the discovery job.

The Oracle VM Manager is considered an independent entity with which Oracle Enterprise Manager Ops Center integrates. The Oracle VM Manager appears in the Administration section of the UI, not as an asset in the Assets tree.

Figure 20–2 shows the Oracle VM Managers displayed in the Administration section of the UI.

Figure 20–2 Oracle VM Manager

When you discover an Oracle VM Manager, all of the Oracle VM Servers managed in it are also discovered and displayed in the Assets section. Existing server pools in the Oracle VM Manager are also displayed in the Server Pools view of the UI.

Discover Oracle VM Servers

When you discover and manage an Oracle VM Manager, all of the known Oracle VM Servers to the manager are also discovered and displayed in the software UI. When you have manually installed Oracle VM Servers and not discovered in the Oracle VM Manager, you can still discover the Oracle VM Server in the Oracle Enterprise Manager Ops Center UI.

When you create a discovery profile for Oracle VM Server, you must define the Oracle VM Manager to which it can be associated. This makes the Oracle VM Server known to the Oracle VM Manager.

You must use SSH discovery of the Oracle VM Server when you want access to the serial console and perform advanced configuration in the Control Domain.

Discovery Profile for Oracle VM Server

The creation of discovery profiles is similar to Oracle VM Manager except for selecting the Discovery Type as Oracle VM Server. When you provide the Oracle VM Server

host name or IP address, select the Oracle VM Manager with which you want to associate the Oracle VM Server.

Figure 20–3 shows the IP Ranges step of the Discovery Profile Wizard in which the user must select the Oracle VM Manager to discover the Oracle VM Server.

Figure 20–3 Select Oracle VM Manager

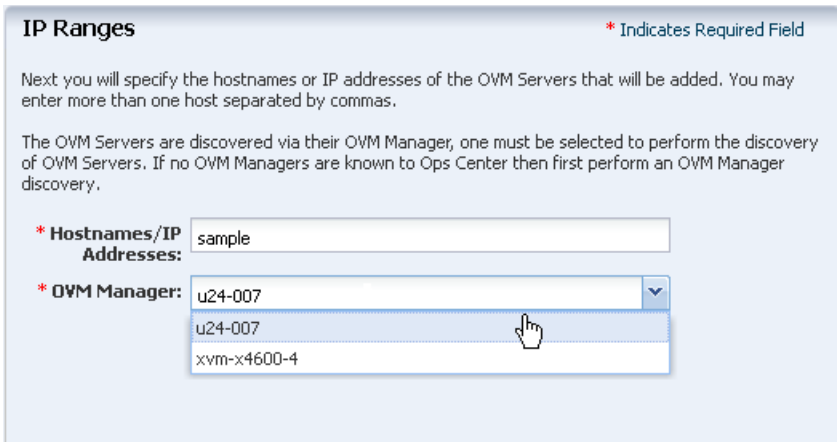

Execute the profile and discover the Oracle VM Server. This is a very basic discovery of the server. The storage and network configuration of the server are not populated in the UI and only basic information is displayed in the UI.

Discovering Oracle VM Servers Using SSH

Use this method when you want to access the serial console of the Oracle VM Server and perform advanced configuration.

To use SSH for discovery, define the credentials for the Oracle VM Server as explained in this procedure:

1. Select Credentials from the Plan Management section.
2. Click **Create Credentials** in the Actions pane.

   The Create Credentials window displays.
3. Select the protocol SSH from the list.
4. Provide a name and description for the credential.
5. Select **Password** for Authentication Type.
6. Enter the root user and the password.
7. Click **Ok** to save the credential.

Create a discovery profile for the Oracle VM Server and select the credential created for ssh discovery. Discover the Oracle VM Server using the profile. Since the discovery of the Oracle VM Server is through the SSH discovery, the network and storage configuration of the server is populated in the UI.

Installing Oracle VM Servers

You can provision Oracle VM Servers using Oracle Enterprise Manager Ops Center. You can download an Oracle VM Server image into the storage libraries, create OS provisioning profiles and then deploy a plan with the profile on the selected target.

When you import an image into Oracle Enterprise Manager Ops Center, a default profile and plan are created automatically. See Uploading or Importing Images to upload or import image into the software library.

Hardware and Software Requirements

Oracle VM Server is a managed virtualization environment based upon the Xen hypervisor technology and includes a small Linux-based management operating system. For hardware and software requirements for installing Oracle VM Server, see *Oracle VM Installation and Upgrade Guide for x86* at http://www.oracle.com/technetwork/documentation/vm-096300.html.

Oracle VM Agent Installation

When you install Oracle VM Server through Oracle Enterprise Manager Ops Center, the Oracle VM Server is installed without Agent Controller. The Oracle VM Agent is installed automatically when you install Oracle VM Server. The Oracle VM Agent enables communication between the Oracle VM Manager and Oracle VM Server for all management tasks.

Oracle VM Server Password Management

The provisioning profile of installing Oracle VM Server is similar to installing a Linux OS profile with additional steps to provide the Oracle VM Server password and the management interface. The Oracle VM Server password is different from the root password. This password is used by the Oracle VM Manager to manage and monitor the Oracle VM Server and the virtual machines running in it. You must use this password while discovering the Oracle VM Server from Oracle VM Manager.

Creating OS Provisioning Profile for Oracle VM Server

You can create OS provisioning profile for provisioning Oracle VM Servers.

To Create a Provisioning Profile for Oracle VM Server

1. Select OS Provisioning profile in the Plan Management section.
2. Click **Create Profile** to create a new profile.
3. Provide a name and description for the profile.
4. Select **Oracle VM Server for x86** in the subtype.

   Click **Next** to specify the OSP parameters.
5. Select the Oracle VM Server from the list.

   Oracle VM Server does not have any specific software groups. Click **Next** to specify the OS setup.

   (Optional) Select Include Custom Scripts to add some scripts to the installation. Select the scripts in the subsequent steps to execute.
6. Specify the time zone, language, terminal type and the root password for the OS. Also, specify the console baud rate and the serial port.

   Click **Next** to specify the installation parameters.
7. Specify the following parameters for the Linux OS:
   - **Installation number**: Enter the installation number that is used to allow installation of all of the software that is included in your subscription.

- **Partition action**: Select whether you want to change the disk partition of the system.
- **Install protocol**: Specify HTTP or NFS as the install protocol.
- **Kernel parameters**: If necessary, enter kernel parameters for the GRUB menu of the target system.
- **Linux packages**: Specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude a package, enter the package name preceded by a dash (-).

The options MD5 Checksum, Reboot action, Initialize Disk label, Use Shadow passwords, and Clear master boot record are enabled by default.

Click **Next** to specify the volume groups when you are using Logical Volume Manager.

8. Select **Use Logical Volume Manager** when you want to specify the volume groups in the Logical Volume Manager for the hard disk drives.

   Click **Next** to specify the disk partitions and file systems.

9. Specify the disk partitions and file systems that you want to create on the target system. For the file system device, you can also select the logical device name.

   Click **Next** to specify the naming service.

10. Specify the name service, domain name and the corresponding name server. Select one of the following name services:

    - **DNS**: Enter the domain name of the DNS server. Provide the IP address of the DNS server in the Name Server field. You can enter up to three IP addresses as the value for the Name Server. Provide the additional domains to search for name service information in the Domain Name Search List. You can specify up to six domain names to search. The total length of each search entry cannot exceed 250 characters.
    - **NIS or NIS+**: Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the option Specify an NIS Server and provide the NIS server host name and the IP address.
    - **LDAP**: Enter the domain name of the LDAP server. Specify the name of the LDAP Profile you want to use to configure the system. Enter the IP address of the LDAP Profile Server. You can also optionally provide the Proxy Bind Distinguished Name and Password.
    - **None**: Select None when there is no naming service configured.

    Click **Next** to specify the network.

11. Select **None** to continue to specify the networks.

    You can bond the interfaces after provisioning the Oracle VM Server. Select the Linux OS and aggregate the links.

12. Click the Add icon to add multiple networks.

    All of the networks that are defined in Oracle Enterprise Manager Ops Center are displayed in the Network list. Enter the following information for each selected network:

    - Select a NIC from the list of available logical interfaces for each network.
    - Select the Address Allocation Method for the selected networks except the boot interface.

When you have selected Use Static IP for Address Allocation Method, then you must provide the IP address when you apply a plan with this profile. The specific IP address is assigned to the target system after provisioning.

Click **Next** to specify the Oracle VM Server parameters.

13. Enter the password to be used by the Oracle VM Agent. This password is used by the Oracle VM Manager to manage and monitor the Oracle VM Server. Re-enter the password to confirm it.

    Figure 20–4 shows the step to enter the Oracle VM Server password.

Figure 20–4 Oracle VM Server Password

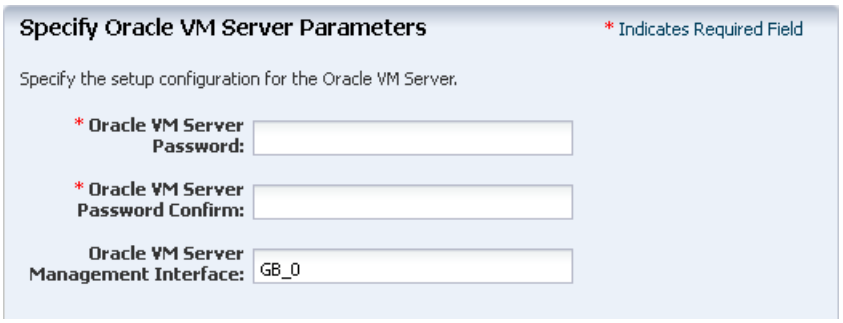

14. Specify the management interface, then click **Next**.
15. Review the information in the Summary, then click **Finish** to create the profile.

Applying Deployment Plan for Oracle VM Server

When you create a profile for provisioning Oracle VM Server, select the option to create a deployment plan with the profile. Either apply the created plan or create your own deployment plan.

To Apply a Deployment Plan for Oracle VM Server for x86

1. Select the plan created for provisioning Oracle VM Server in the Plan Management section.
2. Click **Apply Deployment Plan** in the Actions pane.
3. Specify the network resources defined in the profile. Provide the IP address for the boot interface.
4. (Optional) Provide a host name for the Oracle VM Server.

   Click **Next** to select the Oracle VM Manager.
5. (Optional) Select an Oracle VM Manager from the list to discover and manage the Oracle VM Server. Select **None** when you do not want to associate with an Oracle VM Manager.

   Place the Oracle VM Server in one of the server pools of the selected Oracle VM Manager. After provisioning, a separate job is initiated to add the Oracle VM Server to the server pool.

   Click **Next** to schedule the job.
6. Schedule the provisioning job to run now or at a later time.
7. Review the summary and click **Apply** to launch the provisioning job.

Administration of Oracle VM Manager

The list of Oracle VM Managers discovered and managed in Oracle Enterprise Manager Ops Center are displayed in the Administration section.

Select an Oracle VM Manager and the following information is displayed as shown in Figure 20–5.

Figure 20–5 Oracle VM Manager View

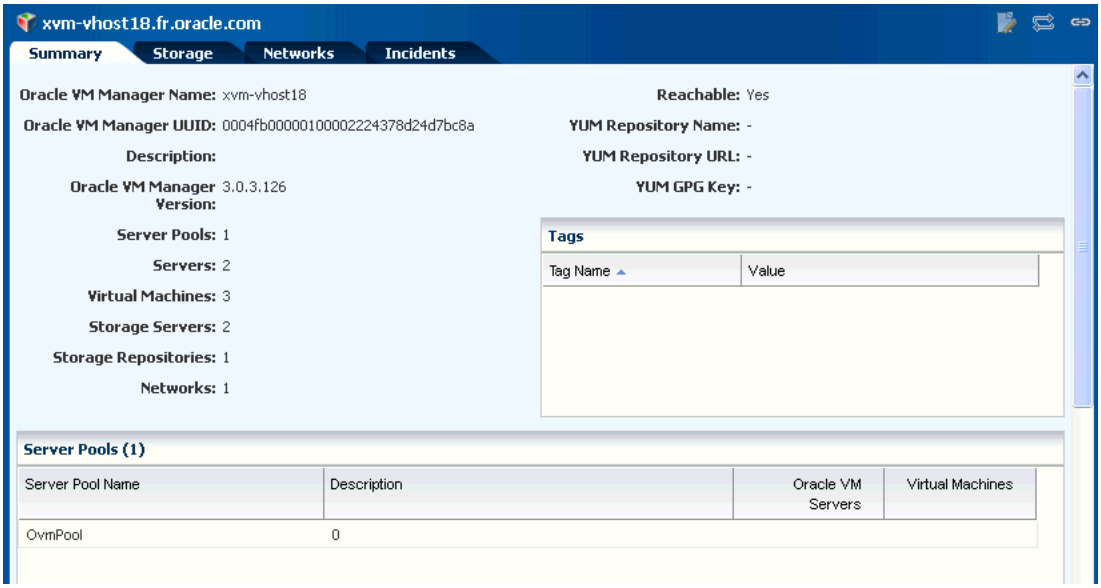

The Summary page displays information about the Oracle VM Manager, the list of server pools, the discovered Oracle VM Servers, and the virtual machines that are created in the corresponding Oracle VM Servers.

Manage Ownership of Oracle VM Servers

To use the Oracle VM Servers in the server pool and create virtual machines, it is required that you take ownership of the Oracle VM Server from the Oracle VM Manager.

Oracle VM Manager must take ownership of the Oracle VM Servers. Either assign ownership from the Oracle VM Server or take ownership through the Oracle VM Manager.

When you select an Oracle VM Server in the Assets section, the **Assign Ownership** and **Release Ownership** actions are available in the Actions pane.

Figure 20–6 shows the options Assign Ownership and Release Ownership displayed in the Actions pane.

Figure 20–6 Assign and Release Ownership From Oracle VM Server

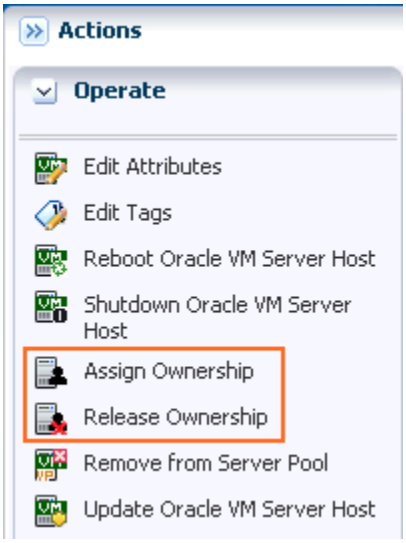

When you select the Oracle VM Manager from the **Administration** section, the **Take Ownership** and **Release Ownership** options are available in the Actions pane.

When you take ownership of the Oracle VM Servers, add them to server pools and create virtual machines in them.

Discover Storage Resources

When you discover an Oracle VM Manager, Oracle Enterprise Manager Ops Center provides the option to discover the storage servers attached to it.

The following discovery options are available from the Oracle VM Manager:

- File system storage
- iSCSI storage
- SAN storage

Figure 20–7 shows the options to discover the storage resources from Oracle VM Manager.

Figure 20–7 Discover Storage Resources from Oracle VM Manager

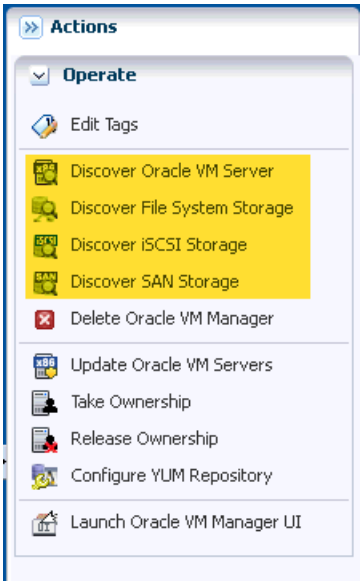

Create discovery profiles for the storage servers and save them.

When you select to discover a storage server, existing discovery profiles are searched for the selected storage type. You can select from an available profile. When there are no profiles available, then the wizard to discover the corresponding storage server appears.

When you discover a storage resource, existing resources such as virtual disks, templates, ISO images, and virtual machine metadata are also discovered and displayed.

See Chapter 16, "Storage Libraries for Virtualization" for more information about the storage servers and the discovering them in Oracle Enterprise Manager Ops Center.

Manage Storage Resources

The following type of libraries setup are supported for storage of Oracle VM resources:

- **Filesystem Storage Libraries**

  Oracle Enterprise Manager Ops Center provides an option to create Oracle VM Storage Repositories. A storage repository is a logical disk space made available through a file system on top of physical storage hardware. The supported types of file system are NFS and OCFS. When the storage repository is created on an NFS file server, it is a NFS based storage repository. The NFS file server consists of NFS file systems. When the storage repository is created on a LUN, it is a LUN-based repository. The OCFS file system is created on the storage server. Create storage repositories on these file servers to be used by Oracle VM Servers to store resources. The resources include virtual machine metadata, templates, assemblies, ISO images and virtual disks.

  To create Oracle VM storage repositories, you must have an Oracle VM Server discovered. Only NFS-based repositories can be shared by multiple server pools.

- **Static Block Storage Libraries**

  The LUNs from the storage arrays that are not managed by Oracle Enterprise Manager Ops Center form the static storage libraries. The LUNs are addressed by iSCSI or Fibre Channel protocols. This forms the iSCSI and SAN static storage

libraries. Also, add LUNs exported from managed storage arrays. The LUNs can belong to one library at a time.

- **Dynamic Block Storage Libraries**

  The storage servers that are discovered and managed in Oracle Enterprise Manager Ops Center are presented under dynamic block storage libraries. The dynamic block storage libraries contains the exported LUNs from the storage array servers.

See Chapter 16, "Storage Libraries for Virtualization" for more information about discovering your storage servers, storage connect plug-ins, creating storage libraries and managing the libraries.

Manage Networks

During the installation of Oracle VM Server, the network interface used for the management is configured as a bonded interface. The bond is created with one interface and named as `bond0`. You can create additional bonds to add redundancy and load balancing of your network environment.

Attach networks to the Oracle VM Server or to the server pool that consists of a group of Oracle VM Servers on an Oracle VM Manager. Configure the network interfaces or the bonds to the network to be attached. You can assign different roles or functions to the networks attached to the Oracle VM Server.

The following are the network roles available for an Oracle VM Server:

- **Server Management**: Manages the Oracle VM Servers in a server pool. The Oracle VM Manager has one Server Management network.
- **Live Migrate**: Migrates the virtual machines from one Oracle VM Server to another in the server pool, without changing the state of the virtual machine.
- **Cluster Heartbeat**: Verifies that the Oracle VM Servers in the server pool are running.
- **Virtual Machine**: Monitors the network traffic between the virtual machines in a server pool.
- **Storage**: Transfers between virtual machines and virtual disks.

The management network created during the installation of Oracle VM Server has the following roles:

- Server Management
- Cluster Heartbeat
- Live Migrate

You can add and remove the roles of this management network, except for the Server Management role.

Depending on the available network interfaces on the Oracle VM Server, you can attach networks to Oracle VM Server and assign different roles to the networks. For example, you can attach the network in which your storage servers are placed and assign the Storage role to that network. You can assign a network with Live Migrate to be used only for migration.

For more information about setting different roles for network, refer to the *Oracle VM User's Guide*.

Attaching Networks

You can attach networks to the Oracle VM Server using the option **Attach Network**. When the Oracle VM Server is placed in a server pool and the server pool is associated with an user-defined network domain, you must add the network to the network domain.

1. Select the Oracle VM Server or the server pool.
2. Select **Attach Networks** from the Actions pane.

   The Attach Network Wizard is displayed. The list of available networks in the network domain are displayed with the current connection to the Oracle VM Server.
3. Select one or more networks that you want to attach to the Oracle VM Server or the server pool.

   Note: You cannot make multiple connections to the same network. You can connect to a network only once.

   Click **Next** to configure the networks.
4. You cannot make more than one connection to the network. Do not increment the number of connections.

   Click **Next** to configure the interfaces to the network.
5. For each selected network, select the NIC and the IP address.
   - **NIC**: Select an available NIC from the list. When the networks have different VLAN IDs, then you can assign the same NIC to another network. Otherwise, you cannot assign the same NIC to different network connection. When supported by the network, you can also select System Allocated for the system to take care of the NIC allocation.
   - **IP Address**: Select Use Static IP to provide the IP address manually or select Use System Allocated IP Address for the system to take care of the IP address allocation.

   Click **Next** to view the summary of the selected information.
6. Review the information provided and click **Finish** to attach the networks to the Oracle VM Server or the server pool.

Manage Oracle VM Servers

When you install or discover the Oracle VM Server through the Oracle Enterprise Manager Ops Center UI, take ownership of the server from the Oracle VM Manager. This helps to place the Oracle VM Servers in the server pool. The following are some functions available for managing Oracle VM Servers:

- Edit Oracle VM Server information
- Reboot the server
- Shutdown the server
- Update the server
- Edit IPMI configuration

Edit Oracle VM Server Information

Select the Oracle VM Server and click **Edit Attributes** in the Actions pane. You can edit the following Oracle VM Server information:

- Name
- Description
- Tags

Editing IPMI Configuration

Intelligent Platform Management Interface (IPMI) allows you to remotely power off an Oracle VM Server, and to send a Wake-on-LAN message to power on an Oracle VM Server without having to physically press the power button. You can edit the IPMI configuration in the following method:

To Edit IPMI Configuration

1. Select the Oracle VM Server and click **Edit IPMI Configuration** in the Actions pane.

   The Edit IPMI Configuration window is displayed as shown in Figure 20–8.

Figure 20–8 IPMI Configuration

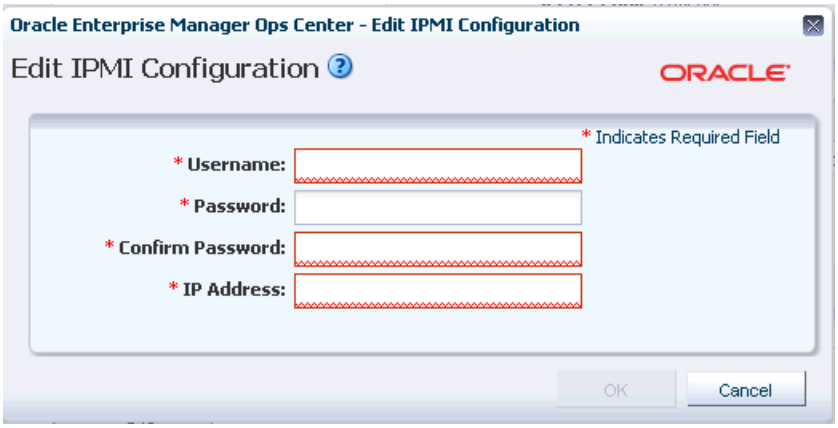

2. Enter the following information for IPMI configuration:
   - Username for the IPMI
   - Password for the IPMI
   - IP address of the IPMI
3. Click **Ok** to enter the IPMI configuration.

Placing Oracle VM Server in Maintenance Mode

To perform hardware or software maintenance, an Oracle VM Server can be placed in maintenance mode. When an Oracle VM Server is placed in maintenance mode, all of the virtual machines running on the Oracle VM Server are automatically migrated to other Oracle VM Servers in the server pool, if they are available, otherwise they are stopped. When the Oracle VM Server is the master Oracle VM Server in the server pool, this role is moved to another Oracle VM Server in the server pool, if available.

1. Select the Oracle VM Server and click **Place in Maintenance Mode** in the Actions pane.
2. Click **Place** to confirm the action.

   The Oracle VM Server is placed in maintenance mode.

When you have finished with the maintenance on the Oracle VM Server and you are ready for it to rejoin the server pool, select the option **Remove From Maintenance Mode**.

Updating Oracle VM Server

Using a YUM repository, update or upgrade Oracle VM Servers. To access patch updates for Oracle VM, contact Oracle to purchase an Oracle VM Support contract and gain access to the Unbreakable Linux Network (ULN) which contains updates for Oracle VM. When you have access to ULN, you can use this to set up your own Yum repository to use when updating your Oracle VM Servers.

Configure YUM Repository

A YUM repository is required to update or upgrade Oracle VM Servers that are attached to an Oracle VM Manager. Set up a YUM repository and configure for automatic updates of the Oracle VM Servers managed by an Oracle VM Manager. Setting up a YUM repository is beyond the scope of this document.

After you set up a YUM repository, select the Oracle VM Manager and click Configure YUM Repository in the Actions pane. Enter the following details of the repository to configure it.

- Yum Repository Name: A name for the Yum repository.
- YUM Base URL: The URL to access the Yum repository, for example, the Oracle public YUM repository is at:

  `http://public-yum.oracle.com/`
- Enable GPG Key: Select whether to use the GPG Key for the Yum repository. The GPG key or GnuPG key is the GNU project's implementation of the OpenPGP key management standard.
- Yum GPG Key: This field is enabled when you select Enable GPG Key. Enter the GPG key for the Yum repository, for example:

  `http://public-yum.oracle.com/RPM-GPG-KEY-oracle-e15`

  The GPG key must be available through HTTP, FTP, or HTTPS protocols. The GPG key for Oracle-signed updates from ULN is pre-installed on the Oracle VM Server at `/etc/pki/rpm-gpg/RPM-GPG-KEY-oracle`. When you want to use this GPG key, enter as:

  `file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle`

To Update Oracle VM Server

1. Configure the Yum repository.
2. Select the Oracle VM Server and click **Update Oracle VM Server Host** in the Actions pane.

When you update the Oracle VM Server, the server is placed into maintenance mode and then the update is performed. When there are virtual machines in the Oracle VM Server, they are migrated to other Oracle VM Servers in the server pool, if they are

available, otherwise they are stopped. After the update, the Oracle VM Server remains in maintenance mode. Select the option **Remove From Maintenance Mode**.

Create Server Pool

You can group one or more Oracle VM Servers in an Oracle VM Manager and create server pools. You must create server pools before you start creating virtual machines. You can apply the virtual machines plans only on Oracle VM Servers that are in a server pool.

This section describes only how to create server pools using Oracle Enterprise Manager Ops Center. For more information about server pool policies, and managing server pools, see Chapter 21, "Server Pools".

When you create a server pool for Oracle VM Server for x86 systems, you must have the following information defined:

- Cluster file system
- Server pool master
- Virtual IP address for the server pool master

Cluster File System

Shared access to the server pool resources is a must for providing for high availability for the virtual machines running in the Oracle VM Servers of the server pool. This is achieved by cluster file system OCFS2 which allows multiple Oracle VM Servers to access the same disk at the same time. OCFS2 ensures that the Oracle VM Servers in a server pool can access and modify resources in the shared repositories in a controlled manner.

When you create a server pool, you must specify the server pool file system for the cluster heartbeat and other cluster information. The file system can be NFS shares or LUNs of iSCSI or SAN-based storage servers. Oracle VM formats the server pool file system as OCFS2 file system.

In Oracle Enterprise Manager Ops Center, the cluster is always enabled by default. The cluster configuration is pushed out to all the Oracle VM Servers in the server pool and the cluster heartbeat starts when the server pool is created. You can set a separate network for this cluster heartbeat. See Manage Networks for more information about setting up the networks and its role for an Oracle VM Server.

Server Pool Master

An Oracle VM Server is internally elected as Server Pool Master. You cannot set the role for an Oracle VM Server. When the elected Oracle VM Server fails, the role is set for another Oracle VM Server in the server pool. The virtual IP address provided while creating the server pool is assigned to the Oracle VM Server that has been elected as server pool master.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for x86 are as follows:

- Set the CPU utilization threshold. Places the virtual machines on the Oracle VM Server that has the lowest relative load.
- Set the CPU utilization threshold. Places the virtual machines on minimum number of Oracle VM servers to minimize the power consumption.

- You can select the automatic load balancing so that the virtual machines can be migrated automatically in the server pool whenever the thresholds are exceeded.

  When you enable automatic load balancing for Oracle VM Server for x86 server pool, the server pool is checked continuously for the selected placement policy. When the threshold exceeds, the virtual machines are migrated from one Oracle VM Server to another.

  When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM Server host servers are freed from the logical domains and powered off to minimize power consumption.

  When the servers are overloaded and there are no other servers in the pool to host the virtual machines, then the policy decides to start a powered-off server using its Wake-on-LAN capability and live migrate the virtual machines to take up the load. The Wake-on-LAN capability must be enabled on the BIOS of the Oracle VM Server.

- When an Oracle VM Server fails, the virtual machines are migrated to another Oracle VM Server in the server pool, when you have selected Enable High Availability while creating the virtual machine.

Creating Oracle VM Server for x86 Server Pools

The following procedure describes the steps in the wizard to create an Oracle VM Server for x86 server pool.

1. Select Server Pool in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.

   The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:
   - Name and description of the server pool.
   - Enter tags for categorizing your server pool.
   - Select Oracle VM Server for x86 from the Virtualization Technology.

   Click **Next** to define the server pool configuration.
4. Define the following configuration details:
   - Select the Oracle VM Manager in which the Oracle VM Servers are discovered and owned.
   - Select the network domain.
   - Select a network from the list of networks available in the network domain.
   - Enter the virtual IP address that is assigned to the server pool master. When the server pool master changes, the IP address is assigned to the new Oracle VM Server.
   - Select the server pool file system to store cluster hearbeat and other cluster information. The file system can be either NFS shares on a NFS file servers or LUNs of SAN and iSCSI based storages. The NFS file servers and the storage servers are known and reachable to the selected Oracle VM Manager.

   Click **Next** to select the members of the pool.
5. Select one or more Oracle VM Servers to add to the server pool.

The list of Oracle VM Servers displayed have the following characteristics:

- Owned by the Oracle VM Manager. Refer to Manage Ownership of Oracle VM Servers to own an Oracle VM Server.
- Not associated with server pool.
- In a healthy state.
- Not placed in maintenance mode.
- Do not have virtual machines in running, shutdown or suspended state.

Select Oracle VM Servers that are compatible for migration of virtual machines. To migrate virtual machines within a server pool, the Oracle VM Servers systems must be identical in model.

Click **Next** to associate the network domain.

6. Select the network domain to associate with the server pool.
   - When you have selected default network domain, then go to Step 7.
   - For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. You cannot bond the interfaces in Oracle VM Server for x86. Bonding can be done in only individual OS of the Oracle VM Server.

   Click **Next** to select the networks and associate with the server pool.
7. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For an user-defined network domain, only the networks assigned to it are listed.

   Select the networks that you want to associate with all the servers in the server pool. You cannot make multiple connections to a network. Limit the number of connections to 1.

   Click **Next** to configure the interfaces of the servers.
8. For each network connection, specify the following details:
   - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, then no rows are displayed for that server.
   - You can assign the same NIC to different network connection when the networks have different VLAN IDs.
   - When supported by the network, you can select System Allocated for the NIC and IP address to be automatically allocated by the system.

   Click **Next** to associate the storage libraries.
9. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store virtual machine metadata, ISO images, and for virtual disks of virtual machine.

   You can select the following type of storage libraries:
   - **Filesystem Storage Libraries**: This includes the Oracle VM storage repositories. You must have at least one Oracle VM storage repository associated with the server pool.
   - **Block Storage Libraries**: This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of

storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

Click **Next** to select the server pool policies.

10. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
    - **Placement Policy**: This policy decides the preferred virtualization host in the server pool to place the logical domains.
    - **Auto Balancing Policy**: This policy performs load balancing of the server pool automatically at set intervals.
11. Select one of the following placement policies:
    - **Lowest relative load**: The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. Based on this, the server with the lowest relative load is considered to place a virtual machine. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and virtual machines are migrated to server with lowest relative load.
    - **Minimize power consumption**: The virtual machines running in a server pool are placed on the minimum number of Oracle VM Servers so that the unused Oracle VM Servers can be powered off. The threshold value set for the server over-utilization ensures that the servers are not over-loaded. Otherwise, the powered off Oracle VM Servers can be powered on to host the virtual machines.
12. Select the automatic load balancing policies:
    - Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the virtual machines automatically. Or send notifications about the approval.
    - You cannot manually balance the resources.

    Click **Next** to view the summary of the server pool details.
13. Review the information to create a server pool for Oracle VM Server for x86. Click **Finish** to create the server pool.

Create Virtual Machines

A virtual machine comprises configurable set of resources and its own operating system. The resources include virtual CPU, memory, network, and virtual disk. You can start, stop, and restart each virtual machine independently. You can create virtual machines using one of the following installation sources:

- Templates
- Existing virtual machines
- ISO images in storage libraries
- Mounted ISO images on NFS, HTTP or FTP server
- Netboot

The operating system provisioned on the virtual machines can be hardware virtualized, paravirtualized and, hardware virtualized with paravirtualized drivers. Based on these virtualization types, you select the appropriate option to create the virtual machines.

Note: You can create virtual machines only in an Oracle VM Server for x86 server pools.

Virtualization Types

The following virtualization types are available for a virtual machine:

- **HVM**

  Hardware virtualization or fully virtualized. For HVM, select an ISO image from the storage library and create a virtual machine. You must activate the hardware virtualization in the BIOS of the Oracle VM Server in which you want to create the virtual machine. The OS might be completely unmodified.

- **PVM**

  Paravirtualized. You must use an ISO file mounted from an NFS share, HTTP, or FTP server. The OS is modified and recompiled to be made aware of the virtual environment. The paravirtualized guests run at near native speed, since most memory, disk and network accesses are optimized for maximum performance.

- **PVHVM or HVM with PV Drivers**

  Hardware assisted virtualization with a paravirtualized driver. Install paravirtualized drivers on the hardware virtualized machines for optimized performance. This is mainly used for running Microsoft Windows guest operating systems.

Creating Virtual Machine Profile

Capture the virtual machine requirements in the form of profile in Oracle Enterprise Manager Ops Center. Create a deployment plan with this profile and apply it on the selected targets. Use these profiles and plans to create more than one virtual machine at a time.

To Create a Virtual Machine

1. In the Plan Management section, expand Profiles and Policies and select Virtual Machine from the list.
2. Click **Create Profile** in the Actions pane.

   The Create Profile - Virtual Machine Wizard is displayed.
3. Enter the name and description of the profile.

   Select Create a deployment plan for this profile to automatically create a plan using this profile. Click **Next** to specify the installation source of the virtual machine.
4. Select an installation source for creating the virtual machine. You have the following options to select from:

   - **Templates**: Templates are pre-configured, pre-installed virtual machine. You must upload the templates into Oracle VM Storage Repositories. In the profile,

the available storage repositories and the corresponding available templates are listed. Select a storage library and template.

- **Existing Virtual Machine**: When there are existing virtual machines, select the storage repository and the virtual machine image stored in it.
- **Install from ISO image (HVM)**: You must have imported ISO images into the storage repository. Select the library and an ISO image from the list.
- **Install from mounted ISO image (PVM)**: Mount the required ISO files to a NFS share, HTTP or FTP server and it must be accessible from the Enterprise Controller. Provide the network path of the mounted ISO file. For example, an ISO image accessed through HTTP server is represented like this.

  ```
  http://example.com/Enterprise-R5-U6-Server-x86_64-dvd.iso/
  ```

- **Install from Network (Netboot)**: Use this option only to create the virtual machine with Network as the boot order. You must provision the OS separately on it.

Select an option for the installation source and click **Next** to identify the virtual machines created.

5. Provide a name for the virtual machine and the number to be appended to the name. You can create one or more virtual machines simultaneously using this profile. To provide unique name for the virtual machines, enter the prefix name and the number to start the series.

   For example, if the prefix name is VMachine, the number to start from is 10 and the number of virtual machines to be created is 3, then the virtual machines are created with the name VMachine10, VMachine11, and VMachine12.

6. Provide the description and tags for the virtual machine. All the virtual machines created with this profile carry the same description and tags. Use the tags to group the virtual machines based on the role.

   Click **Next** to specify the configuration of the virtual machine.

7. When you have the installation source selected as templates or existing virtual machine, the values for all the parameters are set from the template. Otherwise, you must provide the following parameters:

   - **VM Type**: Based on the installation source, select the type as HVM, PVM or HVM with PV Drivers. For HVM, you must have the ISO image in the storage library. For PVM, you must use mounted ISO images.
   - **OS**: Select the OS from the list to be installed on the virtual machine.
   - **Enable High Availability**: High availability is to ensure the uninterrupted availability of a virtual machine. When an Oracle VM Server shuts down or fails in a server pool, all the virtual machines are migrated to other Oracle VM Servers in the server pool. High availability is implemented only in server pool.
   - **CPU Threads**: The number of CPU threads to be used by the virtual machine.
   - **CPU Priority**: The CPU priority of the virtual machine. You can select a high (100), intermediate (50), or low (1) priority for the virtual CPUs. The higher the priority more number of physical CPU cycles are allocated to the virtual machine.

- **CPU Cap**: This parameter defines the maximum percentage to which the virtual CPUs can receive scheduled time. Set this option to keep the low priority virtual machines from using too many CPU cycles.
- **Memory**: The size of the memory the virtual machine is allocated to use when starting it.
- **Boot Property**: Select whether you want to start the virtual machine after creation.

Set these parameters and click **Next** to specify the boot configuration.

8. Select the boot media and order for the virtual machine. The boot options are set according to the installation source.

   The default boot order for different installation source are:

   - Templates – Disk
   - Existing Virtual Machines – Disk
   - HVM using ISO image – CDROM and Disk
   - PVM using mounted ISO image – Network and Disk
   - Netboot – Network

   Select the correct order when you want to modify the boot order. Click **Next** to select the storage resources for virtual machine.

9. Select the storage resources to store the virtual machine configuration, also known as virtual machine metadata, and the virtual disks.

   Use the following libraries for the virtual machine:

   - **Oracle VM Storage Repositories**: The Oracle VM storage repositories are NFS file servers that are used to store virtual machine metadata, ISO images and virtual disk images. You can store the virtual machine metadata only in Oracle VM storage repositories. When you use the storage repository for virtual disk storage, you can either create a virtual disk or select existing virtual disks that are unshared and unused. You can also select ISO images to a virtual machine for HVM and PVHVM types.
   - **Static Block Storage Libraries**: These are the storage array servers that are not managed by Oracle Enterprise Manager Ops Center but reachable to the Oracle VM Servers in the Oracle VM Manager. Select the LUNs that are exported as iSCSI and SAN storage types.
   - **Dynamic Block Storage Libraries**: The exported LUNs that are available from the storage servers that are managed by Oracle Enterprise Manager Ops Center. Select the existing LUNs that are available from the storage servers. You can also create LUNs by specifying the volume group and the size of the LUN.

10. When you select to create virtual machines from existing virtual machines or templates, you are required to select the clone type for the virtual disks. You can select the following clone types:

    - **Sparse type**: This clone type is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size.
    - **Non-sparse type**: This clone type is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

Select the appropriate storage library, specify the virtual disks, disk size, and click Next to select the networks.

11. Select the network domain, and the networks that must be attached to the virtual machine.

    You can connect to one or more networks. The number of connection to a network is limited to 1.

    Click **Next** to view the summary.

12. Review the information provided to create a virtual machine profile. Click **Finish** to create the profile.

Create a deployment plan with this profile and apply the plan to create virtual machines. Except for the installation source Netboot, all the other source types results in creating a virtual machine with the OS installed. For virtual machine created with Netboot installation source, you must provision the OS on it separately using an OS provisioning plan.

The next section describes the different methods of deployment of virtual machine plans to create virtual machines.

Deploying Virtual Machine Plan

This section describes how to create virtual machines and provision OS in an Oracle VM Server using Oracle Enterprise Manager Ops Center.

The two new deployments plans for provisioning and creating multiple virtual machines are:

- Single step deployment plan to create only virtual machines
- Multi-step deployment plan to create virtual machine, provision OS, and install other applications.

In the single step deployment plan, the virtual machine can be created using Oracle VM templates or ISO images stored in Oracle VM repositories. Or only a virtual machine is created with ready to be booted using a separate OS provision plan.

In the complex deployment plan, you can configure and install the virtual machine, provision the OS, install software applications, apply monitoring profiles, and operational plans.

To Apply a Create Virtual Machine Plan

This is a single step virtual machine plan that creates a virtual machine and the OS is installed depending on the installation source selection. When you create HVM or PVHVM type virtual machines, select the ISO image to be used to create the virtual machine. You can select only one ISO image and multiple ISO files are not supported. While selecting the virtual disks, you have the option to select the ISO image. The following procedure describes how to apply the plan for creating a virtual machine:

1. Select the Virtual Machine plan in the Plan Management section.
2. Click **Apply Deployment Plan** in the Actions pane.
3. Select the targets on which you want to apply the plan. Click **Add Targets to List**.

   You can apply the plan only on the Oracle VM Server for x86 server pools. You cannot create virtual machines on Oracle VM Servers that are not part of a server pool.

4. Select to apply the plan with minimal interaction when you do not want to modify the parameters in the profile.

   Click **Next**.

5. You can modify the virtual machine identification, if required.

   Click **Next** to specify the storage resource assignments.

6. The storage resources from the profile are populated. When the storage libraries are not available for a selected target, it is flagged in red. Modify the storage resources accordingly.

7. For the selected storage library, define the virtual disks.

   - For Oracle VM Storage repositories, you can create new virtual disks, select existing disks, or select ISO images. Select ISO image option is available only for HVM and PVHVM type virtual machines. Provide a disk name, select an existing disk or an ISO image.
   - For Static Block Storage, select an available LUN from the list.
   - For Dynamic Block Storage, either select an available LUN from the list or create a new LUN. When you create a new LUN, select the volume group and enter the size of the LUN.

   Click **Next** to provide the network resources.

8. The network selected in the profile is populated. When the network is not available for a selected target, it is flagged in red. Modify the network resource.

9. Specify the network resource and the IP address for the virtual machine.

   Click **Next** to view the summary.

10. Review the information provided and click **Apply** to apply the plan on the selected targets.

Except for the virtual machine created with Netboot as the installation source, all the other virtual machines are installed and provisioned with the OS.

Use the complex plan Configure and Install Virtual Machines to install the virtual machine and provision the OS on it. See Chapter 8, "Plans and Profiles" for more information about complex plans.

Provisioning OS on Virtual Machines

You can install an OS on virtual machines created on the Oracle VM Server. Use OS provisioning profiles to install the OS. See Chapter 12, "Operating System Management" for more information about creating provisioning profiles and plans to install an OS.

Manage Virtual Machines

Oracle Enterprise Manager Ops Center provides options to manage the lifecycle operations of virtual machines. You can start, suspend, resume, shut down, and delete virtual machines on the UI.

When you select the Oracle VM Server in the Assets section, the Summary tab in the center pane lists all the virtual machines that are currently running. All the options required to manage a virtual machine is available on the UI.

Figure 20–9 shows the actions available for a virtual machine from Oracle VM Server.

Figure 20–9 Action Icons for Virtual Machines

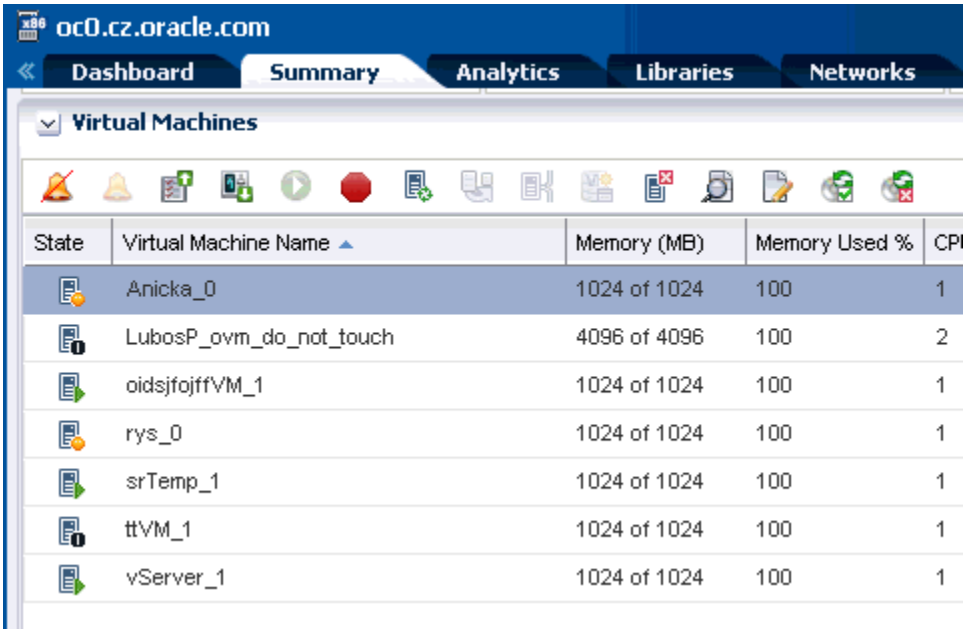

Starting a Virtual Machine

Select the virtual machine and click the **Start Selected Guest** icon. The job is initiated to start the virtual machine.

Suspending a Virtual Machine

Select the virtual machine and click the **Suspend Selected Guest** icon. When you suspend a virtual machine, you cannot use the virtual machine. Click **Ok** to confirm the suspend action.

Rebooting a Virtual Machine

Select the virtual machine in the running state and click the **Reboot Selected Guest** icon. Click **Reboot Guest** to confirm the action.

Shutting Down a Virtual Machine

Select the virtual machine and click the **Shut Down Guest** icon. When you shut down a virtual machine, it is disconnected from its network. Click **Ok** to shut down the virtual machine.

Deleting a Virtual Machine

Select the virtual machine and click the **Delete Selected Guest** icon. Click **Ok** to confirm deleting the virtual machine. When you delete a virtual machine, all the references to the virtual machine such as the metadata and disk images are deleted from the system.

All these options are also available in the Actions pane when you select the virtual machine in the Assets section.

Edit Boot Order

The boot order that you have set during virtual machine installation is displayed. You can modify the boot order. You must specify the appropriate boot order depending on

the installation source of the virtual machine. the following list shows the preferred boot options for different installation source:

- Hardware Virtualized Machine (HVM) – CDROM and Disk
- Paravirtualized Virtual Machine (PVM) – Network and Disk
- Templates and existing virtual machines – Disk
- Netboot – Network

If the virtual machine is in running state, then the modified boot order takes effect on the next reboot. Ensure to select the correct boot order for the virtual machine.

The **Edit Boot Order** option is available in the Actions pane of a selected virtual machine. Figure 20–10 shows the Edit Boot Order window that is displayed to edit the boot order of the virtual machine.

Figure 20–10 Edit Boot Order Window

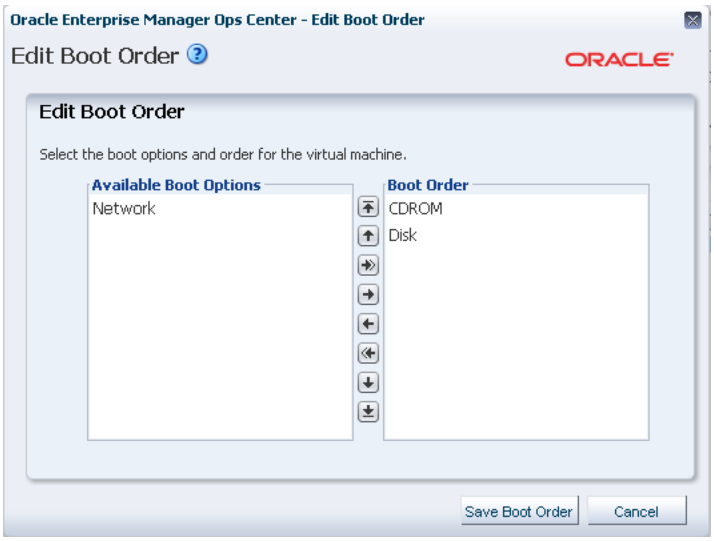

Modify the boot order and click Save Boot Order to save the changes.

Automatic Recovery

You can enable high availability for a virtual machine while creating a virtual machine profile. When the virtual machines are selected to be highly available and the Oracle VM Server shuts down or fails, the virtual machine is restarted on another available Oracle VM Server in the server pool.

You can also use the options **Enable Automatic Discovery** and **Disable Automatic Discovery** for a virtual machine to manage the high availability.

Related Resources for Oracle VM Server for x86

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- Chapter 2, "Asset Management"
- Chapter 21, "Server Pools"
- Chapter 17, "Networks for Virtualization"

- Chapter 16, "Storage Libraries for Virtualization"
- *Oracle Enterprise Manager Ops Center Administration Guide*

21

Server Pools

This chapter includes the following information:

Introduction to Server Pools

A server pool is a group of one or more virtualization hosts with the same processor architecture that have access to the same virtual and physical networks, and storage resources. Server pools provide load balancing, high availability capabilities, and sharing of some resources for all members of the pool.

You can create server pools for the following types of virtualization servers:

- Oracle VM Server for SPARC
- Oracle VM Server for x86
- Oracle Solaris Zones x86
- Oracle Solaris Zones SPARC

A virtualization host can refer to Oracle VM Server for SPARC, Oracle Solaris Zones, or Oracle VM Server for x86 that are managed through Oracle Enterprise Manager Ops Center. A virtual host or guest in a server pool refers to a non-global zone, logical domain or virtual machine running on the virtualization host.

Server pools are resource pools of homogeneous virtualization hosts that allow actions such as balancing load between servers and moving guests between hosts. You can also apply resource configurations and policies to them. The policies that you establish for a server pool manage many of the CPU utilization and resource balancing functions. Operations to the server pool are delegated to the individual virtualization hosts in the server pool.

To manage the guests within a server pool, you can perform warm and live guest migration and you can balance all of the guests' load among the members of the server pool. You can configure a policy to balance the load automatically, based on a schedule that you determine, or you can balance the load manually. When a virtualization host system shuts down, such as a hardware failure, you can start the guests on another host in the same pool.

Virtualization hosts in a server pool share network and storage libraries and several server pools can share the same networks and storage resources.

Server pools form a key concept for virtual datacenter management. Pooling your virtualization supported systems plays an important role and requires proper understanding of your requirements in your environment.

All the requirements, procedures, and different functions available for a server pool depend on the type of virtualization technology selected. See the appropriate sections to create a server pool for different virtualization technology.

- Oracle VM Server for SPARC Server Pool
- Oracle Solaris Zones Server Pool
- Oracle VM Server for x86 Server Pool

When you create a server pool, you must plan the storage and network resources such that they are accessible for all the members of the server pool.

You must set the policies to place the guests on the virtualization hosts and choose how you want to do the load balancing of the server pool. To understand more on this, refer to Server Pool Policies.

Roles for Server Pools

The following table lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See `User and Role Management` in the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 21–1 Server Pool Tasks and Roles

Task	Role
Create, manage, update, and delete server pools	Virtualization Admin
Provision and manage virtualization host	Virtualization Admin
Create, manage, update, and delete guests	Virtualization Admin

Actions for Server Pools

Using Oracle Enterprise Manager Ops Center, you can perform the following tasks:

- Create server pools for different types of virtualization technologies.

- Balance load on the resources in the server pool.
- Manage storage and network resources for the server pool.
- Set placement policies for automatic load balancing.
- Perform automatic recovery of the guests during server failure.
- Edit the attributes of an existing server pool.

Location of Server Pool Information in the User Interface

Table 21–2 lists where to find different information for server pools in the UI.

Table 21–2 Location of Server Pool Information in the UI

To See	Location
Server pools	Expand Assets in the Navigation pane and select Server Pools in the Resource Management Views.
Server pool resources	Expand Assets in the Navigation pane and select Server Pools in the Resource Management Views. Select the server pool and the center pane displays all the details of the server pool resources such as network, libraries, and fabrics.
Server pool actions	Expand Assets in the Navigation pane and select Server Pools in the Resource Management Views. Select the server pool and the Actions pane list the various options such as Associate Libraries, Attach Networks, Create Guests, Associate Network Domains and Edit Attributes for managing the server pool.

Server Pool Capabilities

You can create server pools for different types of virtualization technology. See the appropriate sections to create a server pool for different virtualization technology.

Each time Oracle Enterprise Manager Ops Center adds a new system to the pool, it checks the network connections on all the members of the pool. You will not be allowed to continue if something is wrong.

The following are some server pool capabilities and limitations:

- Mixing Virtualization Technology in a Server Pool
- Performing Maintenance When Using a Server Pool

Mixing Virtualization Technology in a Server Pool

You can create a zones server pool and use the logical domains as the global zones. However, those logical domains should not be part of an Oracle VM Server server pool. The user interface does not prevent you from creating these server pools, but the configuration is not supported and it can cause problems; particularly when automatic recovery is enabled on both server pools.

Performing Maintenance When Using a Server Pool

Before performing maintenance in your datacenter, it is a good idea to place the systems affected in maintenance mode to suppress the creation of alerts and incidents. When the system is part of a server pool, you can decide whether you want to migrate the guests to a different host in the pool before placing the system in maintenance

mode. Maintenance mode does not stop monitoring, it only stops the alerts and incidents.

Note: If you do not relocate the guests before putting the system in maintenance mode, consider disabling the server pool monitoring and disabling automatic recovery for each guest. Beginning with 12.2.2, you can disable automatic recovery at the server pool level instead of disabling it for each guest. See Editing Server Pool Parameters for how to disable automatic recovery for the server pool.

Server Pool Policies

The following information is covered in this section:

- Placement Policy
- Minimize Power Consumption Policy
- Automatic Load Balancing Policy

When you create a server pool, you define guest placement, and automatic load balancing polices. When you create guests, you define the guest resource consumption, including physical and virtual CPUs.

You can edit the server pool policies in the server pool's Summary tab. See Editing Server Pool Parameters for how to edit the pool attributes.

Placement Policy

The placement policy determines the preferred virtualization host for new guests within the server pool, the type of load balancing (automatic or manual) of the server pool, and balances the server pool during any server failure or maintenance. You define the placement policy when you create the server pool; however, you can modify the policy at any time.

The following are the server pool placement policy options:

- Place the guest on the virtualization host with the lowest relative load, based on the lowest CPU resource utilization. The calculation is based on a combination of the average load for the time period defined by the user. The default value is the last 10 minutes. The default threshold value for CPU utilization is 75%.
- Place the guest on the virtualization host with the lowest allocated CPU and memory, that is, the total static resource allocation across all guests on the host. The resource allocation is the sum of the number of vCPUs and virtual memory specified for each guest. This policy first verifies whether the resource allocation is possible and then ranks the server with available resources. The relative weight of CPU and memory resource is 1:1 while calculating the load on the host.
- Place the guest on minimum number of virtualization host and thereby consume less power.

The following is an example of how the placement policy works in a server pool.

Example 21–1 Example of Least Allocated Virtualization Host

A server pool has two virtualization hosts with different CPU and memory allocations. Host A has 8 available CPUs and 16 GB of memory. Host B has 4 available CPUs and 8 GB of memory.

Host A has three guests:

- Guest X has one vCPU and 1024 MB of memory.
- Guest Y has two vCPUs and 2048 MB of memory.
- Guest Z has one vCPU and 1024 MB of memory.

Host A's total static allocation is 4 vCPUs and 4 GB of memory.

Host B has one guest:

- Guest W has 3 vCPU and 4096 MB of memory.

Host B's total static allocation is 3 vCPUs and4 GB of memory.

The allocation percentage for Host A is: CPU allocation is 4 vCPUs/8 physical CPUs, or 50%. Memory allocation is 4 GB/16 GB or 25%.

The allocation percentage for Host B is: CPU allocation is 3 vCPUs/4 physical CPUs or 75%. Memory allocation is 4 GB/8 GB or 50%

Therefore, Host A is the least allocated of the virtualization hosts.

Minimize Power Consumption Policy

The policy to minimize power consumption places the guests on the minimum number of virtualization hosts and powers off the idle servers. The idle virtualization hosts are powered off or set to low-power mode on explicit approval.

For Oracle VM Server for x86, Oracle Enterprise Manager Ops Center will power off unused servers in the pool.

Beginning with Oracle Enterprise Manager Ops Center 12.2.2, you can enable power management for Oracle VM Server for SPARC server pools. When you enable power management, Oracle Enterprise Manager Ops Center will load as much work as it can onto a physical server before adding a virtual machine to a different server in the pool.

When power management is enabled, the software performs the following tasks:

- Places unassigned resources into a lower power state
- Identifies the members with the highest load and determines if those members have enough unallocated resources to support a new virtual machine placement request
- Checks the number of logical domains for each server in the pool and sets the appropriate placement policy
- Adjusts the placement policy every time you create, migrate, or delete a logical domain.

You can enable to policy and time interval to check for resource changes in Administration. The default check.delay time interval checks for resource changes every 5 minutes (300 seconds). See the *Oracle Enterprise Manager Ops Center Administration Guide* for how to set the power consumption policy and edit the property values.

Automatic Load Balancing Policy

Scheduling automatic load balancing is applicable only for Oracle VM Server for SPARC and Oracle Solaris Zones server pool.

Use the Automatic Load Balancing Policy to schedule load balancing within a server pool. You can schedule the automatic balancing to occur weekly, daily, or hourly on a

specific day and time of the week. The default is to balance the load on the servers in the server pool every Saturday at midnight according to the defined placement policy.

Note: The day and time are in the Enterprise Controller's time zone.

When you do not want to balance the server pool's load automatically, schedule a reminder to balance the server pool's load manually.

Maintenance Mode

Automatic load balancing policy does not consider the virtualization hosts that are placed in maintenance mode as a target in the server pool for migrating the guests. When you place a virtualization host in maintenance mode, all of the migratable guests in it are migrated to other servers in the server pool. This action is triggered automatically.

Resource Balancing Enhancements

You can define the resource usage threshold for the servers in the server pool. The servers are considered to be loaded heavily depending on the threshold set for the resources. Depending on the placement policy selected, the CPU or total CPU and memory resource usage is calculated.

When the virtualization host utilization exceeds the threshold, the software identifies the servers with sufficient resources and the guests that can be migrated from the server to balance the load.

When the virtualization hosts are well within the threshold and power minimization policy is selected, the automatic balancing policy attempts to free a virtualization host from its guest so that the administrator can power off the server.

Automatic Recovery

The automatic recovery feature provides the option to recover the guests that are attached to a failed server. The recovered guests are restarted on other servers in the server pool.

This feature is available for all types of supported virtualization technology. You can enable or disable the automatic recovery of zones, logical domains, and the virtual machines. The automatic recovery for virtual machines is managed by Oracle VM Manager. For logical domains and zones, the automatic recovery is managed by Oracle Enterprise Manager Ops Center.

Oracle Enterprise Manager Ops Center follows the recovery process described in this section, and is applicable for zones and logical domains.

By default, automatic recovery is enabled at the server pool level. When enabled at the server pool level, the recovery options selected for zones and logical domains determines whether automatic recovery is attempted.

Oracle Enterprise Manager Ops Center monitors the managed assets by checking the connection between the agent deployed on the assets and the Proxy Controller. You can set the interval of periodic check of the assets reachability while creating the server pool. The default value is 180 seconds. If any of the Oracle Enterprise Manager Ops Center Agent Controllers installed on the server do not respond within the fixed timeout interval, then the server is considered to be a failed server. The timeout interval is a minimum of 20 seconds and a maximum of 60 seconds.

For example, if the interval is set at 180 seconds, the server pool checks the status of the members of the pool every 180 seconds and if there is no response from the Agent Controller in 60 seconds, then the automatic recovery of the zones or the logical domains in the server are started almost after four (4) minutes after the server failure.

If the interval is set at 20 seconds, the server pool checks the status of the members of the pool every 20 seconds and if there is no response from the Agent Controller in 20 seconds, then the automatic recovery of the zones or the logical domains in the server are started almost after 40 seconds after the server failure.

An automatic recovery job is initiated when the global zone or Oracle VM Server for SPARC is:

- Placed in a server pool.
- Has at least a zone or logical domain that is configured to recover automatically.

The job initiated first performs the preliminary check of pinging all the known IP address of the failed server from another member in the server pool and the heartbeat in the metadata that is refreshed by the servers. The recovery job continues if the asset is seen as unreachable from Oracle Enterprise Manager Ops Center.

The recovery job fails if the assets are reachable and only the Agent Controller is down on the server. The attempt to recover the logical domains or zones are stopped as they are reachable.

If an automatic recovery failed or cannot be performed because of non availability of free resources in the server pool, Oracle Enterprise Manager Ops Center periodically checks if there are enough free resources and retries the automatic recovery.

The software check every 60 seconds until it can perform the automatic recovery. Beginning with Oracle Enterprise Manager Ops Center 12.2.2.0.0, a retry counter is available to control how many times the auto-recovery manager will attempt to perform an automatic recovery. Setting the counter to zero (0) results in an unlimited number of times that the automatic recovery is attempted.

Beginning with Oracle Enterprise Manager Ops Center 12.2.2.0.0, you can disable the automatic recovery feature for a server pool or for an individual asset. To disable the automatic recovery feature for a server pool, and all assets in the server pool, select the check box when you create the server pool. You can edit an existing pool to disable or enable the feature. When you disable automatic recover at the server pool level, the software will not check for resources or attempt to recover a guest, regardless of what option is selected at the zone or guest level.

You can enable to policy and time interval to check for resource changes in Administration. The default check.delay time interval checks for resource changes every 5 minutes (300 seconds). See the Administration guide for how to set the power consumption policy and edit the property values.

You can also manually recover the zones and logical domains. See Recovering Zones for more information about manual recovery.

See the sections Oracle Solaris Zones Server Pool, Oracle VM Server for SPARC Server Pool, and Oracle VM Server for x86 Server Pool for more information about automatic recovery options.

Server Pool Libraries

Oracle Enterprise Manager Ops Center uses software libraries to store ISO images and guest metadata. Storage libraries are also used to provide storage disks for the guest operating systems and for guests' data.

The virtualization hosts in a server pool share the libraries associated with the server pool. The type of library you can associate with the server pool depends on the type of virtualization hosts in the server pool and on whether all members of the server pool can access the storage resource.

See Chapter 16, "Storage Libraries for Virtualization" and Chapter 5, "Software Libraries" for more information about setting up these libraries.

Server Pool Networks

Oracle Enterprise Manager Ops Center provides group related networking components, such as fabrics, and networks, such as network domains. By default, all managed and declared networks in Oracle Enterprise Manager Ops Center are placed in the default network domain. Associating a virtualization host or a server pool with the default network domain does not require a fabric connection. All of the networks in the default network domain are available to be attached to the server pool.

You can create network domains. Associating a server pool with a user-defined network domain requires connection between the physical interface of each server in the pool and each fabric in the network domain. Only networks assigned to the user-defined network domain are available to be attached to the server pool.

Note: Before attaching one or more private ethernet networks to a server pool, verify that all members of the pool have access to the private networks. When all members do not have access, you might encounter inconsistent guest creation and network information.

You can use server pools with Oracle SuperCluster systems:

- When you attach public networks to a server pool, the server pool members can belong to different Oracle SuperCluster systems.
- When you attach private (internal) networks to an Oracle SuperCluster server pool, the members must all belong to the same Oracle SuperCluster system.

Note: Do not create server pools using members from more than one Oracle SuperCluster systems and attach private, or internal, networks to the systems.

Before creating a server pool, plan your networks and network domains that you must attach to the server pool. See Chapter 17, "Networks for Virtualization" for more information about creating network domains and managed networks.

Oracle VM Server for SPARC Server Pool

The following information is covered in this section:

- Enhancements in Oracle VM Server for SPARC Server Pool
- Server Pool Policies
- Automatic Recovery
- Creating an Oracle VM Server for SPARC Server Pool

You can pool the Oracle VM Server for SPARC resources and manage your logical domains. The Oracle VM Server for SPARC can be configured and running with I/O domains and root domains.

Oracle VM Servers for SPARC must meet the following requirements to be added to a server pool:

- You are able to place Oracle VM Servers of different CPU types in the server pool; however, you might lose the ability to migrate the logical domains between virtualization hosts with different CPU architectures. When you want to live migrate the logical domains, ensure that you create server pools of compatible CPU types. Cold migration and automatic guest recovery are available by using the shutdown/detach operation of the logical domain. Plan your servers that must be pooled.
- Must be running a supported version of Oracle VM Server for SPARC. The minimum supported version for server pools is Oracle VM Server for SPARC version 1.2.
- Must be running a supported version of Oracle Solaris operating system. The minimum supported version for server pools is Oracle Solaris 10 10/09 OS and must meet specific patch and firmware requirements.
- Plan the network that must be attached to the server pool. Either place those networks in the user-defined network domain or use the default network domain. You can connect to a network multiple times. The multiple network connection allows you to create IPMP or aggregate the links in the logical domain.
- The storage libraries must be associated with the server pool to store logical domain metadata, ISO images, and for virtual disk storage of logical domains. You can associate file system and block storage libraries with the server pool. Filesystem storage includes the NAS storage libraries. Block storage includes the SAN, iSCSI, and Dynamic Storage libraries. Ensure that you have at least one NAS storage library assigned to the server pool. For using migration capabilities, the server pool must be on shared storage facility.

Enhancements in Oracle VM Server for SPARC Server Pool

The following enhancements are available for supporting Oracle VM Server environments created outside Oracle Enterprise Manager Ops Center and advanced I/O domain configurations. Some of the enhancements are available from 3.0 or higher version:

- You can add Oracle VM Servers to the server pool even if they have logical domains configured and running.
- You can select the servers in the server pool to create I/O domains and root domains. The server pool will not be available as a target.
- Discover and manage user-configured Oracle VM Server for SPARC environments. The metadata of logical domains are stored in the local library of the Control Domain. To enable logical domain migration, you must the move the metadata to a NFS storage.
- Oracle VM Servers of different CPU type can be pooled together. The migration is not possible for logical domains between incompatible servers. Still, you can perform cold migration, that is shutting down the domain and starting on other servers in the server pool. Also, the guest recovery can be performed as it is often possible to start the domain on another server even though live migration is not possible.

- You can configure exclusive access of I/O domain and root domain resources to other logical domains. The exclusive access prevents creation of zones in the I/O domain or root domain OS and the global zone is not available to placed in a zones server pool.
- You can delete a server pool even if the logical domains are attached to the Oracle VM Servers.
- Placing an Oracle VM Server in the server pool does not remove any existing network and storage resource that are already attached to it.
- Deleting a server pool also does not unconfigure all the Oracle VM Servers of the attached network and storage resources. All the attached resources are retained and only the server pool is deleted.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for SPARC are as follows:

- **Placement Policy**

  Set the CPU utilization threshold. Place the logical domains on the Oracle VM Server for SPARC that has the lowest relative load.

  Set the CPU and memory allocation threshold. Place the logical domains on the Oracle VM Server for SPARC allocated with lowest CPU and memory resources.

  Set the CPU utilization threshold. Place the logical domains on the minimum number of Oracle VM Servers to minimize the power consumption.

- **Automatic Balancing**

  Set the automatic load balancing. Logical domains are migrated automatically in the server pool whenever the thresholds are exceeded. Schedule the automatic load balancing to take place at definite time interval.

  When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM Server host servers are freed up from the logical domains and powered off to minimize power consumption.

  Set whether administrator approval is required for migrating the logical domains.

Automatic Recovery

You can enable automatic recovery for a logical domain in the following scenarios:

- During logical domain creation, you can enable automatic recovery option and set the value for priority of recovery. See Creating a Guest Domain Profile for more information about setting the priority of recovery while creating a logical domain profile.
- For an existing logical domain, use the options Enable Automatic Recovery and Disable Automatic Recovery to set the automatic recovery. Also, you can edit the attributes of a logical domain to set the value for priority of recovery. See Automatic Recovery of Logical Domains for more information.
- See Automatic Recovery for the recovery process of the logical domains when a server fails in the server pool.

In an Oracle VM Server for SPARC server pool, the option to power-off the failed server and recover the logical domains is enabled by default. When a server fails,

Oracle Enterprise Manager Ops Center tries to power-off the failed server, ensure that it is stopped and then initiates the automatic recovery of the logical domains. The logical domains that are configured for automatic recovery are recovered and restarted on other servers in the server pool.

When the option to power-off the failed server is enabled and Oracle Enterprise Manager Ops Center cannot power-off the failed server, that is the Service Processor is also unreachable, the automatic recovery of the logical domains does not take place. This is to avoid any potential data corruption during the recovery process of the logical domains.

When Oracle Enterprise Manager Ops Center detects that the failed server is unreachable, the logical domains that were running in the server disappear from the Navigation pane in the UI until they are recovered to other servers in the server pool.

The metadata of disappeared logical domains is still available in the storage libraries. The logical domains that are not configured for automatic recovery or the logical domains that cannot be recovered re-appear in the UI in the following scenario:

- When the failed server is restarted, the logical domains are booted and attached to the server again.
- When you have followed the manual recovery procedure and deleted the failed server from Oracle Enterprise Manager Ops Center, the logical domains re-appear in the list of Shutdown Guests in the Server Pool. You can start them on other server in the server pool. See Recovering Logical Domains for the logical domains recovery.

Table 21–3 provides a quick view of different conditions that exist for recovering a logical domain.

Table 21–3 Recovery of Logical Domains

	Conditions	Result
Logical domains	▪ Failed Oracle VM Server for SPARC ▪ The server pool option to power-off the failed server on automatic recovery is enabled ▪ Oracle Enterprise Manager Ops Center can power-off the failed server ▪ Automatic recovery is enabled for the logical domain	The logical domain is restarted and placed in another server in the server pool. If there are no server available in the server pool, then you must manually recover the logical domains. See Recovering Logical Domains for more information.

Table 21–3 (Cont.) Recovery of Logical Domains

	Conditions	Result
Logical domains	■ Failed Oracle VM Server for SPARC ■ The server pool option to power-off the failed server on automatic recovery is enabled ■ Oracle Enterprise Manager Ops Center cannot power-off the failed server ■ Automatic recovery is enabled for the logical domain	The automatic recovery of the logical domain does not take place. This is to avoid any data corruption during the recovery procedure. You must manually recover the logical domains.
Logical domains	■ Failed Oracle VM Server for SPARC ■ The server pool option to power-off the failed server on automatic recovery is not enabled	Oracle Enterprise Manager Ops Center proceeds to recover the logical domains configured to automatically recover without trying to power-off the failed server from its processor. This can lead to data corruption if the failed server is just isolated from the network point of view, and has still logical domains writing to disks.
Logical domains	■ Failed Oracle VM Server for SPARC ■ Automatic recovery is not enabled for the logical domain	Follow the procedure in Recovering Logical Domains to recover the logical domains. Only after this the logical domains appear in the Shutdown Guests list in the UI. You can restart them on the available server in the server pool.

Network Tagging Mode Conditions

There are certain conditions in which the network tagging mode must be selected while creating a server pool. Review the following conditions and scenarios before deciding on the network tagging mode when you create an Oracle VM Server for SPARC server pool:

- You can select networks without VLAN ID. The UI does not provide the option to select Tagged or Untagged mode.
- You can select to associate and configure the networks with VLAN ID in Tagged mode.
- You can select to associate and configure the networks with VLAN ID in Untagged mode.
- You can select to configure the networks in mixed tagging mode in the server pool. For example, you can attach the network N1 with VLAN ID = 100 in tagged mode with the server S1 and in untagged mode for server S2. Refer to Mixed Network Tagging Mode Configurations in Server Pool for more detailed information.
- You can attach networks whose VLAN ID is similar to another network already connected to the servers. For example, a server S1 is already connected to network N1 with VLAN ID = 100, then while creating the server pool with S1 as the member of the pool, you can also attach a network N2 with VLAN ID =100.
- You can edit the VLAN ID of a network when you are attaching the network in Tagged mode for the first time.

- When you can edit the VLAN ID of the network, you cannot enter -1 as the value for the VLAN ID.
- If the selected network with a VLAN ID is already connected to the selected assets in Tagged mode, then you cannot edit the VLAN ID and make another connection.
- You cannot make multiple network connections to a member of the server pool over the same network in both tagged and untagged modes. The mode can be either in tagged or untagged mode only. For example, if you attach network N1 with VLAN ID =100 for the first time to server S1 in Tagged mode, then you cannot make another connection to the same network N1 in Untagged mode. Every other connection with network N1 must always be in Tagged mode for server S1.
- If the selected members of the pool are already connected to network N1 with VLAN ID =100, then you cannot select the same network with different VLAN ID to be connected for the server pool.

Mixed Network Tagging Mode Configurations in Server Pool

In an Oracle VM Server for SPARC server pool, you must group the servers that are homogenous in the network VLAN tagging mode to avoid any network outage. When you create a server pool, you can select the network tagging mode for the networks configured with VLAN ID to be attached to the Oracle VM Servers. You must group the servers in a way that all the Oracle VM Servers in the pool are either untagged or tagged mode for each network connection.

If the server pool is mixed with servers attached to networks in untagged and tagged mode, then there is a possibility that you will lose the network configuration of the logical domain OS.

When a logical domain is using an untagged VLAN, you cannot migrate the guest to a server that uses a tagged VLAN.

For example, in the following scenario, there is a server pool with two Oracle VM Servers CD1 and CD2. CD1 has a logical domain LD1. A network 192.0.2.0/24 with VID=100 is attached to CD1 in tagged mode. When you attach the network on the CD1, a VNIC is created with VID=100 and configured with an IP address to reach the network 192.0.2.0/24. The same network is attached to CD2 in untagged mode. LD1 is also connected to the network 192.0.2.0/24. When you migrate the LD1 from CD1 to CD2 or when you start the shutdown-detached LD1 on CD2, the LD1 will successfully start on the CD2. However, the OS of the LD1 is not able to reach the network 192.0.2.0.24 with VID=100 because, the VNIC is created with VID=100. To re-establish the network connection, you must change the configuration of the VNIC in the OS of the Oracle VM Server CD2.

To avoid such issues, create server pools with homogenous network tagging mode. Or to maintain mixed network configurations, check with your administrator for maintaining the network connection of the logical domain operating system.

Creating an Oracle VM Server for SPARC Server Pool

Oracle Enterprise Manager Ops Center launches a wizard that collects information about the servers, storage and network resources, and placement policies to create a server pool.

With the support of I/O domains in Oracle VM Server for SPARC, you can assign network interfaces from I/O domains or root domains to connect to the network. The

Create Server Pool wizard provides the option to select the network interfaces from other domains.

The wizard provides the option to select the SR-IOV enabled network interfaces to connect to the network.

To Create a Server Pool for Oracle VM Server for SPARC

1. Select Server Pools in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.

   The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:
   - Name and description of the server pool.
   - Enter tags for categorizing your server pool.
   - Select Oracle VM Server for SPARC from the Virtualization Technology.

   Click **Next** to select the members of the server pool.
4. Select one or more Oracle VM Servers to add to the server pool.

   The list of Oracle VM Servers displayed have the following characteristics:
   - Not associated with a server pool
   - In a healthy state
   - Not placed in maintenance mode

   The list displays the details of CPU frequency, type, and architecture of Oracle VM Servers. You can add Oracle VM Servers of different CPU type. If you want to live migrate the logical domains, then select Oracle VM Servers that are compatible for migration.

   Click **Next** for configuring the I/O domains and root domains I/O resources for exclusive access.
5. When Oracle VM Servers have I/O domains or root domains, you can select to use the I/O resources exclusively for logical domains only. The exclusive access does not allow to create zones on them.

   Click **Next** to select the network domain.
6. Select the network domain to associate with the server pool.
   - When you have selected default network domain, then go to Step 9.
   - For a user-defined network domain, select the physical interfaces of each server to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 9.
7. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

   Click **Next** to configure the bonding parameters.
8. In the Configure Bonding step, specify the following parameters for Link Aggregation:
   - Load balancing policy.

- LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
- LACP timer.
- MAC address policy and provide the MAC address.

Click **Next** to select the networks and associate with the server pool.

9. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For an user-defined network domain, only the networks assigned to it are listed. Select the following details:
   - Select the networks that you want to associate with all the servers in the server pool.
   - Select the network tagging mode as Tagged or Untagged for networks configured with VLAN ID.
   - You can edit the VLAN ID if you are attaching the network for the first time in Tagged mode with the selected server.
   - Enter the number of connections for each network. You can make multiple connections to a network.
   - Select the network that you want to use for migration of logical domains.

   Click **Next** to configure the interfaces of the servers.
10. For each network connection, provide the connection details:
    - **Service Domain**: Select the domain that provides the network interface. The domain can be primary, I/O domain or the root domain.
    - **SR-IOV**: Select this option if you want to assign SR-IOV enabled network interface for the network connection.
    - **Mode**: For networks configured with VLAN ID, you can modify the option to select the network to be attached in Tagged or Untagged mode.
    - **NIC**: The network interfaces that are available in the selected domain are listed. If you have selected SR-IOV, the network interfaces that are SR-IOV enabled are listed.

      You can assign the same NIC to different network connection when the networks have different VLAN IDs and every connection is in tagged mode.
    - **Switch Name**: For SR-IOV enabled networks, there is no virtual switch creation. Instead, select the physical functions listed in the column.

      For non SR-IOV enabled networks, a virtual switch is created. Enter a name for the virtual switch for a new network connection. You can also leave it blank for Oracle Enterprise Manager Ops Center to create a name using the default naming pattern.

      When any of the selected members for the server pool is already connected to the selected network, then the existing connection details is displayed. The Connected option is selected and the virtual switch name is displayed. You can either keep the existing connection or modify the network connection to make an additional connection to the same network. To modify the network connection, you can change the service domain that provides the NIC or select another NIC from the list.
    - **Address Allocation Method**: You can select the following options:

- **Assign by DHCP**: Select this option to automatically allocate the IP address by the system.
- **Use Static IP**: Select this option and provide the IP address for the network connection.
- **Do not Plumb Interface**: Select this option when you do not want to plumb an interface. This option allows you to avoid creating a virtual NIC in the control domain for networks used by logical domains. This option is available for non SR-IOV Ethernet connections beginning with Oracle Enterprise Manager Ops Center 12.2.2.
- **Do not Allocate**: Select this option when you do not want to allocate any IP address for the network connection. You can allocate the IP address later in the logical domain OS. When you select the network interfaces from domains other than primary, then the Address Allocation Method is set to this option automatically and you cannot change it.

Click **Next** to associate the storage libraries.

11. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store logical domain metadata, ISO images, and for virtual disks of logical domain.

    You can select the following types of storage libraries:

    - **Filesystem Storage Libraries**: This includes the NAS storage libraries. Associate at least one NAS storage library with the server pool to store the logical domain metadata.
    - **Block Storage Libraries**: This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries consists of the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries consists of the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

    Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

    Click **Next**. If there are I/O domains and root domains in the selected Oracle VM Server for SPARC, then you are directed to define the association details for the libraries in the next step. Otherwise, you are directed to select the server pool policies in Step 13.

12. For each selected member of the server pool, you must select at least one domain to be associated with the storage library. The domains that are already associated with the library are displayed with the **Associate** option selected. You can deselect to remove the association and select another domain in the server.

    The action to associate the Control Domain or the I/O domains is repeated for all the selected storage libraries in the previous step.

    Click **Next** to select the server pool policies.

13. Select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:

    - **Placement Policy**: This policy decides the preferred virtualization host in the server pool to place the logical domains.
    - **Auto Balancing Policy**: This policy performs load balancing of the server pool automatically at set intervals.

14. Select one of the following placement policies:
    - **Lowest relative load**: The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. The server with the lowest relative load is considered to place a logical domain. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and logical domains are migrated to server with lowest relative load.
    - **Lowest allocated CPU and memory resources**: The total number of virtual CPU and memory resources allocated for all the logical domains in an Oracle VM Server are calculated. The server with the lowest allocated resources is considered to place a logical domain in the pool. Provide the threshold values for CPU and memory allocation. When the allocation is exceeded, the server is considered to be over-allocated and logical domains are migrated to server with lowest allocated CPU and memory resources.
    - **Minimize power consumption**: The logical domains running in a server pool are placed on the minimum number of Oracle VM Servers and the unused Oracle VM Servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded.
15. Select the automatic load balancing policies:
    - Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the logical domains automatically. Or send notifications about the approval.
    - You can also select to manually balance the resources. See Balancing Resources for more information.
16. For Automatic Recovery, you can specify whether an attempt to power-off the failed server must be performed before initiating the automatic recovery of the logical domains.

    When Oracle Enterprise Manager Ops Center has the capability to power-off the Service Processor of a failed server and you have selected the automatic recovery option, the failed server is powered off and then the logical domain recovery is started. If Oracle Enterprise Manager Ops Center cannot power-off the failed server, the recovery of the attached logical domains does not take place. This is to avoid any data corruption during the recovery process.

    Enter the time interval to check the state of the server pool members reachability. The time value is entered in seconds. The minimum interval for checking the status must be 20 seconds and the default value is 180 seconds.

    Click **Next** to view the summary of the server pool details.
17. Review the information to create a server pool for Oracle VM Server for SPARC. Click **Finish** to create the server pool.

Oracle Solaris Zones Server Pool

The following information is covered in this section:

You can now pool Oracle Solaris Zones resources and manage your zones. Oracle Solaris Zones must meet the following requirements to be added to a server pool:

- All the assets in the pool must have compatible architecture for supporting guest migration. For a global zone to be added to the server pool, it must have the same release or must be at least Oracle Solaris 10 10/08 version. This ensures that the zones can be migrated to a global zone in the server pool.
- The global zones must be in healthy state.
- The global zones must not be associated with a server pool.
- The global zones must not be placed in maintenance mode.

Note: 'Oracle Solaris 11 shared IP zones should not be part of a server pool. You cannot perform migration, connect networks or add storage resources to shared IP zones.

IP Stack in Server Pool

The network deployment for zones server pool vary depending on the Oracle Solaris OS version.

For a stand-alone Oracle Solaris 10 OS, you can attach the network in either shared IP or exclusive IP mode. Whereas, in a server pool, you can attach a network only in shared IP mode for Oracle Solaris 10 OS. You cannot make multiple connections to a network.

For Oracle Solaris 11 OS, the network is always attached in exclusive IP mode. In a server pool, the network for Oracle Solaris 11 OS is always deployed in exclusive IP mode. You can make multiple connections to a network.

When you want to have a server pool with mixture of Oracle Solaris 10 and 11 OS, then you cannot make multiple connections to a network. So, plan your server pools with compatible Oracle Solaris OS versions.

Network Tagging Mode Conditions

There are certain conditions in which the networking configuration that must be selected while creating a server pool. Refer to the following scenarios for deciding on the networking configuration when you create an Oracle Solaris Zones server pool:

- You can select networks without VLAN ID. The UI does not provide the option to select Tagged or Untagged mode.
- You can select to associate and configure the networks with VLAN ID in Tagged mode.
- You can select to associate and configure the networks with VLAN ID in Untagged mode.
- You can select to configure the networks in mixed tagging mode in the server pool. For example, you can attach the network N1 with VLAN ID = 100 in tagged mode with the server S1 and in untagged mode for server S2. Refer to Mixed Network Tagging Mode Configuration for more detailed information.
- You can attach networks whose VLAN ID is similar to another network already connected to the servers. For example, a server S1 is already connected to network N1 with VLAN ID = 100, then while creating the server pool with S1 as the member of the pool, you can attach a network N2 with VLAN ID =100.

- You can edit the VLAN ID of a network when you are attaching the network in Tagged mode for the first time.
- When you can edit the VLAN ID of the network, you cannot enter -1 as the value for the VLAN ID.
- If the selected network with a VLAN ID is already connected to the selected assets in Tagged mode, then you cannot edit the VLAN ID and make another connection.
- You cannot make multiple network connections to a member of the server pool over the same network in both tagged and untagged modes. The mode can be either in tagged or untagged mode only. For example, if you attach network N1 with VLAN ID =100 for the first time to server S1 in Tagged mode, then you cannot make another connection to the same network N1 in Untagged mode. Every other connection with network N1 must always be in Tagged mode for server S1.
- If the selected members of the pool are already connected to network N1 with VLAN ID =100, then you cannot select the same network with different VLAN ID to be connected for the server pool.

Mixed Network Tagging Mode Configuration

When you attach network to zones server pool, you can specify the network tagging mode for the networks configured with VLAN ID. You must ensure that the server pools have global zones either in tagged or untagged mode. If the server pool has mixed network tagging modes for the global zones, there might be any network outage issues for the zones depending on the action performed upon them.

If you want to maintain mixed network tagging configuration in your server pool, check with your administrator for re-establishing the network connection of the zones when the network outage occurs.

Server Pool Policies

The server pool policies that are applicable for Oracle Solaris Zones are as follows:

- **Placement Policy**

  Place the zones on the global zone that has the lowest relative load. Set the CPU utilization threshold.

  Place the zones on the global zone allocated with lowest CPU and memory resources. Set the CPU and memory allocation threshold.

  Place zones on minimum number of global zones to minimize the power consumption. Set the CPU utilization threshold.

- **Automatic Balancing**

  You can select the automatic load balancing so that the zones can be migrated automatically in the server pool whenever the thresholds are exceeded. You can schedule the automatic load balancing to take place at definite time interval.

  When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some servers are freed from the zones and powered off to minimize power consumption.

  You can set whether the approval is required from the administrator for migrating the zones.

- **Automatic Recovery**

You can set the automatic recovery for the zones in the following scenarios:

– During zone creation, you can enable automatic recovery of the zones. You can set the value for priority of recovery. Zone with a highest priority is migrated first.

– For existing zones, use the option **Enable Automatic Recovery** and **Disable Automatic Recovery** to set the automatic recovery. You can edit the zone configuration to modify the value for priority of recovery.

In an Oracle Solaris Zones server pool, you can specify that the automatic recovery of a failed server must start first with an attempt to power off the failed server, ensure that it is stopped and then initiate the automatic recovery of the zones.

Creating a Zones Server Pool

Oracle Enterprise Manager Ops Center takes you through a series of steps to collect information for creating a zones server pool.

To Create a Zones Server Pool

1. Select **Server Pools** in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.

   The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:
   - Enter a descriptive name for the server pool.
   - (Optional) Enter a description for the server pool.
   - (Optional) Enter tags for categorizing your server pool.
   - Select **Oracle Solaris Zones – SPARC** or **Oracle Solaris Zones –x86** from the Virtualization Technology menu.

     You cannot have a zones server pool with a mixture of SPARC and x86 architectures.

   Click **Next** to select the members of the server pool.
4. Select one or more global zones to add to the server pool.

   The list of global zones displayed have the following characteristics:
   - The same architecture – SPARC or x86
   - Oracle Solaris 10 10/08 OS or higher version
   - Not associated with a server pool
   - In a healthy state
   - Not placed in maintenance mode
   - Not running in an I/O domain or root domain that is configured to provide its resources exclusively to other logical domains

   The list includes the CPU architecture and OS version of the global zone. Select global zones that are compatible for migration of zones within the pool.

   Click **Next** to associate the network domain.
5. Select the network domain to associate with the server pool from the available network domains in the list, then click **Next**.

- Default network domain: If you select the default network domain, go to Step 8.
- User defined network domain: If you select a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 8.

6. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

   Click **Next** to configure the bonding parameters.

7. In the Configure Bonding step, specify the following parameters for Link Aggregation:
   - Load balancing policy.
   - LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
   - LACP timer.
   - MAC address policy and provide the MAC address.

   Click **Next** to select the networks and associate with the server pool.

8. For a default network domain, all of the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For a user-defined network domain, only the networks assigned to it are listed.

   Select the networks that you want to associate with all the global zones in the server pool. Depending on the Oracle Solaris OS versions selected for the server pool, you can make multiple connections to a network. See IP Stack in Server Pool for more information about IP stack mode for networks in the zones server pool.

   For networks configured with VLAN ID, you can specify whether the network must be attached in Tagged or Untagged mode. If you are attaching the network in Tagged mode for the first time, then you can edit the VLAN ID of the network, provided the VLAN ID is not used by any of the networks attached to the asset.

   Enter the total number of connections and click **Next** to configure the interfaces of the servers.

9. For each network connection, provide the connection details:
   - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.

     If required, you can modify the network tagging mode specified in the previous step.

     You can assign the same NIC to different network connection when the networks have different VLAN IDs and every connection is in tagged mode.
   - When supported by the network, you can select System Allocated for the NIC and IP address to be automatically allocated by the system.
   - You can select **Do Not Allocate IP** to skip the option of providing the IP address for the network connection.

   Click **Next** to associate the storage libraries.

10. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store zone metadata, ISO images, and for virtual disks of zones.

    You can select the following type of storage libraries:

    - **Filesystem Storage Libraries**: This includes the NAS storage libraries. Associate at least one NAS storage library with the server pool to store the zone metadata.
    - **Block Storage Libraries**: This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries consists of the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries consists of the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

    Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

    Note: When you have unmanaged storage attached to the non-global zones, then upload the script to be placed in all the global zones in the server pool. See Script to Migrate a Zone With Dependencies for more information about migrating zone with unmanaged storage.

    Click **Next** to select the server pool policies.

11. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:

    - **Placement Policy**: This policy decides the preferred virtualization host in the server pool to place the zones.
    - **Auto Balancing Policy**: This policy performs load balancing of the server pool automatically at set intervals.

12. Select one of the following placement policies:

    - **Lowest relative load**: The recent lowest memory and CPU utilization for the global zones in the server pool are calculated. The server with the lowest relative load is considered to place the zone. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and zones are migrated to server with lowest relative load.
    - **Lowest allocated CPU and memory resources**: The total number of virtual CPU and memory resources allocated for all the zones in an Oracle VM Server are calculated. The server with the lowest allocated resources is considered to place a zone in the pool. Provide the threshold values for CPU and memory allocation. When the allocation is exceeded, the server is considered to be over-allocated and zones are migrated to server with lowest allocated CPU and memory resources.
    - **Minimize power consumption**: The zones running in the server pool are placed on the minimum number of global zone and the unused servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded. Otherwise, the powered off servers can be powered on to host the zones.

13. Select the automatic load balancing policies:

- Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the zones automatically. Or send notifications for the approval.
- You can also select to manually balance the resources. See Balancing Resources for more information.

14. For Automatic Recovery, you can specify whether an attempt to power-off the failed server must be performed before initiating the automatic recovery of the zones.

    When Oracle Enterprise Manager Ops Center has the capability to power-off the Service Processor of a failed server and you have selected the automatic recovery option, the failed server is powered off and then the zone recovery is started.

    Enter the time interval to check the state of the server pool members reachability. The time value is entered in seconds. The minimum interval for checking the status must be 20 seconds and the default value is 180 seconds.

    Click **Next** to view the summary of the server pool details.

15. Review the information to create a server pool for Oracle Solaris Zones. Click **Finish** to create the server pool.

Oracle VM Server for x86 Server Pool

You can group one or more Oracle VM Servers in an Oracle VM Manager and create server pools. When you create a server pool for Oracle VM Server for x86 systems, you must have the following information defined:

- Cluster file system
- Server pool master
- Virtual IP address for the server pool master

Cluster File System

Shared access to the server pool resources is a must for providing for high availability for the virtual machines running in the Oracle VM Servers of the server pool. This is achieved by cluster file system OCFS2 which allows multiple Oracle VM Servers to access the same disk at the same time. OCFS2 ensures that the Oracle VM Servers in a server pool can access and modify resources in the shared repositories in a controlled manner.

When you create a server pool, you must specify the server pool file system for the cluster heartbeat and other cluster information. The file system can be NFS shares or LUNs of iSCSI or SAN based storage servers. Oracle VM formats the server pool file system as OCFS2 file system.

In Oracle Enterprise Manager Ops Center, the cluster is always enabled by default. The cluster configuration is pushed out to all the Oracle VM Servers in the server pool and the cluster heartbeat starts when the server pool is created. You can set a separate network for this cluster heartbeat. See Manage Networks for more information about setting up the networks and their role for an Oracle VM Server.

Server Pool Master

An Oracle VM Server is internally elected as server pool master. You cannot set the role for an Oracle VM Server. When the elected Oracle VM Server fails, the role is set for another Oracle VM Server in the server pool. The virtual IP address provided while

creating the server pool is assigned to the Oracle VM Server that has been elected as server pool master.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for x86 are as follows:

- **Placement Policy**

  Places the virtual machines on the Oracle VM Server that has the lowest relative load. Set the CPU utilization threshold.

  Places the virtual machines on minimum number of Oracle VM servers to minimize the power consumption. Set the CPU utilization threshold.

- **Automatic Balancing**

  You can select the automatic load balancing so that the virtual machines can be migrated automatically in the server pool whenever the thresholds are exceeded.

  When you enable automatic load balancing for Oracle VM Server for x86 server pool, the server pool is checked continuously for the selected placement policy. When the threshold exceeds, the virtual machines are migrated from one Oracle VM Server to another.

  When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM Server host servers are freed from the logical domains and powered off to minimize power consumption.

  When the servers are overloaded and there are no other servers in the pool to host the virtual machines, then the policy starts a powered-off server using its Wake-on-LAN capability and live migrate the virtual machines to take up the load. The Wake-on-LAN capability must be enabled on the BIOS of the Oracle VM Server.

- **Automatic Recovery**

  When an Oracle VM Server fails, the virtual machines are migrated to another Oracle VM Server in the server pool. The automatic recovery for a virtual machine is set in the following scenarios:

  – During virtual machine creation, select the option Enable High Availability. This ensures that the virtual machines are migrated when an Oracle VM Server fails.

  – For an existing virtual machine, use the option Enable or Disable Automatic Recovery to set the automatic recovery.

Creating an Oracle VM Server for x86 Server Pool

The following procedure describes the steps in the wizard that is launched to create an Oracle VM Server for x86 server pool.

To Create an Oracle VM Server for x86 Server Pool

1. Select Server Pool in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.

   The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:

- Name and description of the server pool.
- Enter tags for categorizing your server pool.
- Select Oracle VM Server for x86 from the Virtualization Technology.

Click **Next** to define the server pool configuration.

4. Define the following configuration details:
   - Select the Oracle VM Manager in which the Oracle VM Servers are discovered and owned.
   - Select the network domain.
   - Select a network from the list of networks available in the network domain.
   - Enter the virtual IP address that is assigned to the server pool master. When the server pool master changes, the IP address is assigned to the new Oracle VM Server.
   - Select the server pool file system to store cluster hearbeat and other cluster information. The file system can be either NFS shares on a NFS file servers or LUNs of SAN and iSCSI based storages. The NFS file servers and the storage servers are known and reachable to the selected Oracle VM Manager.

   Click **Next** to select the members of the pool.
5. Select one or more Oracle VM Servers to add to the server pool.

   The list of Oracle VM Servers displayed have the following characteristics:
   - Owned by the Oracle VM Manager. Refer to Manage Ownership of Oracle VM Servers to own an Oracle VM Server.
   - Not associated with a server pool
   - In a healthy state
   - Not placed in maintenance mode
   - Have no virtual machines in the running, shutdown or suspended state

   Select Oracle VM Servers that are compatible for migration of virtual machines. To migrate virtual machines within a server pool, the Oracle VM Servers systems must be identical in model.

   Click **Next** to associate the network domain.
6. Select the network domain to associate with the server pool.
   - When you have selected default network domain, go to Step 7.
   - For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. You cannot bond the interfaces in Oracle VM Server for x86. Bonding can be done in only individual OS of the Oracle VM Server.

   Click **Next** to select the networks and associate with the server pool.
7. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For a user-defined network domain, only the networks assigned to it are listed.

   Select the networks that you want to associate with all the servers in the server pool. You cannot make multiple connections to a network. Limit the number of connections to 1.

Click **Next** to configure the interfaces of the servers.

8. For each network connection, specify the following details:
   - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.
   - You can assign the same NIC to different network connection when the networks have different VLAN IDs.
   - You can select System Allocated for the NIC and IP address to be automatically allocated by the system.

   Click **Next** to associate the storage libraries.
9. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store virtual machine metadata, ISO images, and for virtual disks of virtual machine.

   You can select the following type of storage libraries:
   - **Filesystem Storage Libraries**: This includes the Oracle VM storage repositories. At least one Oracle VM storage repository must be associated with the server pool.
   - **Block Storage Libraries**: This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

   Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

   Click **Next** to select the server pool policies.
10. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
    - **Placement Policy**: This policy decides the preferred virtualization host in the server pool to place the logical domains.
    - **Auto Balancing Policy**: This policy performs load balancing of the server pool automatically at set intervals.
11. Select one of the following placement policies:
    - **Lowest relative load**: The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. The server with the lowest relative load is considered to place a virtual machine. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and virtual machines are migrated to server with lowest relative load.
    - **Minimize power consumption**: The virtual machines running in a server pool are placed on the minimum number of Oracle VM Servers and the unused Oracle VM Servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded. Otherwise, the powered off Oracle VM Servers can be powered on to host the virtual machines.
12. Select the automatic load balancing policies:

- Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the virtual machines automatically. Or send notifications about the approval.
- You can also select to manually balance the resources. See Balancing Resources for more information.

Click **Next** to view the summary of the server pool details.

13. Review the information to create a server pool for Oracle VM Server for x86. Click **Finish** to create the server pool.

Manage Server Pools

You must monitor the requirements of the server pool so that there are resources to host the guests and run them efficiently. You might need to add more storage and network resources to run the guests. Also, the server pools gives the infrastructure support for the virtual datacenters. Oracle Enterprise Manager Ops Center provides the following management functions to manage the resources in the server pool:

- Editing Server Pool Parameters
- Adding Virtualization Hosts
- Associating Network Domains
- Attaching Networks
- Associating Libraries
- Creating Guests
- Migrating Multiple Guests
- Migrating Zones
- Migrating Logical Domains
- Migrating Virtual Machines
- Balancing Resources
- Monitoring Server Pool Incidents
- Deleting Server Pool

Editing Server Pool Parameters

You can edit the following parameters and attributes of a server pool:

- Name
- Description
- Placement policy, including CPU Utilization
- Auto balancing policy
- Enable or disable automatic recovery option.
- Automatic Recovery Authorization.
- Automatic Recovery Number of Retries.

For Oracle Solaris Zones server pool, you can edit the following additional parameters:

- Check server reachability interval.

- Upload scripts to manage the file systems on unmanaged storage. See Step 10 in Oracle Solaris Zones Server Pool for more information.

For Oracle VM Server for SPARC server pool, you can edit the following additional parameters:

- Migration networks.
- Enable or disable automatic recovery to power off a failed server from Service Processor, given capabilities, before automatic recovery of attached logical domains.
- Check servers reachability.

Perform the following to modify the parameters of a server pool:

1. Expand **Assets** in the Navigation pane.
2. Select **Server Pools** in the Resource Management view.
3. Select the server pool in the Navigation pane, then click the **Summary** tab.
4. To edit the parameters, click **Edit Attributes** in the Actions pane. To add, remove or modify the tags added to the server pool, click **Edit Tags** in the Actions pane.
5. Click the **Save** icon.

Adding Virtualization Hosts

You can add more virtualization hosts to a server pool when you require more CPU and memory resources to take up the load in the server pool. Depending on the virtualization type, you can add global zones, and Oracle VM Servers of x86 and SPARC architecture.

Oracle Solaris Zones

The list of global zones that are displayed to be added to the server pool has the following characteristics:

- Same release as the existing global zones in the pool or at least Oracle Solaris 10 10/08 OS or higher version
- Not associated with a server pool
- In a healthy state and not placed in maintenance mode

When the server pool has Oracle Solaris 11 OS and multiple network connections, then the list of available global zones is limited to only Oracle Solaris 11 OS. You cannot add Oracle Solaris 10 OS as you cannot make multiple network connections to it.

Load balancing in the server pool requires the zones to be migrated within the pool. Therefore, ensure compatibility for migration within the pool when you select the members.

Refer to Network Tagging Mode Conditions for more information about the network tagging conditions while adding an asset to the server pool.

To Add Global Zones to the Server Pool

1. Select the server pool and click **Add Global Zones** in the Actions pane.

   The Add Global Zones to Server Pool Wizard is displayed.
2. Select one or more global zones to be added to the server pool.

When the server pool is associated with the default network domain, you must proceed to configure the interfaces. Otherwise, you must associate the network domain with the selected assets.

3. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces for global zones, go to Step 4.
4. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select **Bond Interface** in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

   Click **Next** to configure the bonding parameters.
5. In the Configure Bonding step, specify the following parameters for Link Aggregation:
   - Load balancing policy.
   - LACP mode. If the Ethernet switch to which the physical interface connects supports aggregation, then specify the LACP mode.
   - LACP timer.
   - MAC address policy and provide the MAC address.

   Click **Next** to define the network connection.
6. For each network connection of the global zones, specify the following details:
   - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.

     You can assign the same NIC to different network connection when the networks have different VLAN IDs and every connection is in tagged mode.
   - For networks configured with VLAN ID, you can select the network to be attached in Tagged or Untagged mode. If you are attaching the network in Tagged mode for the first time, then you can edit the VLAN ID of the network, provided the VLAN ID is not used by any of the networks attached to the asset.
   - You can select System Allocated for the NIC and IP address to be automatically allocated by the system.

   Click **Next** to view the summary.
7. Review the information provided and click **Finish** to add the selected global zones to the server pool.

Oracle VM Server for SPARC

The list of Oracle VM Servers displayed to be added to the server pool have the following characteristics:

- Not associated with a server pool
- In a healthy state and not placed in maintenance mode

Refer to Network Tagging Mode Conditions for more information about the networking configurations when you add Oracle VM Server to the server pool.

To Add Oracle VM Servers to the Server Pool

1. Select the server pool and click **Add Oracle VM Servers** in the Actions pane.

The Add Oracle VM Servers to Server Pool Wizard is displayed.

2. Select one or more Oracle VM Servers to be added to the server pool.

   When the server pool is associated with the default network domain, you must proceed to configure the interfaces. Otherwise, you must associate the network domain with the selected assets.

3. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces for Oracle VM Server for SPARC server pool, go to Step 4.

4. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

   Click **Next** to configure the bonding parameters.

5. In the Configure Bonding step, specify the following parameters for Link Aggregation:

   - Load balancing policy.
   - LACP mode. If the Ethernet switch to which the physical interface connects supports aggregation, then specify the LACP mode.
   - LACP timer.
   - MAC address policy and provide the MAC address.

   Click **Next** to define the network connection.

6. For Oracle VM Server for SPARC, you can select the network interfaces from I/O domains and root domains to connect to the network. You can also select SR-IOV enabled network interfaces that are available from primary and root domains. Select the following information for the network connection:

   - **Service Domain**: Select the domain that provides the network interface. The domain can be primary, I/O domain, or the root domain.
   - **SR-IOV**: Select this option if you want to assign SR-IOV enabled network interfaces for the network connection. SR-IOV enabled network interfaces are available only from primary and root domain.
   - **Mode**: For networks configured with VLAN ID, you can select the network to be attached in Tagged or Untagged mode. If you are attaching the network in Tagged mode for the first time, then you can edit the VLAN ID of the network, provided the VLAN ID is not used by any of the networks attached to the asset.

     Note: It is recommended to attach the network in the same tagging mode as the other Oracle VM Servers of the server pool to avoid any network outage.

   - **NIC**: The network interfaces that are available in the selected domain are listed. If you have selected SR-IOV, the network interfaces that are SR-IOV enabled are listed.

     You can assign the same NIC to different network connection when the networks have different VLAN IDs and every connection is in tagged mode.

- **Switch Name**: For SR-IOV enabled networks, there is no virtual switch creation. Instead, select the physical functions listed in the column.

  For non SR-IOV enabled networks, a virtual switch is created. Enter a name for the virtual switch for a new network connection. You can also leave it blank for Oracle Enterprise Manager Ops Center to create a name using the default naming pattern.

  When any of the selected members for the server pool is already connected to the selected network, then the existing connection details is displayed. The Connected option is selected and the virtual switch name is displayed. You can either keep the existing connection or modify the network connection to make an additional connection to the same network.

- **Address Allocation**: Select **Do not Allocate** to skip the option of providing the IP address. Select **Use Static IP** to provide the IP address, or select **System Allocated** to have the system automatically allocate an IP address.

  Do Not Plumb Interface: Select this option when you want to avoid creating a virtual NIC in the control domain for networks used by logical domains. This option is available for non SR-IOV Ethernet connections beginning with Oracle Enterprise Manager Ops Center 12.2.2.

Click **Next** to view the summary.

7. Review the information provided and click **Finish** to add the selected assets to the server pool.

Oracle VM Server for x86

The list of Oracle VM Servers displayed to be added to the server pool have the following characteristics:

- Owned by the Oracle VM Manager
- Not associated with a server pool
- In a healthy state and not placed in maintenance mode
- Have no virtual machines in the running, shutdown or suspended states

To Add Virtualization Hosts to the Server Pool

1. Select the server pool and click **Add Oracle VM Servers** in the Actions pane.

   The Add Oracle VM Servers to Server Pool Wizard is displayed.

2. Select one or more virtualization hosts to be added to the server pool.

   When the server pool is associated with the default network domain, you must proceed to configure the interfaces. Otherwise, you must associate the network domain with the selected assets.

3. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. For Oracle VM Server for x86, you cannot bond the interfaces and click **Next**.

4. For each network connection of Oracle VM Server for x86 server pool, specify the following details:

   - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.

     You can assign the same NIC to different network connection when the networks have different VLAN IDs.

- You can select System Allocated for the NIC and IP address to be automatically allocated by the system.

Click **Next** to view the summary.

5. Review the information provided and click **Finish** to add the selected assets to the server pool.

Removing Virtualization Hosts from Server Pool

You can remove the virtualization hosts from the server pool. For zones and Oracle VM Server for SPARC, you can remove them from the server pool even when there are zones and logical domains running on them. All the network and storage resources are not unconfigured. Instead, the resources are retained as it is in the virtualization hosts.

To remove Oracle VM Server for x86 from a server pool, you must stop or migrate the virtual machines that are running in it.

Associating Network Domains

All networks in a network domain that is associated with a server pool are available to the members of the server pool. To associate a user-defined network domain with a server pool, there must be a connection between the virtualization servers and the fabrics in the network domain. The fabric connection is not required for the default network domain.

The default network domain includes all the networks in Oracle Enterprise Manager Ops Center. When you associate a server pool with the default network domain, all the networks in Oracle Enterprise Manager Ops Center are available for attaching with the server pool.

For a user-defined network domain, you must define the physical interfaces of each virtualization server in the server pool to connect to each fabric in the network domain. You can also bond the interfaces of the virtualization servers.

Note: You cannot the bond the interfaces of Oracle VM Server for x86 servers.

To Associate the Network Domain to a Server Pool

1. Select the server pool.
2. Click **Associate Network Domain** in the Actions pane.

   The Associate Network Domain Wizard is displayed.
3. Select the network domain from the list.
4. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 7. For Oracle VM Server for x86, you cannot bond the interfaces and go to Step 7.

   For default network domain, you are taken to the Summary step.
5. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

   Click **Next** to configure the bonding.

6. In the Configure Bonding step, specify the following parameters for Link Aggregation:
   - Load balancing policy.
   - LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
   - LACP timer.
   - MAC address policy and provide the MAC address.

   Click **Next** to view the summary.
7. Review the information and click **Finish** to associate the network domain with the server pool.

When you associate with a user-defined network domain, only the assigned networks in the network domain are available for attaching with the server pool. Therefore, ensure that you assign the required networks to the domain to attach to the server pool.

Attaching Networks

After the network domain is associated with a server pool, you can attach the networks in the network domain. Depending on the type of virtualization, the multiple connection to a network is allowed.

For Oracle VM Server for SPARC, you can make multiple connections to a network. Refer to the Chapter 19, "Oracle VM Server for SPARC" for more information.

For Oracle Solaris Zones, you can make multiple connections to a network when the server pool contains only Oracle Solaris 11 OS. For Oracle Solaris 10 OS, you can make only one connection. Also, when the server pool has mixture of Oracle Solaris 10 and 11 OS, then also you can make only single connection.

Refer to Network Tagging Mode Conditions for more information about selecting the correct tagging modes when you attach the network to Oracle Solaris Zones and Oracle VM Server for SPARC.

To Attach Networks to Server Pool

1. Select the server pool.
2. Click **Attach Networks** in the Action pane.

   The Attach Network Wizard is displayed. All the networks assigned to the network domains are listed. For default network domain, all the networks are listed. The table also displays existing number of connections to the server pool.
3. Select one or more networks to attach to the server pool.

   Click **Next** to specify the number of connections.
4. Enter the total number of connections required as per the virtualization type:
   - For Oracle VM Server for SPARC, you can enter multiple connections. For each network connection, a virtual switch is created. Virtual switch creation is not applicable for SR-IOV enabled networks. For each network connection, you can assign the network interfaces from the I/O domains and root domains. Specify the tagging mode as Untagged or Tagged for the networks configured with VLAN ID.

- For Oracle VM Server for x86, you can enter only one connection per network. You cannot make multiple connections.
- For Oracle Solaris Zones, the Oracle Solaris OS version plays an important role. For Oracle Solaris 10 OS, the network is always deployed in shared IP mode. You cannot make multiple connections. For Oracle Solaris 11 OS, the network is always deployed in exclusive IP mode. You can make multiple connections to the network. For each network connection, a VNIC is created when you boot the zone. For a server pool of mixture of Oracle Solaris 10 and 11 OS, the network connection is limited to single connection.

Note: The number of connections do not limit the number of guests to be connected to be network.

Click **Next** to configure the interfaces.

5. When one of the virtualization hosts are connected to the network, the corresponding row for network configuration for that host is not displayed. If not, specify the information for each connection:
   - Specify the NIC and IP address for each connection.

     You can assign the same NIC to different network connection when the VLAN IDS are different for the networks and every connection is in tagged mode.
   - Select **Do Not Allocate IP** to connect to the network without specifying the IP address. This option is available for zones and Oracle VM Server for SPARC systems.
   - Select **System Allocated** for the system to take care of the NIC and IP address allocation.

   For Oracle VM Server for SPARC and zones, you can select the network interfaces from I/O domains and root domains to connect to the network. You can also select SR-IOV enabled network interfaces that are available from primary and root domains. Select the following information for the network connection:
   - **Service Domain**: Select the domain that provides the network interface. The domain can be primary, I/O domain or the root domain.
   - **SR-IOV**: Select this option if you want to assign SR-IOV enabled network interfaces for the network connection. SR-IOV enabled network interfaces are available only from primary and root domain.
   - **Mode**: For networks configured with VLAN ID, you can modify the option to select the network to be attached in Tagged or Untagged mode.
   - **NIC**: The network interfaces that are available in the selected domain are listed. If you have selected SR-IOV, the network interfaces that are SR-IOV enabled are listed.

     You can assign the same NIC to different network connection when the networks have different VLAN IDs and every connection is in tagged mode.
   - **Switch Name**: For SR-IOV enabled networks, there is no virtual switch creation. Instead, select the physical functions listed in the column.

     For non SR-IOV enabled networks, a virtual switch is created. Enter a name for the virtual switch for a new network connection. You can also leave it blank for Oracle Enterprise Manager Ops Center to create a name using the default naming pattern.

When any of the selected members for the server pool is already connected to the selected network, then the existing connection details is displayed. The Connected option is selected and the virtual switch name is displayed. You can either keep the existing connection or modify the network connection to make an additional connection to the same network.

- **Address Allocation**: Select **Do not Allocate** to skip the option of providing the IP address. Select **Use Static IP** to provide the IP address, or select **System Allocated** to have the system automatically allocate an IP address.
- Beginning with Oracle Enterprise Manager Ops Center release 12.2.2, select **Do Not Plumb Interface** when you do not want to plumb the interface.

Click **Next** to view the summary of the information.

6. Review the information provided and click **Finish** to attach the networks to the server pool.

Detaching Networks from Server Pool

You can detach the networks that have been attached to the server pool. Detaching the networks results in unplumbing the network connection to the host's NIC from all the virtualization hosts in the server pool.

To Detach a Network from Server Pool

1. Select the server pool.
2. Select the **Network** tab in the center pane.
3. Select the network that you want to remove.
4. Click the icon **Unbind Network** from Server Pool.

   The Unbind Network from a Server Pool window is displayed.
5. Click **Unbind** to confirm the network removal.

Associating Libraries

Table 21–4 lists the types of storage libraries that you can associate with the server pool.

Table 21–4 Supported Libraries

	Oracle VM Storage Repositories	NAS Libraries	SAN Fibre Channel Libraries	SAN iSCSI Libraries
Oracle VM Server for SPARC	No	Yes	Yes	Yes
Oracle VM Server for x86	Yes	Yes	Yes	Yes
Oracle Solaris Zones	No	Yes	Yes	Yes

When you associate libraries with the server pool, only the libraries that are reachable from the virtualization hosts in the server pool are listed. Depending on the type of library, you either provide virtual disks or LUNs to the guests' storage.

Except for Oracle VM Server for SPARC server pool, the Associate Libraries action for zones and Oracle VM Server for x86 server pool displays the available libraries. Select

the libraries and click Associate to associate the libraries with the zones or Oracle VM Server for x86 server pool.

For Oracle VM Server for SPARC server pool, the procedure to associate and disassociate the storage libraries are described as follows.

To Associate Libraries with Oracle VM Server for SPARC Server Pool

1. Select the Oracle VM Server for SPARC or server pool from the Server Pools list.
2. Click **Associate Libraries** in the Actions pane.

   The Associate Library Wizard is displayed.
3. Select the libraries that you want to associate with the server pool, then click **Next**.
4. Select **Associate** action for the Control Domain, I/O domain, or both. You must select at least one domain per server.

   When the storage library is already associated with domains, then the Associate option is selected. You can deselect to disassociate the storage library from the domain.

   This action is repeated for each selected storage library. Click **Next**.
5. Review the summary of the association and click **Finish** to associate the library with the selected domains in the server pool.

To Disassociate Storage Libraries from Oracle VM Server for SPARC Server Pool

1. Select the Oracle VM Server for SPARC server pool in the Server Pools list.
2. Select the **Libraries** tab in the Center pane.

   All the storage libraries that are associated with the server pool are displayed.
3. Select the storage library that you want to disassociate from the server pool.

   The Disassociate Library window is displayed.
4. When the storage library is associated with the domains are displayed with the **Disassociate** option selected. You can deselect the Disassociate action if you do not want to remove the association.
5. Click **Finish** to disassociate the storage library from the selected domains in the Oracle VM Server for SPARC server pool.

Creating Guests

The guests refer to the logical domains, zones or the virtual machines that can be created in the virtualization hosts. According to the virtualization type of the server pool, you have the following options:

- Oracle VM Server for SPARC – Create Logical Domains
- Oracle Solaris Zones – Create Zones
- Oracle VM Server for x86 – Create Virtual Machines

These options trigger the deployment plans for the corresponding guests. Refer to the following chapters to refer to the profile and plan creation for the guests:

- Chapter 18, "Oracle Solaris Zones"
- Chapter 19, "Oracle VM Server for SPARC"
- Chapter 20, "Oracle VM Server for x86"

Migrating Multiple Guests

When you place a virtualization host in a server pool, the option to migrate one or more guests is enabled. The following options become available for the virtualization hosts:

- Oracle Solaris Zones – Migrating Zones
- Oracle VM Server for SPARC - Migrating Logical Domains
- Oracle VM Server for x86 – Migrating Virtual Machines

Migrating Zones

You can migrate multiple zones from a global zone which is in a server pool. Migrate the zones to an individual global zone or other zones server pool. When you migrate zones, the source and the target global zone must be compatible. The target global zones must have the following characteristics:

- Running at least Oracle Solaris 10 8/07 OS.
- Can access all the storage libraries associated with the zone.

Apart from compatibility, the target and the source global zone might have differences in the patches and packages installed on them. Choose to update the zone to match the patches and packages of the target global zone.

Note: You cannot downgrade the patches and packages of the zone. The migration fails in such scenario.

You can also force the migration of the zone without updating the patches and packages. Select the update options while migrating the zone.

To Migrate Multiple Zones

1. Select the global zone from which you want to migrate the zones.
2. Click **Migrate Zones** in the Actions pane.

   The Migrate Zones Wizard is displayed.
3. The list includes the zones running in the global zone. Select one or more zones from the list.

   Click **Next**.
4. Select an individual global zone or server pool to which you can migrate the zones.

   The table displays the list of eligible global zones and server pool to which you can migrate the zones. The target global zone in the server pool depends on the server pool placement policy.
5. Select an update option to continue with migration.

   The source and the target global zones might not be in the same patch level. Either you can select to update the patches and packages of zone to match the target global zone or continue migration without updating the zone.
6. Review the details and click **Finish** to migrate the zones.

Migrating Logical Domains

Migrate multiple logical domains from an Oracle VM Server to another Oracle VM Server in the same server pool. You cannot do cross server pool migration.

Only from Oracle VM Server for SPARC 2.1 version, live migration of logical domains is supported. For earlier releases, it is cold migration. The logical domains are shut down and then restarted on the target server.

See Migrate Logical Domains in Chapter 19, "Oracle VM Server for SPARC" for more conditions in migrating a logical domain.

To Migrate a Logical Domain

1. Select the Oracle VM Server in a server pool.
2. Click **Migrate Logical Domains** in the Actions pane.
3. The list of logical domains running in the Oracle VM Server are listed. Select one or more logical domains from the list.
4. Select an Oracle VM Server from the list.

   The table lists the Oracle VM Servers that have enough resources to host the logical domains in the same server pool. The Oracle VM Server are listed in the decreasing order of preference.
5. Review the information and click **Finish** to migrate the logical domains.

Migrating Virtual Machines

For Oracle VM Server for x86 systems, you can migrate virtual machines only within a server pool. Live migration of virtual machines is supported. The eligible Oracle VM Servers have the following characteristics:

- Owned by the Oracle VM Manager.
- Placed in the same server pool as the source Oracle VM Server.
- Have required resources to host the virtual machines.
- Identical in machine make and model of the source Oracle VM Server.

To Migrate a Virtual Machine

1. Select the Oracle VM Server in the server pool.
2. Click **Migrate Virtual Machines** in the Actions pane.
3. Select the virtual machines that you want to migrate.
4. Select the Oracle VM Server to which you want to migrate the virtual machines.

   The table lists the Oracle VM Servers that are eligible to host the virtual machines.
5. Review the summary and click **Finish** to migrate the virtual machines.

Balancing Resources

When you have selected to balance server pool resource manually, use the Balance Resources action to check for the under-utilization or over-utilization of resources. Select a server pool and click **Balance Resources** in the Actions pane to display the current utilization of the virtualization hosts in the server pool. When balancing is required according to the placement policy, Oracle Enterprise Manager Ops Center

displays a list of target server that can accept migrated guests. Click **Balance Resources** to start the migration job.

Monitoring Server Pool Incidents

When the server pool does not meet the policies and rules set, warning or critical incidents are raised accordingly. For example, when the servers in the pool are overloaded and need to migrate a guest to another server.

Incidents also provide information about the result of a server pool status on automatic load balancing.

The guests are recovered according to the policy set in the server pool for automatic recovery. The incidents are raised that describes the guests that are recovered and in which servers they are recovered in the server pool.

Incidents are also raised when there are no resources or servers available in the server pool for recovering the guests. Incidents provide more information about the scenario.

Refer to Chapter 9, "Incidents" for more information about handling and viewing incident messages.

Deleting Server Pool

Use the option **Delete Server Pool** to release the virtualization hosts back to stand-alone state. The type of server pool imposes some limitations to delete a server pool.

Oracle VM Server for SPARC

You can delete the Oracle VM Server for SPARC server pool even when the logical domains are attached to it and running. There is no restriction for deleting an Oracle VM Server for SPARC server pool to shut down the logical domains. All the network and storage resources are not unconfigured from the Oracle VM Servers in the server pool. The network and storage resources are retained by the servers.

Oracle Solaris Zones

For a zones server pool, you can delete the server pool and the zones continue to be attached to the global zone. Oracle Enterprise Manager Ops Center does not provide any restrictions to shut down the zones before deleting a zones server pool.

Oracle VM Server for x86

Shut down all the virtual machines and remove all Oracle VM Servers except for the Master Oracle VM Server before deleting the server pool. The virtual machine metadata and virtual disks are stored in the storage library. The virtual machine details are not lost and you can start the virtual machine.

Related Resources for Server Pools

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- Chapter 20, "Oracle VM Server for x86"
- Chapter 19, "Oracle VM Server for SPARC"
- Chapter 18, "Oracle Solaris Zones"
- Chapter 17, "Networks for Virtualization"

- Chapter 2, "Asset Management"
- Chapter 16, "Storage Libraries for Virtualization"
- Chapter 22, "Virtual Datacenters"
- *Oracle Enterprise Manager Ops Center Administration Guide*

See the following how to documentation in the Deploy How To tab in the library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` for end-to-end examples.

- *Oracle Enterprise Manager Ops Center Exploring Your Server Pools*
- *Oracle Enterprise Manager Ops Center Creating a Server Pool for Zones*
- *Oracle Enterprise Manager Ops Center Creating a Server Pool for Oracle VM Server for SPARC*

See the *Oracle Enterprise Manager Ops Center Exploring Your Server Pools* how to documentation in the library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm` for an end-to-end example of how to manage server pools.

22

Virtual Datacenters

This chapter includes the following information:

Introduction to Cloud Management

Oracle Enterprise Manager Ops Center provides the platform to consolidate the physical resources in your data center, and to build and operate your cloud services. The physical resources such as virtualization servers, storage, and network are pooled which are accessed by users to build their applications.

Oracle Enterprise Manager Ops Center provides comprehensive management solution to deploy, configure, and manage the virtualization servers, storage resources, and network fabrics. Also, pool the virtualization resources that share the storage and network resources. All these features are leveraged and can be deployed as Infrastructure-as-a-Service (IaaS) cloud platform service model in Oracle Enterprise Manager Ops Center.

In Oracle Enterprise Manager Ops Center, the provision to setup the IaaS cloud platform service is available as vDC Management. Virtual Datacenter or vDC is a collection of the server pools that share the common storage and network resources. A cloud administrator sets up the infrastructure and provide access to cloud users. Cloud users use the allocated resources to create guests with an OS installed, deploy applications, monitor, and manage the applications. The cloud user is provided the access to Oracle Enterprise Manager Ops Center UI to view and manage their applications.

Oracle Enterprise Manager exposes APIs and command–line interface (CLI) to enable the access to a subset of the Virtual Datacenter functionality. The cloud user can also manage the allocated resources programmatically by calling the IaaS web services directly.

The features accessible through the Cloud Infrastructure API and CLI are available for a Cloud User. The features that are restricted to the Cloud Administrator, such as infrastructure configuration and setup are not available through the Cloud Infrastructure API and CLI except for listing the accessible accounts and to create key based access to cloud users who have access to the accounts already. All the functions that are available for a Cloud User are also available for the Cloud Administrator.

When you log in to Oracle Enterprise Manager Ops Center, the views and access for different sections of the UI are different for a cloud administrator and a cloud user.

Cloud User

If you are a cloud user looking for more information about creating and managing virtual servers, then go to the section Overview of Cloud Users.

See *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide* for more information about using Cloud Web Service and Cloud Infrastructure API and CLI.

Cloud Administrator

If you are a cloud administrator, continue with the following sections that describe how to create and manage the virtual datacenters in Oracle Enterprise Manager Ops Center.

Oracle Engineered Systems

If you want to manage virtual datacenters on engineered systems such as Oracle SuperCluster, then refer to Chapter 23, "Oracle Engineered Systems" for more information.

Roles for Managing Virtual Datacenter

The following table lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 22–1 Virtual Datacenter Tasks and Roles

Task	Role
Create and Manage Virtual Datacenters	Cloud Admin
Create and Manage Accounts	Cloud Admin
Create and Manage vServer Types	Cloud Admin
Add and Manage Cloud Users	Cloud Admin

Required User Role

A cloud administrator requires the Cloud Admin role to create and manage the virtual datacenters, accounts, and cloud users. The Cloud Admin role has the necessary permissions of Asset Admin and Network Admin roles to setup and manage the virtual datacenter infrastructure.

Specify vDC Privileges

By default, the addition of the Cloud admin role to a user does not give full vDC privileges. To access an existing vDC as a cloud administrator, you must have the proper privileges set to view and manage it. Adding only the cloud admin role is not enough, you must have the privileges for the vDCs to be managed.

Contact your Ops Center administrator to manage your user's role to specify proper vDC privileges. See *Oracle Enterprise Manager Ops Center Administration Guide* for more information about how to manage user's roles and privileges.

Actions Available in vDC Management

The virtual datacenter configuration involves the setup of the infrastructure, managing the resource workloads, and understanding cloud user requirements. The user with the cloud administrator role has the following actions:

- Create and manage virtual datacenters
- Create and manage accounts
- Manage access to accounts
- Manage physical resources of virtual datacenter
- Create and manage vServer types

Location of Virtual Datacenter Information in the User Interface

Table 22–2 lists where to find different information for virtual datacenter in the UI.

Table 22–2 Location of Virtual Datacenter Information in the UI

To See	Location
Virtual Datacenter	Expand vDC Management in the Navigation pane.

Table 22–2 (Cont.) Location of Virtual Datacenter Information in the UI

To See	Location
Network resources for a vDC	Expand vDC Management in the Navigation pane and select the vDC. Expand the selected vDC and select Network. The center pane displays information about network resources in the vDC.
Storage resources for a vDC	Expand vDC Management in the Navigation pane and select the vDC. Expand the selected vDC and select Storage. The center pane displays information about storage resources in the vDC.
Server Pools for a vDC	Expand vDC Management in the Navigation pane and select the vDC. Expand the selected vDC and select Server Pools. The center pane displays information about server pools in the vDC.
Accounts in a vDC	Expand vDC Management in the Navigation pane and select the vDC. Expand the selected vDC and select Accounts. All the accounts created in the vDC are listed.
Options for managing vDCs	Expand vDC Management in the Navigation pane and select the vDC. The Actions pane list the options for managing the vDCs and its accounts.

Overview of Virtual Datacenter

In Oracle Enterprise Manager Ops Center, consolidate the virtualization servers, storage, and network resources, and enable those resources to be utilized optimally and securely for mixed and dynamic workloads. This optimization of your resources is done by creating a Virtual Datacenter (vDC) in Oracle Enterprise Manager Ops Center.

The physical resources allocated of a vDC are entitled to accounts as virtual resources under quotas set by the cloud administrator. Accounts provide the required capabilities to manage the allocated resources. An account entitles designated cloud users the right to use its virtual computing, network, and storage resources.

A cloud user with access to different accounts can create virtual servers, known as vServers, and host or access applications.

A cloud administrator is involved in setting up the infrastructure for the vDC, creating and managing vDCs, creating and managing accounts, and managing access to cloud users.

Creating Virtual Datacenters

Virtual datacenter (vDC) is a consolidation of your physical resources that share the network and storage resources. The physical resources that form a virtual datacenter are:

- **Homogenous server pools**

  The server pools are groups of virtualization supported servers that share compatible chip architecture. The supported virtualization types are:

  – Oracle Solaris Zones

    A vDC based on Oracle Solaris Zones can not use a server pool with virtualization host with a mix of operating systems. For creating a vDC using server pools based on Oracle Solaris Zones, all virtualization hosts of the server pool must have the same configuration and only one from the following:

    * Oracle Solaris 10 Zones for x86

    * Oracle Solaris 11 Zones for x86
    * Oracle Solaris 10 Zones for SPARC
    * Oracle Solaris 11 Zone for SPARC
  - Oracle VM Server for x86
  - Oracle VM Server for SPARC

  Server pools are of single virtualization technology. For example, you can have a server pool of Oracle VM Server for SPARC servers only and not a mixture of Oracle VM Server for x86 and Oracle VM Server for SPARC servers. You cannot have server pools that have mixture of different virtualization technology. A vDC is based on any one of the supported virtualization technology. This infrastructure difference brings in some major differences in using the vDC. The cloud user might not be aware of the background infrastructure but the cloud administrator ensures that the vDC resources are always highly available to a cloud user. See Setting Up the Server Pool for more information about the server pool setup required for virtual datacenter.

- **Storage**

  The storage is inherited from the server pools. The storage is used for storing root disk of the virtual server created in the accounts, virtual server configuration data, templates, ISO images, FLAR images, and also used as volumes. While creating a vDC, you can allocate some storage resources that are used for volumes and root disks. Volumes are used to attach to the vServers. See Setting Up Storage Resources for more information about the storage resources required for virtual datacenter.

- **Networks**

  Plan the network requirements and attach the network to the server pools. vDC inherits the network resources from the server pools. Server pools in a vDC are associated with the same user-defined network domain. See Setting Up the Server Pool and Setting Up Network Resources for more information about setting up your network resources for virtual datacenter.

Setting Up the Server Pool

You can create vDC with server pools based on the following types of virtualization technology:

- Oracle VM Server for x86
- Oracle Solaris Zones
- Oracle VM Server for SPARC

Ensure all the required networks are assigned to the user-defined network domain that is associated with the server pool. The vDC creation wizard displays the list of available server pools in Oracle Enterprise Manager Ops Center. The list includes the server pools that conform to the following characteristics:

- The server pool is not empty and there are virtualization servers placed in the server pool.
- The server pool is associated with a user-defined network domain. Each network in the user-defined network domain must be connected to every server in the server pool.

- Server pools of a single supported virtualization technology.

> **Note:** For Oracle Solaris Zones server pools, the pool members must not be a mixture of Oracle Solaris 10 and Oracle Solaris 11, such server pools are not listed in the vDC creation wizard.

- Server pools have compatible storage resources to be used as root disk and volumes.
- Compatible storage resources for saving vServer metadata.

To add a server pool to a vDC, you must verify that each network in the network domain is connected to every server in the server pool when creating a server pool. See Chapter 21, "Server Pools" for more information about how to create server pools.

> **Note:** For Oracle VM Server for SPARC based server pools, the option to select the SR-IOV enabled network interfaces to connect to the network is not supported for vDCs. VLAN tagging support for networks configured with VLAN ID is not used in the vDC context.

Setting Up Storage Resources

The vDC inherits the storage resources allocated for the server pool. The virtualization type of the server pool defines all the possible types of storage resources available for the vDC. Oracle VM Server for x86 based server pools can have the following type of storage resources:

- Oracle VM Storage Repositories
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

For Oracle VM Server for x86 server pools, the Oracle VM Storage repositories must be associated with the server pool.

The Oracle Solaris Zones server pools can have the following types of storage resources:

- NAS libraries
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

The Oracle VM Server for SPARC server pools can have the following types of storage:

- NAS libraries
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

You require the Oracle Solaris Zones and Oracle VM Server for SPARC server pool to be associated with NAS libraries if they are associated only with Block Storage libraries. NAS libraries are required for storing guest metadata details.

When you create a vDC, you allocate the storage resources to be used for all the accounts created in the vDC. You can select storage resources to be used for root disks of the virtual servers (vServers) and for volumes. Volumes are additional storages that you can attach to the vServers. For volumes, the storage is allocated from the storage

libraries that are associated with the server pools. A cloud user might also import external volumes into the account.

Setting Up Network Resources

Oracle Enterprise Manager Ops Center provides comprehensive network management options that enables secure management of the virtual datacenter.

The server pools must be associated with the same user-defined network domain to be added to the vDC. Each network in the user-defined network domain must be connected to every server in the server pool.

Understanding Network Domain Implementation

The network domain is a logical grouping of related networking components in Oracle Enterprise Manager Ops Center. The networks are built on the following type of fabrics:

- **Fully-managed switched fabrics**

  The switches and the fabrics they support are discovered and managed in Oracle Enterprise Manager Ops Center and you can create VLANs or InfiniBand partitions. The network domain created from fully-managed fabrics provides the option to create dynamic private networks.

- **Host-managed fabrics**

  Set VLAN IDs to unmanaged Ethernet fabrics and make it host-managed fabrics. Though the switches are not managed in Oracle Enterprise Manager Ops Center, you can assign VLAN IDs to the Ethernet fabrics. You must enable the VLAN IDs manually on the switch ports connected to the hosts in that fabric. Host-managed fabrics allow you to create dynamic private networks.

  Oracle Enterprise Manager Ops Center uses the VLAN IDs set for the creation of dynamic private networks.

- **Unmanaged fabrics**

  Neither the switches are managed nor the VLAN IDs are available in Oracle Enterprise Manager Ops Center to create the network. The networks are declared or discovered while discovering an asset in Oracle Enterprise Manager Ops Center. The network domain created out of these networks does not provide the option to create dynamic private networks. You can set aside the networks created on the unmanaged fabrics as private while creating a network domain.

  You must ensure that these networks assigned as private are not in use and there is no routing between the selected networks and other networks. These private networks are defined as static private networks. For vDC, a network domain with unmanaged fabrics can be assigned to a server pool only if there are static private networks available in it.

Dynamic Private Networks

Each dynamic private network created is allocated a subnet address according to the standards specified in IETF RFC1918. While creating network domains, you can select the option to select the fabric on which you want to create the network, limit the number of networks to be created on the fabric, and the network addresses to be excluded from use.

The number of private networks that you can create on a fully-managed fabric depends on the type of physical fabric. Each port on an Ethernet switch can support 128 logical fabrics. Each partition on an InfiniBand switch can support 32000 logical

fabrics. The number of logical fabrics determines how many private networks you can create. For an InfiniBand fabric, each P-key can support a private network. For an Ethernet fabric, each VLAN ID can support a private network.

When you create a network domain, you can set the network creation limit for each selected fabric. This determines how many private vNets that can be created in an account. You can spread the number of private networks across all of the accounts in the vDC.

Plan Your Network Domains

The network setup must be implemented in such a way that the cloud user has the required network resources available to allocate to the virtual servers created in the accounts. The cloud user must be able to create private networks. The networks must be created upon either fully-managed switched fabrics, host-managed fabrics or unmanaged fabrics.

Create a network domain, assign the managed networks, and associate it with the server pool which must be added to the vDC. The vDC inherits the network domain from the server pool. From the available networks in the network domain, attach the required networks to the server pool. These managed networks form the public external networks for the vDC. You can assign these networks to the accounts created in vDC.

Note: All the server pools in the vDC must be associated with the same user-defined network domain.

The vDC enables cloud users to create virtual machines and run applications. The cloud user provides networking to its virtual machines from the public networks that are allocated to each account or by creating private networks. The cloud user can use the public networks, or create a private vNet for use in the account. When the cloud users create a private vNet, either a dynamic private network is created or the static private network is available for use in that account.

Note: You must allocate the IP addresses in the public networks and static private network so that cloud users can use the IP addresses.

When the cloud user creates a private vNet, it is listed under the network domain of the vDC.

Avoid to delete any EoIB network resource when a cloud user has just deleted a vServer. When you delete a EoIB network resource, wait for approximately 5 minutes for the VLAN maps to be refreshed automatically.

See Chapter 17, "Networks for Virtualization" for more information about managing fabrics, creating network domains, and private networks.

CPU Oversubscription

You can oversubscribe the CPU resources allocated to a vDC for an increased utilization of the resource. CPU oversubscription is applicable only for vDCs based on Oracle Solaris Zones and Oracle VM Server for x86 virtualization technologies. For Oracle VM Server for SPARC, this is not applicable as there is a one to one relationship between vCPU and physical CPU thread.

You define two parameters that define the CPU oversubscription:

- The ratio of number of virtual CPUs (vCPUs) to physical CPU Threads.

  For example, if the vCPU to physical CPU threads is set to 2, each virtual CPU receives at least 50% of the cycles of a physical CPU thread.

  Note: As the CPU oversubscription ration increases, the performance might be affected, but the utilization of the CPU resources improve. The CPU oversubscription ratio that you might want to use is at most 3:1. At extremely high ratios, the risk of instability of the system increases.

- CPU cap that defines the maximum share of the physical CPU thread's cycle that can be allocated to a vCPU. This parameter is only applicable for vDCs based on Oracle VM Server for x86 virtualization technology.

  For example, if the CPU Cap is set to 50, a vCPU can use up to 50% of the cycles of a physical CPU thread.

Provide appropriate CPU cap for the ratio of vCPU to physical CPU thread to get accountable vServer performance and balanced system.

Creating a Virtual Datacenter

Before you create the vDC, you must have your server pools, network, and storage set up. When they are set up properly, you can select the required resources for the vDC.

To Create a Virtual Datacenter

1. Select vDC Management in the Navigation pane.
2. Click **Create Virtual Datacenter** in the Actions pane.

   The Create Virtual Datacenter Wizard is displayed.
3. The first step provides a introduction to the vDC and the prerequisites for creating a vDC.

   You can select to skip this step in the future when you create a vDC again. Click **Next** to specify identification details for the vDC.
4. In the Specify Virtual Datacenter Details, enter the following information:
   - Provide a name and description for the vDC.
   - Enter tags for better identification and classification of the vDC in Oracle Enterprise Manager Ops Center.
   - Select Password Required to force setting vServer credentials when creating vServers in the vDC.

     When this option is enabled, cloud users must specify a root password, with or without an SSH key. Additionally to the root password, cloud users must specify either an SSH key or a remote user credentials when creating vServers using Oracle Solaris 11 OS in a vDC based on Oracle Solaris Zones or Oracle VM Server for SPARC.

   Click **Next** to select the server pools.
5. Select a server pool from the list.

6. Select the Add more compatible server pools option to add other compatible server pools to the vDC. Then, use the Left and Right arrow keys to select the server pools that must be assigned to the vDC.
7. Select Oracle SuperCluster Support if the vDC is in an Oracle SuperCluster.

   When selecting the Oracle SuperCluster Support option, specific actions for Oracle SuperCluster are enabled or disabled in the vDC.

   Click **Next** to continue.
8. For Oracle VM Server for SPARC based vDCs, select one or more boot networks from the list. The list contains the available public networks, you must select at least one network.

   Note: Boot networks are required for vServer OS deployment. You must also add at least one boot network to each account in the vDC for creating vServers in the account. You can also assign a boot network to an account and use it as a regular public network for that account.

   Click **Next** to specify the vCPU sizing.
9. Define the following values for vCPU sizing:
   - Enter the ratio of number of vCPUs to a physical CPU thread. The value must be greater or equal to 1.0. Decimal values are also supported. The total number of vCPUs can be higher than the number of physical CPU threads. You can run more vCPUs than the existing physical CPU threads through timesharing of the CPU cycles.
   - Specify the CPU cap which is the maximum share of the physical CPU thread's cycle that can be allocated to a vCPU. The CPU cap parameter is only applicable for vDCs based on Oracle VM Server for x86 virtualization technology.

   The following information is displayed in the vCPU sizing screen:
   - **Total number of vCPUs**: The total number of vCPU in the vDC for the updated vCPU to physical CPU Threads ratio.
   - **Avg memory per vCPU**: The average memory per vCPU in GB. The total memory available for the vDC by the total number of vCPUs. When you want to use the CPU and memory resources to the full extent, then the vServers must be using this amount of resources.

   When vServers are created in a zone based vDC, the Fair Share Scheduler is the default scheduler. The physical CPU allocation for each vServer is provided by the Fair Share Scheduler.

   This is not applicable for vDCs based on Oracle VM Server for SPARC virtualization technology.

   Click **Next** to configure the storage resources.
10. Select the root disk storage type from the list. Designate one of the following storage types for root disks:
    - Oracle VM Server for x86 based vDC
      - Oracle VM Storage Repositories

  - Dynamic Block Storage Libraries
- Oracle Solaris Zone based vDC
  - NAS Libraries
  - Block Storage Libraries that include both Dynamic and Static
- Oracle VM Server for SPARC based vDC
  - NAS Libraries
  - Block Storage Libraries that include both Dynamic and Static

11. Select the storage resources to be used for root disk from the list. The list contains the storage resources associated to the selected server pools and storage type.
12. Select the volume and template storage type from the list. Designate the following storage types for volumes and templates:
    - Oracle VM Server for x86 based vDC
      - Oracle VM Storage Repositories
      - Dynamic Block Storage Libraries
    - Oracle Solaris Zone based vDC
      - NAS Libraries
      - Block Storage Libraries that include both Dynamic and Static
    - Oracle VM Server for SPARC based vDC
      - NAS Libraries
      - Block Storage Libraries that include both Dynamic and Static
13. Select the storage resources to be used for volumes and templates from the list. The list contains the storage resources associated to the selected server pools and storage type.

    Click **Next** to view the summary.
14. The summary lists the total resources that are allocated for the vDC. The summary lists the following:
    - **vDC Name**: The name given to the vDC.
    - **Virtualization type**: The virtualization technology of the server pool.
    - **Oracle SuperCluster Support**: Flag indicating if the vDC must behave as part of an Oracle SuperCluster.
    - **Total Physical CPU Threads**: The physical CPU threads available from all the virtualization servers in the selected server pools for the vDC.
    - **Required Password**: Flag to indicate if the cloud user must specify vServer credentials when creating vServers in the vDC.
    - **Total Memory (GB)**: The total RAM allocated to the management domains of the selected server pool. For example, Oracle Solaris Zones based vDC, it is the total memory of the global zones in the server pool. For Oracle VM Server for x86 based server pools, it is the total memory of the Oracle VM Servers in which the virtual machines run.
    - **Total Disk Space (GB)**: The sum of all the disk spaces allocated from the storage libraries associated with the selected server pools.

- **vCPU to Physical CPU Threads ratio**: The number of vCPUs to physical CPU Threads.
- **vCPU cap**: The maximum share of the physical CPU Thread's cycle that can be allocated to a vCPU. This parameter is only available for vDCs based on Oracle VM Server for x86 virtualization technology.

Confirm the vDC configuration and click **Finish** to create the vDC.

Managing a Virtual Datacenter

Perform the following actions for managing a vDC:

- Updating the Configuration
- Managing Resources
- Deleting a Virtual Datacenter

Updating the Configuration

Update the following details of vDC:

- vDC identification and tags.
- Assign more server pools.
- Assign more storage for volume or root disks. You cannot remove the storage that is currently used by the vDC.
- Modify the vCPU sizing.
- Disable or enable the Password Required option.
- Disable or enable the Oracle SuperCluster Support option.
- Add public networks to be used as boot networks. This option is only available for Oracle VM Server for SPARC virtual datacenters.

Note: After upgrading from Oracle Enterprise Manager Ops Center version 12.1 to version 12.2, cloud administrators must update the configuration of all Oracle VM Server for SPARC vDCs created in a 12.1 version to select a public network as a boot network. Cloud administrator might also need to update the account configuration to add at least a boot network to the account. If a boot network is not added to the account, cloud users won't be able to create new vServers.

Server templates created in a 12.1 version are not usable after upgrading to 12.2, cloud users need to create new server templates for creating new vServers in the vDC.

Select a vDC and click the **Edit Virtual Datacenter** option in the Actions pane. Modify the required attributes and complete the wizard.

To add storage, networks, or servers to the vDC, you must first assign the resources to the server pool of the vDC. When adding networks or servers to a sever pool, you must verify that each network in the network domain is connected to every server in the server pool. See Attaching Networks and Adding Virtualization Hosts sections of the Server Pools chapter for more information.

Managing Resources

Oracle Enterprise Manager provides a complete view of the resources in the vDC and monitors the utilization of the virtualization servers, network, and storage. The cloud administrator gets a good picture of the utilization of the resources which helps to manage the under-utilized and over-utilized resources.

The Dashboard shows the total resources available for a vDC and the amount allocated for different accounts. The committed resources indicate the total resources entitled to all the accounts in the vDC.

For example, if you have 4 physical threads in the vDC and set the vCPU to physical CPU threads ratio to 2, then you have 8 virtual CPUs. You can oversubscribe the vCPUs and allocate 5 vCPUs to one account and another 5 vCPUs to another account. The total committed resources to all the accounts in the vDC is 10 vCPUs.

The Network tab displays detailed network usage in the vDC. The page lists the fabric and public networks of the vDC, and all of the private networks created by the cloud user. The fabric on which the private vNets are created are also displayed in this page.

Figure 22–1 Network Usage in a vDC

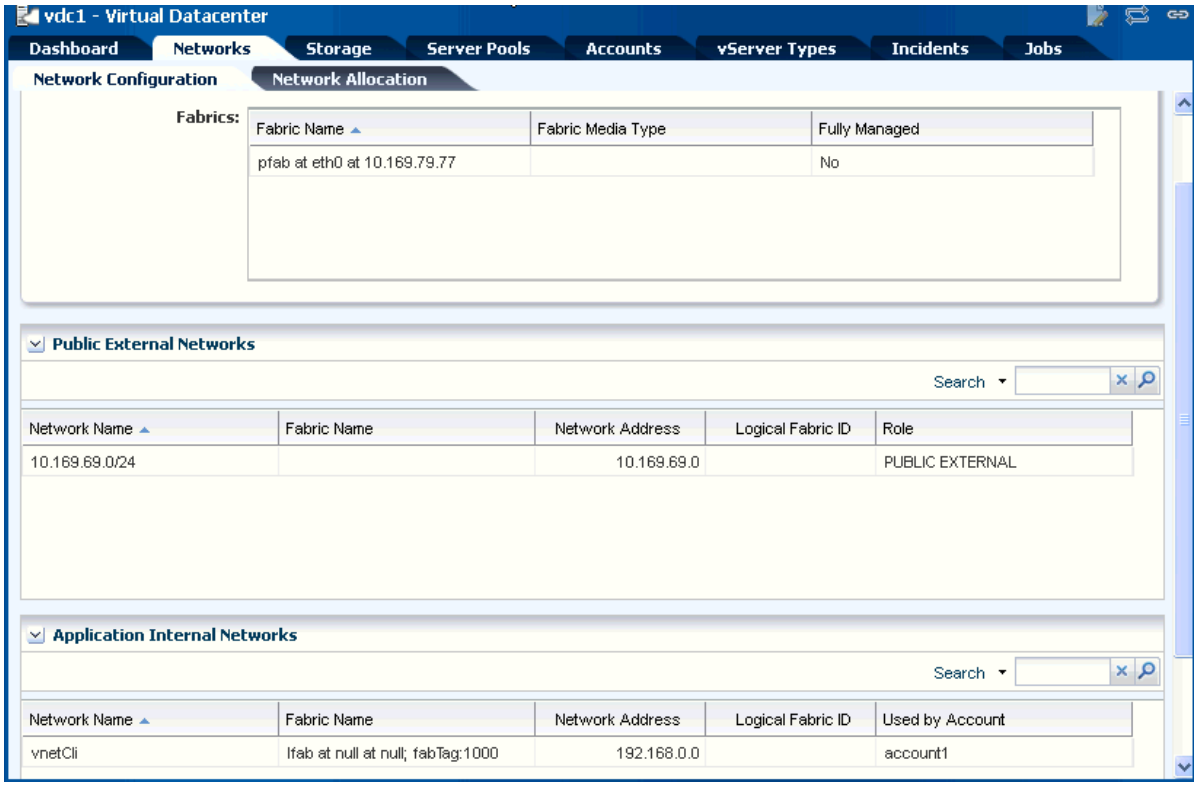

The Server Pools tab shows the usage of the server pool resources. This lists the server pools that are in the vDC. The total and committed resources of vCPU, memory and storage are displayed. This helps you to plan whether you must add more resources to the vDC.

Figure 22–2 vDC Server Pool Usage

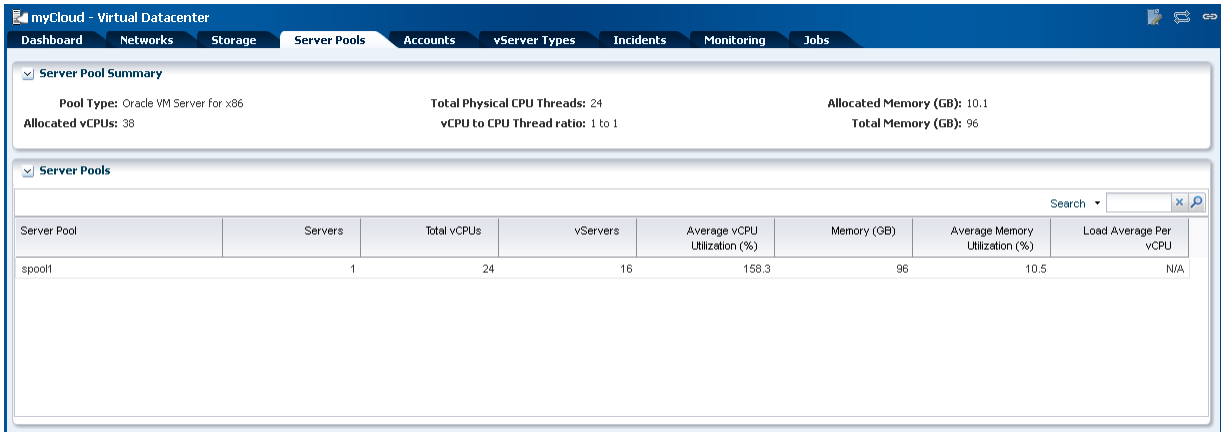

The Incident tab lists all the incidents reported from all the infrastructure components in the vDC. For each incident, view the source of the incident which helps you to identify the component and the source of the problem. The incidents also covers the warnings issued when the vDC resources are over-utilized.

Figure 22–3 vDC Incidents Tab

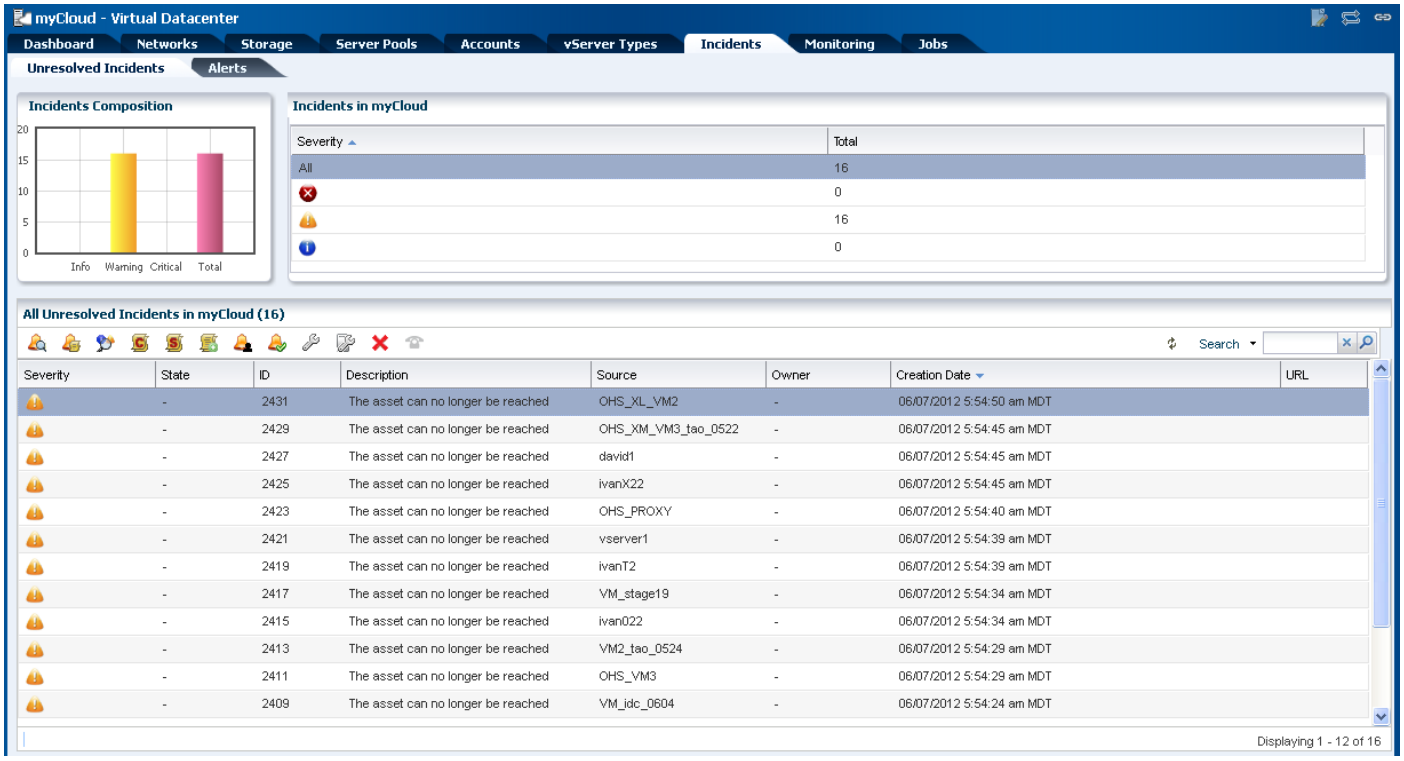

Deleting a Virtual Datacenter

You can use the **Delete Virtual Datacenter** action to delete a virtual datacenter and release all of the resources attached to it. Before deleting a virtual datacenter, you must delete all the accounts in the vDC.

When you delete an account, ensure that you do not have running virtual servers or running jobs creating resources in the account. Once the account is deleted, all of the

suspended or shut down vServers are deleted and the resources are released back to the vDC.

Creating and Managing Cloud Users

You can create cloud users and provide access to the accounts so that they can utilize the computing resources allocated to them. The cloud user is associated with an existing user on the OS under the Enterprise Controller. The ways to create this operating system user varies by the type of OS and type of name services, for example, file based, NIS, or LDAP.

Use the following Oracle Solaris OS command for a simple local file based user account.

```
useradd <cloud user name>
passwd <cloud user name>
```

Adding and Removing Cloud Users to an Account

You can add cloud users to the account and give them the right to use the resources allocated to an account. You can add the cloud users when you create an account or using the option **Add Users**.

To Add a Cloud User

1. Select the account and click **Add Users** in the Actions pane.

   The Add Users Wizard is displayed.

2. The lists of available cloud users and cloud users assigned to the account are displayed. Select users you want to assign to the account and move them to the Assigned Users.

   Click **Next** to view the summary.

3. View the summary and click **Finish** to add the user to the account.

To Remove a Cloud User

1. Select the account.
2. Select the **Users** tab in the center pane.

   The list of users assigned to the account are displayed.

3. Select the cloud user whom you want to remove.
4. Click the **Delete** icon.
5. Click **Remove** to confirm the action.

Managing Roles for Cloud User

By default, the cloud user is provided with default privileges to use all the profiles, plans, and vServer Types available in the selected vDC. Ensure that the cloud user has appropriate privileges to the vServer types. Otherwise, the users cannot create vServers in their environment.

Refer to the *Oracle Enterprise Manager Ops Center Administration Guide* for more information about managing user roles.

Figure 22–4 shows the step to set the privileges for a cloud user.

Figure 22–4 Specify Privileges

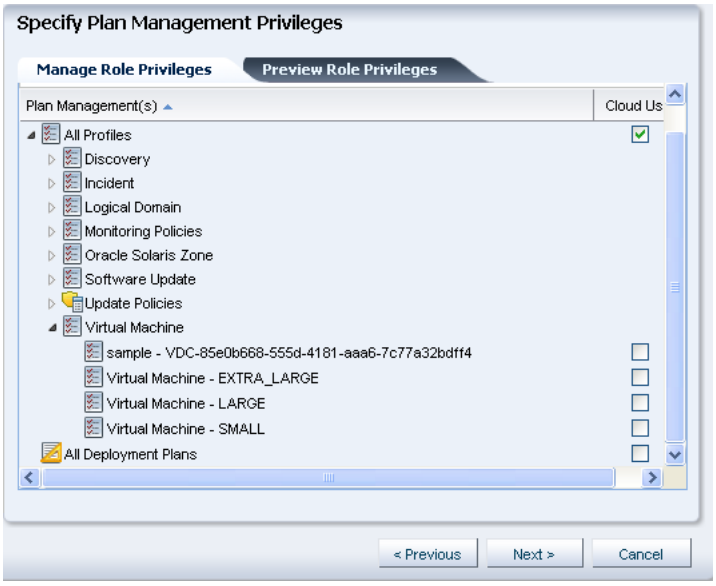

Creating Accounts

An account entitles designated cloud users the right to use computing, network, and storage resources of vDC. The account provides the required capabilities to manage these resources.

The prerequisites for creating an account are:

- Estimate the resource quotas to be allocated for the account
- Identify the cloud users to be assigned to the account

The quota for vCPU, memory and storage resources are defined during account creation. The Resource Quota Information display in the account wizard creation indicates how much of the corresponding vDC resources are subscribed. The resource usage indicates whether the vDC resources are oversubscribed or undersubscribed.

You can create a maximum of 4096 private vNets in an account. You can set the limit of number of private vNets that can be created in an account.

During account creation, the public networks that are available in the vDC are listed. You can set the number of public IP addresses allocated to the account from the resource. A cloud user can assign the public IP address to a vServer, as needed.

Assign the cloud users to the account during account creation or separately. Cloud users have access to only specific accounts. As a cloud administrator, you can manage the access of the cloud users to all the accounts.

Creating an Account

You provide an entitlement to the virtual resources for an account. You allocate the resources from the vDC to an account. The resource allocation for all the accounts in a vDC can be more than the actual resources in a vDC. This oversubscription of the resources must be identified and planned for a vDC. You must configure the virtual resources for an account properly and update the resource configuration when the requirement increases.

As a cloud administrator watch the resource usage and properly configure the resources for an account.

To create an account:

1. Select the vDC in which you want to create the account.
2. Click **Create Accounts** in the Actions pane.

   The Create Account Wizard displays.
3. The first step provides a introduction to accounts and the prerequisites for creating an account.

   You can skip this step in the future when you create an account again. Click **Next** to specify identification details for the account.
4. Enter the name and description for the account.

   Add tags for categorizing and identifying the account. Click **Next**.
5. Specify the resource quotas for vCPU, memory, and storage for the account.

   The available resources from vDC are allocated to the vDC. You can oversubscribe the resources. The Resource Quota Information displays whether the vDC resources are undersubscribed or oversubscribed. This gives a complete picture of the vDC resource usage.
6. Select the number of private vNets a cloud user can create in an account.

   The maximum number of private vNets that you can create in an account is 4096.
7. Select the public networks from the list. The list contains all available public networks from the vDC that you can allocate to the account.

   For each network, the available number of IP addresses is displayed. Enter the number of IP addresses to be allocated to the account. Click **Next**.

   Note: For vDCs using Oracle VM Server for SPARC as virtualization technology, the list includes the boot networks of the vDC. A boot network must be added to the account for vServer OS deployment.
8. Select the cloud users who can access the account. You must assign a cloud user while creating the account. Thereafter, use **Add Users** option to add more users.
9. Review the account information and click **Finish** to create the account.

Managing Accounts

You have the following options to manage the account configuration in a vDC:

- Update account resource configuration
- Assign cloud users to accounts
- Delete an account

Updating Accounts

As a cloud administrator, you can modify the identification, and resources allocated to an account. When you want to allocate more resources to an account, you can modify by selecting an account and click **Update Account** in the Actions pane. Modify the resources and complete the wizard.

You can remove the resources allocated to an account if they are not used by the vServers. Similarly, you can remove unused public networks from an account or reduce the limit of public IP addresses if they are unused by vServers.

Note: For accounts in vDCs using Oracle VM Server for SPARC as virtualization technology, a boot network must be added to each account. If a boot network is not added to the account, cloud users won't be able to create new vServers.

Managing Account Resources

When an account does not have enough resources, the cloud user receives notifications that they cannot create vServers due to unavailability of resources. Also, when there are no physical resources available from the corresponding vDC, the cloud user cannot create vServers even if the account resources show that they are not fully used up.

For example, assuming that you have 100 GB physical storage allocated to a vDC and an account with 125 GB storage. The vDC storage is oversubscribed. The cloud users create vServers in the account and use the storage. When the storage usage exceeds 100 GB, the cloud user cannot create vServer as the vServer job fails with the message that enough space is not available.

Therefore, the cloud administrator must the watch the resource usage and add more resources to the vDC.

Maintaining OS Images

For vDCs based on Oracle VM Server for SPARC virtualization technology, you must manage OS provisioning profiles and OS configuration profiles that the cloud user can use for creating server templates to deploy vServers in their environment.

Note: For handling IPS package dependencies properly, when a server template for Oracle Solaris 11 uses an OS provisioning profile that contains a Solaris 11 update profile for a package, the latest version of the package gets installed regardless of the version selected in the Solaris 11 update profile.

The Oracle Solaris 10 ISO or FLAR images must be imported in the NAS libraries associated with the server pools placed in the vDCs.

Configure your Oracle Solaris 11 Software Update Library in the Enterprise Controller to synchronize with the Oracle Solaris 11 Package Repository. The cloud user can use all the profiles for Oracle Solaris 11 OS.

Ensure to configure DHCP in your environment so that the cloud user can install Oracle Solaris 10 ISO image.

OS provisioning profiles that are based on Oracle Solaris 10 FLAR images and Oracle Solaris 11 OS, use WAN boot and do not require DHCP configuration.

See Chapter 5, "Software Libraries" for more information.

Deleting an Account

Ensure that you do not have running virtual servers or running jobs creating resources in the account. Use the option Delete Account to remove the account. Once the account

is deleted, all of the suspended or shut down vServers are deleted and the resources are released back to the vDC.

When there are vServers in running state, the following message is displayed:

Figure 22–5 Delete Account

Creating and Managing vServer Types

vServer Type are profiles of virtual machines that defines the computing resources such as virtual CPU, memory, and storage size. A cloud user can use this to implement and create vServers.

> **Note:** For vDCs based on Oracle VM Server for x86, the storage size is not defined in the vServer type. The storage size is defined from the server template of the vServer.

A vServer type created is available for all the accounts in a virtual datacenter.

As a cloud administrator, you can capture the user requirements and create different vServer Types. Cloud users can use this to create vServers. By default, there are three system-defined vServer types that define the computing resources.

System-Defined vServer Types

The three system-defined vServer types provided for all vDCs are:

- Extra Large
- Large
- Small

The vServer type configuration varies depending on the virtualization technology on which the vDC is constructed.

Figure 22–6 vServer Types

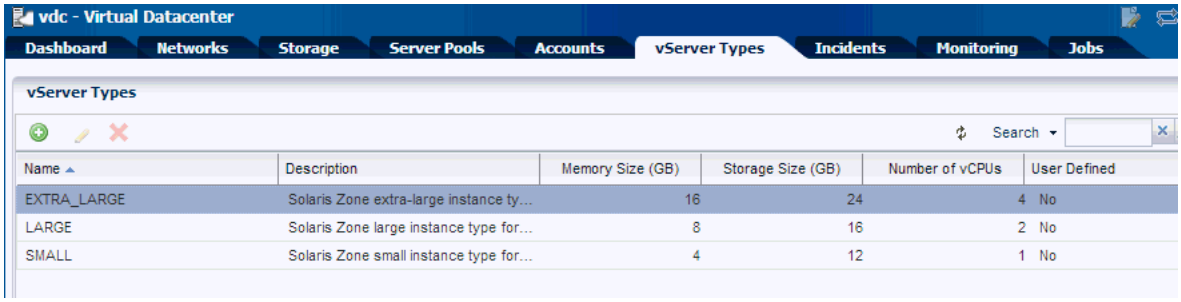

For Oracle VM Server for x86 based vDCs, the vServer Type is displayed as shown in Figure 22–7.

Figure 22–7 vServer Type for Oracle VM Server for x86 based vDCs

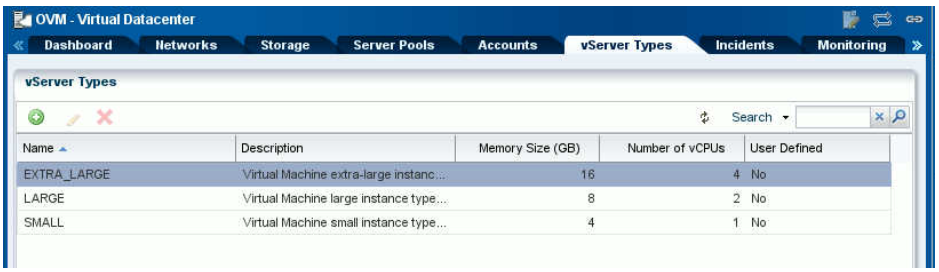

Creating a vServer Type

When you create a vServer type, the VM hosting details display the following information in the wizard based on the resources defined:

- The number of virtualization servers in the vDC that have sufficient physical resources to host a vServer with the selected resources.
- An estimation of number of vServers that can be hosted with the total number of physical resources of the vDC.
- A warning when the current value of the memory size exceeds the selected storage size.

To Create a vServer Type

1. Select the vDC and click **Create vServer Type** in the Actions pane.

   The Create vServer Type Wizard is displayed.

2. Enter the name and description for the vServer type.

   Add tags for identification and classification of the vServer type. Click **Next**.

3. Specify the vCPU, memory and storage resources for the vServer type.

   As you enter the values, the VM Hosting displays the following information:

   - The number of virtualization servers that have sufficient resources in the vDC to host a vServer of this type.
   - The number of vServers that the vDC can host with the total number of physical resources of the vDC.

A warning message is displayed when the current memory size is more than the storage size. The existing vServer types are displayed with the configuration details. This helps to avoid creating duplicate vServer types.

For vDCs based on Oracle VM Server for x86, the storage size is not defined in the vServer Type.

Provide the values for the resources and click **Next**.

4. Review the information provided and click **Finish** to create the vServer type.

Updating vServer Types

Modify the following details of a vServer type:

- Name and description
- Modify, or add new tags
- Modify the resource configuration

Use the option **Update vServer Type** to launch the Update vServer Type Wizard. Modify the configuration and complete the wizard.

Deleting vServer Type

You cannot the delete the system-defined vServer types. You can delete only the vServer types that you have created.

Overview of Cloud Users

A cloud user with access to an account is entitled to manage and use computing, network, and storage resources allocated in a vDC within the limits of the account quotas.

Cloud users can create and manage the life cycle of vServers for their applications. Creation and management of vServers involve the setup of virtual resources, the management of virtual resource workloads, and understanding application requirements.

Cloud Users can manage the following virtual resources:

- **Virtual Networks (vNets)**: Used to connect and restrict network access of vServers.
- **Server Templates**: Designates the operating system and how it is installed while creating a vServer.
- **Virtual Storage**: Includes volumes that you can attach to vServers, and snapshots to capture the current state of a volume for different purposes.
- **vServers**: An entity that provides the outward interface of a stand-alone operating system. A vServer has its own identity, local storage, interfaces, and configuration that exist for the full lifetime of the vServer.

Availability and management of some virtual resources vary depending on the virtualization technology of the vDC or resource type supported by the vDC.

Oracle Enterprise Manager Ops Center offers to cloud users the option to perform their tasks using its browser interface or through:

- **APIs**: A Web service API and a Java API to programmatically manage the allocated resources in a vDC account for automation or integration purposes.

- **CLI**: To manage allocated resources in a vDC account from a text-based console that can be also used for automation or integration purposes.

The following sections in this guide describe the tasks a cloud user can perform in Oracle Enterprise Manager Ops Center using the UI. For more information about the use of APIs and CLI, refer to *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide*.

Roles for Cloud User Tasks

Table 22–3 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 22–3 Cloud User Tasks and Roles

Task	Role
View Account Quotas and Virtual Resources	Cloud User
Create and Manage vServers	Cloud User
Create and Manage Virtual Networks	Cloud User
Create and Manage Volumes	Cloud User
Create and Manage Snapshots	Cloud User
Create and Manage Server Templates	Cloud User

Actions Available for a Cloud User

A cloud user can perform different actions to use and manage the allocated resources in a vDC account. This section lists the actions a cloud user can perform for each type of resource and also provides information about the actions that are available for the different vDC account types.

- View account quotas and virtual resources
- Create and manage vServers
- Create and manage virtual networks
- Create and manage volumes
- Create and manage snapshots
- Create and manage server templates

Based on the virtualization technology on which the vDC and its accounts are created, there might be differences in the availability of the options in the cloud user view. The following tables provide a detailed list of options that are available for a cloud user when the vDC is based on different virtualization technology.

Account

Table 22–4 Account Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
View account quotas and resources	Yes	Yes	Yes
View account Incidents	Yes	Yes	Yes
View Account Jobs	Yes	Yes	Yes

vServers

Table 22–5 vServers Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create vServer	Yes	Yes	Yes
Stop and start vServer	Yes	Yes	Yes
Pause and Resume vServer	Yes	No	No
Update vServer	Yes	Yes	Yes
Launch Virtual Console	Yes	No	No
Shutdown and Start All vServers	Yes	Yes	Yes
Attach and Detach Volume	Yes	Yes	Yes
Enable and Disable HA	Yes	Yes	Yes
Delete vServer	Yes	Yes	Yes

Networks

Table 22–6 Networks Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create private vNet	Yes	Yes	Yes
Update private vNet	Yes	Yes	Yes
Delete private vNet	Yes	Yes	Yes
Allocate vIP	Yes	Yes	Yes
Deallocate vIP	Yes	Yes	Yes
View vIPs	Yes	Yes	Yes

Volumes and Snapshots

Table 22–7 Volumes and Snapshot Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create Volume	Yes	Yes	Yes
Create Volume from Snapshot	Yes	No	No
Import Volume	Yes	No	No
Update Volume	Yes	Yes	Yes
Delete Volume	Yes	Yes	Yes
Create Snapshot	Yes	No	No

Server Templates

Table 22–8 Server Templates Option Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Upload Server Template	Yes	No	No
Create Server Template	No	Yes	No
Update Server Template	Yes	Yes	No
Delete Server Template	Yes	Yes	No
Register and Unregister Server Template	Yes	Yes	No
Save vServer as Template	Yes	No	No

Distribution Groups

Table 22–9 Distribution Groups Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create Distribution Group	Yes	Yes	No
Update Distribution Group	Yes	Yes	No
Delete Distribution Group	Yes	Yes	No

Location of Account Quotas and Virtual Resources in the User Interface

Account quotas limit the creation of new virtual resources in the account. A cloud user must be aware of the account quotas and manage the workloads of the virtual resources.

To see the quotas and virtual resources of an account, expand vDC Management in the Navigation pane.

This lists all the accounts to which a cloud user has access under vDC Accounts.

Figure 22–8 vDC Management

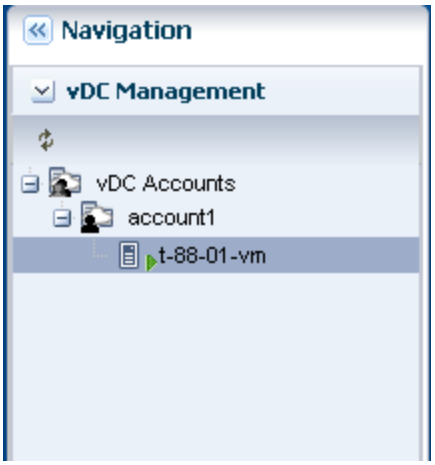

You can select a listed account to display general account information and quotas in the Dashboard tab of the center pane. The Dashboard tab displays general information about the selected account and a summary of the virtual resources status and usage.

Account resources details appear in the other tabs across the center pane:

- **Networks**: Shows the vNets quotas. Lists and displays information for each private vNet and public network available for the account. Also displays the actions bar for the actions that you can perform for a vNet.
- **Storage**: Shows the storage quotas. Lists and displays information for each vServer root disk, volume, and snapshot available for the account. Also displays the actions bar for the actions that you can perform for a storage resource.
- **vServers**: Shows the CPU quotas. Lists and displays information for each vServer available for the account. Also displays the actions bar for the actions that you can perform for a vServer.
- **Server Templates**: Lists and displays information for each server template available for the account. Also displays the actions bar for the actions that you can perform for a server template.

Creating vServers

A vServer is an entity that provides the outward interface of a stand-alone operating system that consumes CPU, storage, and memory resources. A vServer has its own identity, local storage, interfaces, and configuration that exist for the full lifetime of the vServer.

You determine the creation of new vServers according to the account quota limits and applications requirements.

Before You Begin

When creating a vServer, the following account resources are required:

- **A vServer type**: vServer Types are vServer profiles that defines the computing resources such as virtual CPU, memory, and storage size. After memory, storage, and number of vCPUs are defined for the vServer, you can select the best suitable vServer type from those available for the account. vServer types are visible to cloud users during the vServer creation process. Contact your cloud administrator when you require a vServer type.
- **A server template**: Server templates designate the OS and how it is installed while creating a vServer. You can either select a server template from those that exist for the account or create a new server template. Actions to use or create server templates vary based on the virtualization technology used in the vDC. For more information about server templates, see Creating Server Templates.
- **One or more virtual networks**: For the vServer network connectivity you must choose one or more vNets from the available vNets or create new ones before creating a vServer. vServers are only assigned to virtual networks at vServer creation time. For more information about vNets, see Creating vNets.

> **Note:** For accounts in vDCs using Oracle VM Server for SPARC as virtualization technology, at least a boot network must be assigned to the account. If a boot network is not added to the account, cloud users won't be able to create vServers. Contact your cloud administrator if you receive an error message when creating new vServers.

Depending on the needs of the user and the virtualization type of the vDC, a cloud user can also specify the following resources:

- **One or more volumes**: Volumes provide additional storage for vServers. You might be requested to attach volumes to vServers at vServers creation time, in that case, volumes must exist before creating the vServers. You can also attach volumes after creating the vServer. For more information about creating volumes, see Creating Volumes.
- **A distribution group**: Distribution groups are only available for vDCs based on Oracle VM Server for SPARC virtualization technology. Distribution group enforces that two vServers are not running in the same Oracle VM Server. If you choose to assign a vServer to a distribution group, the distribution group must exist before creating the vServer. For more information about distribution group, see Creating Distribution Groups.

Additionally to the resources listed, a cloud user specifies the following values when creating a vServer:

- **IP address assignment method**: You can select from wo different methods for IP address assignment to vServers:
  - **Static method**: Used to assign a specific IP address to a vServer. When using this method you must allocate in advance an IP address from the selected virtual networks. You must consider using this method when creating a single vServer at a time.
  - **Automatic method**: This method dynamically assign an IP address from each selected virtual network. When creating a multiple vServer at a time, only the use of the automatic IP address assignment method is allowed.
- **Number of vServer**: You can create single or multiple vServers at a time. When creating multiple vServers at a time, vServers are created with the same configuration and a suffix is added to each vServer's name. You must also

consider the IP address assignment method when defining the number of vServers.

- **High Availability**: When a high availability is enabled, the vServer is available uninterruptedly. The vServer is not shutdown when the backend infrastructure fails and it is migrated to another server and started. When high availability is not enabled, the vServer starts only after the backend infrastructure is available. Cloud administrator takes care of server pool settings to provide high availability feature in a vDC.
- **vServer control mechanism**: You can supply credentials for remote SSH root access to the vServer. An option available as an alternative secure method is to use a public key to authenticate SSH root access to the vServer without the use of a password. This option is available to all cloud users, other options to create local user account or provide password for remote SSH access are also available.

Creating a vServer

Oracle Enterprise Manager Ops Center provides a wizard that walks you through a series of steps to collect information and resources to create vServers.

To create a vServer:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create vServer** in the Actions pane.

   The Create vServer Wizard is displayed.
4. Enter the following information in the **vServer Details** step:
   - Name and description for the vServer.
   - Tags for better identification and classification of the vServer.
   - Number of vServers to create.
   - Select the High Availability Support option to enable it for the vServer.

   Click **Next** to select a server template.
5. Select a server template from the list.

   Click **Next** to select a vServer type.

   Note: Server templates created in a 12.1 version are not usable in 12.2. After upgrading from Oracle Enterprise Manager Ops Center version 12.1 to 12.2, you must create new server templates to create new vServers for Oracle VM Server for SPARC vDCs created in a 12.1 version.
6. Select a vServer type from the list and click **Next**.
7. (Optional) When the step to select volumes is present, select one or more volumes from the Available Volumes list. Use the arrow keys to move the selected volumes to the Attached Volumes list.

   If not required, skip this step as the volumes can also be attached after creating a vServer. Click **Next** to continue.
8. Select one or more vNets from the list.

Click **Next** to select the IP address assignment method.

> **Note:** For accounts in vDCs using Oracle VM Server for SPARC as virtualization technology, at least a boot network must be assigned to the account. If a boot network is not added to the account, cloud users won't be able to create vServers. Contact your cloud administrator if you receive an error message when creating new vServers.

9. Select the IP address assignment method for each vNet and click **Next**.

   When the static method is selected, choose an IP address from the list.

10. (Optional) When the Distribution Group Selection step is present, select a distribution group from the list if required.

    If a distribution group is not required, skip this step. Click **Next** to provide a public key.

11. (Optional) In the vServer Access Control step, you can supply a public key as an alternative to authenticate remote SSH root access to the vServer. By suppling a public key, root user will not be able get remote SSH access without the corresponding private key.

    You can choose one of the following options:

    – Paste the public key directly into the Public Key text area.

    – Click Browse to select a local file containing the public key, and then click **Upload Public Key**. The size of the file containing the public key is limited to 2GB.

    If not required, skip this step and click **Next**.

12. In the Login Credentials step, provide the following information to authenticate the remote SSH access to the vServer:

    - Enter the root password and confirm it.
    - For vDCs based on Oracle Solaris Zones or Oracle OVM for SPARC, you can enter a local user name.
    - For vDCs based on Oracle Solaris Zones or Oracle OVM for SPARC, you can enter a local user name password and confirm it.

    > **Note:** If the Password Required option is enabled for the vDC, you must specify a root password, with or without an SSH key. Additionally to the root password, you might require to specify either an SSH key or a remote user credentials when creating vServers using Oracle Solaris 11 OS.

    You can also skip this step and click **Next** to view the summary if the vDC is configured to allow the creation of vServer without setting credentials.

13. Confirm the vServer information provided in the Summary and click **Finish** to launch the job to create the vServer.

    After the job completes, the vServer is created and listed in the Navigation pane. By default, the DNS and other naming information is taken from the selected vNet or server template and added in the `/etc/resolv.conf` file of the vServer.

If the job fails, you can look at the job details to identify the issue. Once the issue is solved, open the Create vServer wizard to try again. Re-running an existing job for creating vServers is not supported.

Managing vServers

After creating a vServer, you can manage the vServer life cycle by executing the available actions for vServer management.

To locate all the actions available for managing a vServer:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Select the vServer tab in the center pane.
4. You can then select a vServer from the vServers list displayed in the center pane and then choose one of the actions displayed.

Figure 22–9 vServer Management Options

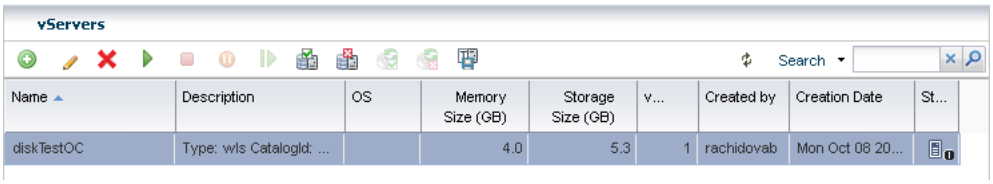

Cloud users can perform the following actions to manage a vServer life cycle. The availability of some actions depends on the cloud infrastructure. See Actions Available for a Cloud User for available actions under different cloud infrastructure:

- **Update a vServer**: To modify the current name, description, memory size, or number of vCPUs of a vServer. You can only modify the memory size or number of vCPUs of a vServer in shutdown or shutdown/detached status. The Update vServer action also allows the creation or deletion of vServer tags.
- **Stop a vServer**: To stop a running vServer. When a vServer is stopped, the guest operating system is shutdown. A stopped vServer can be started later.
- **Start a vServer**: To start a stopped or shutdown vServer. The restarted vServer might not get the same IP address as the original vServer.
- **Pause a vServer**: To suspend a running vServer. This action is only available for vDCs based on Oracle VM Server for x86.
- **Resume vServer**: To start a suspended vServer. This action is only available for vDCs based on Oracle VM Server for x86.
- **Shutdown all vServers**: To shutdown all vServers listed for an account. This action is only available in the Actions pane.
- **Start all vServers**: To start all vServers listed for an account. This action is only available in the Actions pane.
- **Attach vServer volumes**: To attach one or more new volumes to a vServer. Do not stop the vServer to attach volumes. You might require to stop the vServer for vDCs based on zones.
- **Detach vServer volumes**: To detach one or more volumes from a vServer. Do not stop the vServer to detach volumes. You might require to stop the vServer for vDCs based on zones.

- **Enable HA**: To enable high availability of the vServer. This ensures that the vServer is migrated and restarted on another Oracle VM Server when the current Oracle VM Server fails.
- **Disable HA**: To disable high availability of the vServer. If the Oracle VM Server fails, the vServer is shut down and restarted when the Oracle VM Server becomes available.
- **Save vServer as Template**: To save the vServer as a server template and use for new vServer creation. This action is only available for vDCs based on Oracle VM Server for x86.
- **Delete a vServer**: To delete a vServer from a vDC account. Deleting a vServer results in shutting down the vServer followed by the deletion of the vServer.

Creating Server Templates

Server templates designate the operating system and how it is installed while creating a vServer. Server templates are specific to processor architecture of the server pool and virtualization type and can be pre-built images or identifies the OS distribution.

Server templates are loaded into the storage libraries associated with the vDC and cannot be changed later. By default, a server template is bound to a specific account.

The proper server template must exist before creating vServers. Depending on the virtualization type of the vDC account, cloud users can have the following options:

- Use a default server template provided by the system for creating vServers. This is available only for Zones based cloud infrastructure.
- Upload a new server template to be used for creating vServers. This is available only for Oracle VM Server for x86 based cloud infrastructure.
- Create a new server template for installing vServers. This is available only for Oracle VM Server for SPARC based cloud infrastructure.
- Save a server template from an existing vServer to be used for creating vServers, see Managing vServers.

Uploading Server Templates

To upload a server template, you can choose from two different server templates subtypes:

- **Template**: Single virtual machine template that is ready to be deployed into virtualized platforms. Templates can be of format `.tar` or other file types. This option allows you to upload a multi-file template that is stored as a single server template for the account.
- **Assembly**: Collection of interrelated software appliances that can include a configuration of multiple virtual machines with their virtual disks and their inter connectivity. An assembly is contained in a single `.ova` (Open Virtualization Format Archive) file. When uploading an assembly, you can also create snapshots when registering an .ova file with shared virtual disks. Those snapshots are tagged with the template ID and the assembly ID.

Either select the server templates stored in the local host machines or from other locations that you can access. Use FTP, HTTP, or HTTPS protocols to upload the server templates from other locations.

Before uploading a server template, you must:

- Decide on a suitable server template subtype.
- Make sure the file for the server template is of the correct format.
- Check whether the file is accessible when uploading the server template from a file in other location.

To Upload a Server Template

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Upload Server Template** in the Actions pane.

   The Upload Server Template Wizard is displayed.
4. Enter the following information in the Identify Server Template step:
   - Name and description for the server template.
   - Tags for better identification and classification of the server template.

     Click **Next** to specify server template details.
5. Choose a server template subtype.
6. Choose an option to upload the server template file.

   When the URL option is selected, enter the complete URL. When using the subtype template, you can also specify multiple URLs when uploading a multi-file template.

   Click **Next** to view the summary.
7. Confirm the server template information and click **Finish** to create the server template.

Creating a Server Template

The option to create a server template is available only when the vDC is built on Oracle VM Server for SPARC virtualization technology. The server templates define the OS installation media and version to be used during vServer creation.

> **Note:** Server templates created in version 12.1 are not usable in 12.2. After upgrading from Oracle Enterprise Manager Ops Center version 12.1 to 12.2, you must create new server templates to create new vServers in Oracle VM Server for SPARC vDCs created in a 12.1 version.

You can create server templates from existing OS provisioning profiles and OS configuration profiles. The list of profiles is filtered to display only the distribution which are available to vDC. If a required profile is not available, contact your cloud administrator for required OS provisioning profiles to install your vServers. The OS provisioning and configuration profiles are created when importing an Oracle Solaris OS image or when configuring the Oracle Solaris 11 Software Update Library in Oracle Enterprise Manager Ops Center.

Note: For handling IPS package dependencies properly, when a server template for Oracle Solaris 11 uses an OS provisioning profile that contains a Solaris 11 update profile for a package, the latest version of the package gets installed regardless of the version selected in the Solaris 11 update profile.

To Create a Server Template

1. Select the account in the vDC Management section.
2. Click **Create Server Template** in the Actions pane.

   The Create Server Template Wizard is displayed.
3. Skip the introduction and click **Next**.
4. Enter the name and description for the server template.

   If required, enter tags for easy classification and categorization.

   Click **Next**.
5. Select a profile that provisions Oracle Solaris OS, then click **Next**.

   You can install Oracle Solaris 10 or Oracle Solaris 11 OS on the vServer.
6. Select an OS configuration profile from the list, then click **Next**.

   The OS configuration profile collects information about control domain parameters and network configuration.
7. Review the information and click **Finish** to create the server template.

The new template is created and available to create vServers.

Managing Server Templates

Perform the following actions for server templates management:

- **Update Server Template**: To modify the name or description of a server template. This action also allows the creation or deletion of server template tags.
- **Register Server Template**: To register a server template for public use. Registering a server template make the server template available to other accounts in the vDC.
- **Unregister Server Template**: To unregister a server template is accessible for public use. Unregistering a server template make the server template available to only to the vDC account from which the server template was created.
- **Delete Server Template**: To delete a registered or unregistered server template. You can only delete a registered server template from the account where the server template was originally created. Deletion of a server template does not influence a vServer that was created based on the server template. Deleting a server template that was created based on an assembly causes the deletion of all snapshots associated with that template.

These options are disabled when the vDC is built on Zones virtualization technology.

Creating vNets

vServers are bound to one or more vNets to restrict network connectivity. The different types of vNets that a cloud user can use are:

- **Public networks**: Created by cloud administrators. Cloud users cannot create, update, or delete this type of vNet. Cloud administrators can also share this type of vNet among a number of accounts in a vDC. vServers that are members of public vNets have also external communication beyond vDCs. You can also use the vServer to host public services.
- **Private vNets**: Created by cloud users according to their requirements and within the limits of the account quota. A private vNet is created based on the private network from the network domain of the vDC. Private vNets are only accessible within an account. All vServers that have membership to a private vNet in common can communicate freely through that subnet.

Cloud users define which vNets are associated with a vServer. You can specify the membership of a vServer to one or more vNets when you create a vServer. Once a vNet is associated with a vServer, the association persists until the vServer is deleted.

A cloud user can release a reserved IP address that is not allocated to a vServer. Allocated IP addresses are only released when the vServer is deleted.

Before You Begin

Cloud users can create private vNets. To create a private vNet, you must:

- Plan vServers connectivity.
- Define the number of vServer that can be part of a private vNet.

Creating a Private vNet

Private vNet is a private virtual network set up exclusively for an account. The vServers associated with this vNet have private virtual IP address for internal communication.

To create a private vNet:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Private vNet** in the Actions pane.

   The Create Private vNet Wizard is displayed.
4. Enter the following information in the Private vNet Details step:
   - Name and description for the private vNet.
   - Tags for better identification and classification of the private vNet.

   Click **Next** to configure the private vNet.
5. Select the number of elements for the private vNet.

   This is the maximum number of vServers that can be part of this vNet. Use the slide bar to set the value. The values entered are rounded to the next value of 1, 5, 13, 29, 61, 125, 253, 509, 1021, 2045, 4093, and 8189.

   Click **Next** to view the summary.
6. Confirm the private vNet information and click **Finish** to create the private vNet.

Managing vNets

Virtual network management involves the necessary actions to connect and restrict network access to vServers. After a vNet is created, cloud users can perform the following actions for virtual networks management:

- **Allocate vIP**: To allocate one or more IP addresses from a private or public vNet. IP addresses are dynamically allocated from those available IP addresses that have not been assigned to a vServer or allocated previously. You can use allocated vIP addresses for static assignment to vServers.
- **Deallocate vIP**: To release an IP address that was previously allocated from a public or private vNet. IP addresses assigned to a vServer are not listed and they cannot be deallocated. Once an IP address is deallocated, the IP address is available to the account.
- **Update a Private vNet**: To modify the name or description of a private vNet. This action also allows the creation or deletion of vNet tags.
- **Delete Private vNets**: To delete a private vNet. You cannot delete vNets associated with a vServer.
- **View Reserved IP Addresses**: To see a complete list of the reserved IP addresses for a public or private vNet. Lists all IP addresses of the vNet, the list includes allocated and not allocated IP addresses.
- **View Allocated IP Addresses**: To see a complete list of the allocated IP addresses from a public or private vNet. This option lists all allocated IP addresses, regardless of whether they are assigned to a vServer.
- **View Used IP Addresses**: To see the list of IP addresses that are in use in the public or private vNet. The list includes the IP addresses that are used by the vServers in that account.

Creating Volumes

A volume is a virtual block storage device that you can attach or detach from vServers. Volumes are bound to an account. Storage space for volumes is limited by the account's quota.

You can attach one or more volumes to a vServer at vServer creation time or at a later time.

To create a volume, you can:

- Create a new empty volume specifying only the size.
- Create a volume from a snapshot. A volume created from a snapshot can be empty or not. The size of the volume is defined by the snapshot. Availability of this action depends of the virtualization type of the vDC account. The action might be disabled for some cloud users.
- Import volume from another location. Volumes must be contained in a file of the format `.img`. Volume files must be accessible by the Enterprise Controller, using HTTP, HTTPS or FTP protocols. The action is available only when the vDC is built on Oracle VM Server for x86 based virtualization technology.

Volumes can be shared at volume's creation time. When a volume is shared, the volume is available for all the cloud users of the account.

Before You Begin

Before creating a volume a cloud user must:

- Verify the storage space available for the account and plan accordingly.
- Decide for a suitable option to create the volume.

  When importing a volume, check for file accessibility and format.
- Define whether the volume must be shared.

You can create or import volumes. The procedures to create and import volumes are described in this section.

Creating an Empty Volume

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Volume** in the Actions pane.

   The Create Volume Wizard is displayed.
4. Enter the following information in the Volume Details step:
   - Name and description for the volume.
   - Tags for better identification and classification of the volume.

   Click **Next** to configure the volume.
5. Check the **Shared** option if required.
6. Enter the size of the volume.

   Click **Next** to view the summary.
7. Confirm the volume information and click **Finish** to create the volume.

Creating a Volume from a Snapshot

Creating a volume from a snapshot is only available for vDCs based on Oracle VM Server for x86. To create a volume from a snapshot:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Select the **Storage** tab in the center pane.
4. Select the **Snapshot** sub tab in the center pane.
5. Select a snapshot for the Snapshots list.
6. Click the **Create Volume from Snapshot** action.

   The Create Volume From Snapshot Wizard is displayed.
7. Enter the following information in the Volume Details step:
   - Name and description for the volume.
   - Tags for better identification and classification of the volume.
8. Check the **Shared** option if required.
9. Click **Create** to create the volume.

Importing a Volume

Importing a volume is only available for vDCs based on Oracle VM Server for x86. To import a volume:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Import Volume** in the Action action.

   The Import Volume Wizard is displayed.
4. Enter the following information in the Volume Details step:
   - Name and description for the volume.
   - Tags for better identification and classification of the volume.
5. Enter the URL in which the external volume resides.
6. Check the **Shared** option if required.
7. Click **Import** to import the volume.

Managing Volumes

After a volume is created, perform the following actions for volume management.

- **Update Volume**: To modify the name or description of a volume. The action also allows the creation or deletion of volume tags.
- **Delete Volume**: To delete a volume from a vDC account. Only volumes that are not currently attached to a vServer can be deleted. Deleting a volume results in releasing storage space. The process of deletion of a volume does not influence a snapshot that has been created previously based on that volume.

Creating Snapshots

A snapshot is an image of a volume at a given time. A snapshot captures the current state of the volume and is immutable. You can create snapshots for the following purposes:

- Backup of data stored on a volume
- Creation of new volumes based on a snapshot.

Snapshots availability and management are possible only when supported by the storage of the vDC and for vDCs based on Oracle VM Server for x86. Contact a cloud administrator for verification.

You can create a snapshot from:

- An existing volume.
- Uploading an assembly. See Creating Server Templates.

  In an assembly's template, there are four types of exposed disks. Three of them are presented to cloud users as snapshots, this allows cloud users to optionally create a shared volume from the snapshot. The types of disk created as snapshots are:
  - **Public Populated**: Extra disk with present data
  - **Private Raw**: Dynamically created as empty disk, no disk data, just a size
  - **Shared Raw**: Dynamically created as empty disk, no disk data, just a size

You can create a volume from a snapshot and attach those volumes to vServers.

Before You Begin

Before creating a volume, you must:

- Verify the storage space available for the account and plan accordingly.
- Decide for a suitable option to create the snapshot: from a volume or from an assembly.

To Create a Snapshot from Volume

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Snapshot** in the Actions pane.

   The Create Snapshot Wizard is displayed.
4. Enter the following information in the Snapshot Details step:
   - Name and description for the snapshot.
   - Tags for better identification and classification of the snapshot.

   Click **Next** to select a volume.
5. Select a volume from the list.

   Click **Next** to view the summary.
6. Confirm the snapshot information and click **Finish**.

Managing Snapshots

Perform the following actions for storage management

- **Update Snapshot**: To modify the name or description of a snapshot. The action also allows the creation or deletion of snapshot tags.
- **Create Volume from a Snapshot**: To create a volume base on a snapshot, see Creating Volumes
- **Delete Snapshot**: To delete a snapshot from a vDC account. Deleting a snapshot does not affect volumes, snapshots exist independently of the volume.

Creating Distribution Groups

Distribution groups are necessary for properties similar to anti-affinity scaling. Distribution groups are available for vDCs based on Oracle VM Server for SPARC virtualization technology. The action to create a distribution group might be disabled to some cloud users.

A distribution group is bound to a specific account. You can assign a vServer to a distribution group only when you create the vServer.

You can create a distribution group and define its size. The size of the distribution group represents the minimum amount of anti-affinity desired for running vServers in the distribution group.

Distribution group rules are applied every time a vServer assigned to a distribution group is created or started. Distribution group rules enforce the following behavior for vServers assigned to a specific distribution group:

- When the number of running vServers in the distribution group does not exceed the size of the distribution group, then all running vServers are distributed in a separate server of the vDC. If a spare server is not available in the vDC, then the create or start task fails.
- When the number of running vServers in the distribution group exceeds the size of the distribution group, then running vServers are distributed in the servers of the vDC according to the normal account resources capacities.

At all times, distribution rules satisfy one of the following conditions for vServers assigned to a specific distribution group:

- If the number of running vServers in the distribution group exceeds the size of the distribution group, then the total number of servers in the vDC hosting running vServers is greater or equal to the distribution group size.
- If the number of running vServers in the distribution group does not exceed the size of the distribution group, then the number of servers in the vDC hosting running vServers is greater or equal to the total of vServers of the distribution group in running status.

If the distribution group rules are broken, then an alert is generated and it is displayed on the account and vDC level incidents report. Events that can break the distribution group rules are stopping, deleting, or migrating a vServer.

Before You Begin

Before creating a distribution group, you must:

- Be aware of the size limit for creating a distribution groups. Contact a cloud administrator to verify it.
- Plan the creation and vServers distribution for the distribution groups.

To Create a Distribution Group

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Distribution Group** in the Actions pane.

   The Create Distribution Group Wizard is displayed.
4. Enter the following information in the Distribution Group Details step:
   - Name and description for the distribution group.
   - Tags for better identification and classification of the distribution group.

   Click **Next** to configure the distribution group.
5. Define the size of the distribution group.

   Click **Next** to view the summary.
6. Confirm the distribution group information and click **Finish**.

Managing Distribution Groups

Perform the following actions for distribution groups management:

- **Update a distribution group**: To modify the name or description of a distribution group. The action also allows the creation or deletion of distribution group tags.
- Delete a distribution group.

If the distribution group rule is broken, then an alert is generated and displayed on the account and vDC level incidents report. The vServer that has violated the distribution group rule must be restarted. When a vServer is started, the accounts resources and distribution group rules are checked to place the vServer.

Related Resources for Virtual Datacenters

For instructions in performing actions or to learn more about the role of this feature, go to one of the following resources:

- Chapter 21, "Server Pools"
- Chapter 20, "Oracle VM Server for x86"
- Chapter 18, "Oracle Solaris Zones"
- Chapter 17, "Networks for Virtualization"
- Chapter 16, "Storage Libraries for Virtualization"
- *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide*
- *Oracle Enterprise Manager Ops Center Administration Guide*

See the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm` and the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm` for workflows and end-to-end examples.

Part V

Engineered Systems

Part V contains the following chapters:

23

Oracle Engineered Systems

The following information is included in this section:

Introduction to Oracle Engineered Systems

Traditionally, IT departments have been required to install and configure software that is hosted on hardware that was itself installed and configured separately. Oracle Engineered Systems relieve this burden from the IT department by providing a pre-defined combination of hardware and software. Oracle Engineered Systems is a complete set of integrated hardware and software that is designed to reach a predetermined level of capability, capacity, and scale. This relatively new concept of an integrated hardware and software solution, dedicated to provide a specific service, changes the way one must think about architecture in IT environments.

In addition to removing the burden of installation and configuration, Oracle Engineered Systems provide significant cost savings, this is one of the many advantages delivered by Oracle Engineered Systems. An even larger advantage is the optimization that pre-defined hardware and software enables. As the hardware and software are engineered together to form a complete system, there are multiple opportunities to improve the overall system performance. The unified and integrated monitoring and management of the Oracle Engineered Systems also provide cost savings through simplification of the overall environment. Thus, Oracle Engineered Systems allow datacenter services to be delivered more efficiently through modular or dedicated systems. This greatly simplifies the entire purchase, deploy, configure, monitor, and manage lifecycle of the provided services.

The whole impetus for cloud computing is to provide specific services such as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS) as efficiently as possible; therefore, the increased efficiency delivered by Oracle Engineered Systems make them ideal candidates for the building blocks of a cloud computing environment.

In addition to creating a cloud computing environment, the building blocks address a variety of enterprise architecture requirements to provide a path for enterprise maturation alongside the move to cloud adoption. Each building block is selected as

needed to meet specific maturation needs, and thus, ultimately it provides an increase in overall architectural flexibility. For example, an architect may choose to integrate an Exadata Storage Server and Oracle SuperCluster or Oracle VM Blade Cluster early on to begin the process of consolidation into a scalable architecture, and then expand the system later to meet future demand. Shifting the overall environment into a standardized architecture opens possibilities for future shift into cloud adoption or continuation down the traditional enterprise computing maturation continuum while providing for future scale.

The following are the different Oracle Engineered Systems that are supported by Oracle Enterprise Manager Ops Center.

Oracle SPARC SuperCluster T4-4

Oracle SPARC SuperCluster T4-4 includes a complete stack of hardware and software, computing, storage, and network, all engineered to work optimally together to provide a consolidated platform for running database, middleware, or third party applications. Oracle Enterprise Manager Ops Center is closely integrated with SPARC SuperCluster and provides hardware management, provisioning, and virtualization management.

Oracle SuperCluster T5-8

Oracle SuperCluster T5-8 includes a complete stack of hardware and software, computing, storage, and network, all engineered to work optimally together to provide a consolidated platform for running database, middleware, or third party applications. Oracle Enterprise Manager Ops Center is closely integrated with Oracle SuperCluster T5-8 and provides hardware management, provisioning, and virtualization management.

Oracle SuperCluster M6-32

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

Understanding User Roles

In Oracle Enterprise Manager Ops Center, users are assigned several roles such as Asset Admin, Cloud Admin, SuperCluster Systems Admin, and many more. Each role grants the user a set of permissions; a particular permission can be granted by more than one role such as Asset Management, Network Management, and other management roles.

You can add users to Oracle Enterprise Manager Ops Center from the local authentication subsystem of the Enterprise Controller's operating system. Each user is given a different role which grants or denies access to the different functions of Oracle Enterprise Manager Ops Center.

The following user roles are described in this section:

Ops Center Administrator Role

The Ops Center Administrator user role is only used for initial discovery of the Oracle SuperCluster system. As an OpsCenter Admin user, though you have permissions, do not perform actions that are disabled for the SuperCluster Systems Admin Role. All discovery operations on Oracle SuperCluster systems must be started using OpsCenter Administrator account only.

When using Oracle Engineered Systems in Oracle Enterprise Manager Ops Center, some actions might be potentially dangerous; hence those actions are disabled.

See **Asset Protection** for more information.

SuperCluster Systems Admin Role

The SuperCluster Systems Administrator is responsible for overall monitoring and management of all associated Oracle SuperCluster systems. The SuperCluster Systems Administrator also has privileges to manage Virtual Pools, Storage, and Oracle Solaris Cluster. As a SuperCluster Systems Administrator, you can perform most of the operations that are allowed in the context of the Oracle SuperCluster system.

The SuperCluster Systems Administrator role is a default role recommended for the management of the Oracle SuperCluster system.

Prerequisites: The user must be familiar with the use of Oracle Enterprise Manager Ops Center and be familiar with hardware management and OS management in general.

Creating a SuperCluster Systems Administrator Role

You can create a SuperCluster Systems Administrator user to manage an Oracle SuperCluster system. You can create one or more users that have access to the same Oracle SuperCluster system.

Note: Before you create a SuperCluster Systems Administrator role, create a user on your OS or connect to the LDAP sever.

To create a SuperCluster Systems Administrator role, perform the following steps:

1. In the Navigation pane, click **Administration**.
2. Under Enterprise Controller, click **Local Users**.

Figure 23–1 Local Users

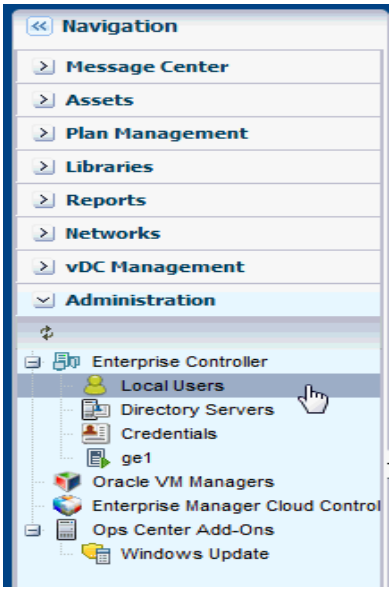

3. In the Actions pane, click **Add User**.

Figure 23–2 Add User

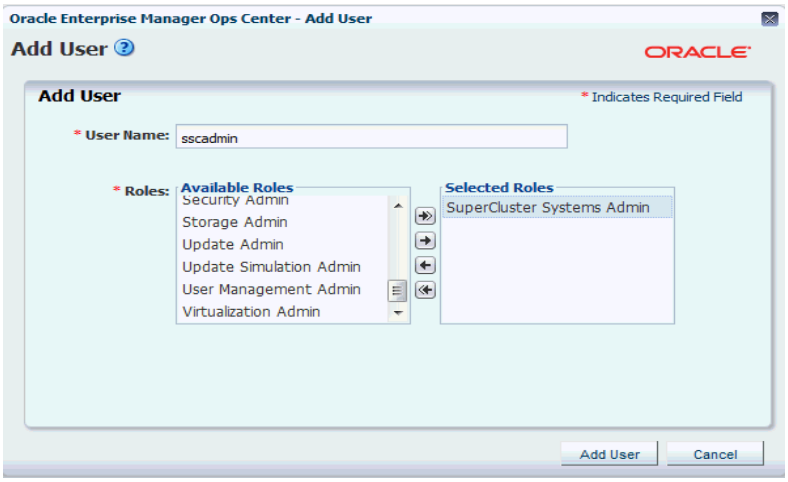

4. In the User Name field, enter a name for the user role (for example, sscadmin, provided such local user exists on the OS or directory server if configured). The sscadmin user requires a UNIX user on the Enterprise Controller system where Oracle Supercluster system will be managed and monitored.
5. Select **SuperCluster Systems Admin** in the Available Roles section and click the right arrow to move the role to the Selected Roles section.
6. Click **Add User**. The new user role is added.

Assign Permissions to the Role

After you add the role, you must assign permissions to the new role that you created.

1. In the Navigation pane, click **Administration**.
2. Under Enterprise Controller, click **Local Users**.

3. In the center pane, click the **Users** tab.
4. Select the sscadmin user, then click the **Manage User Roles** icon.

Figure 23–3 Manage User Roles

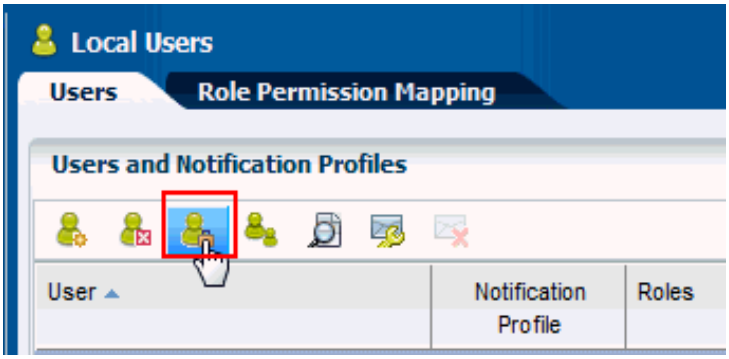

5. In the User Roles screen, uncheck **Use the default Role associations**, then click **Next**.

Figure 23–4 Manage Roles

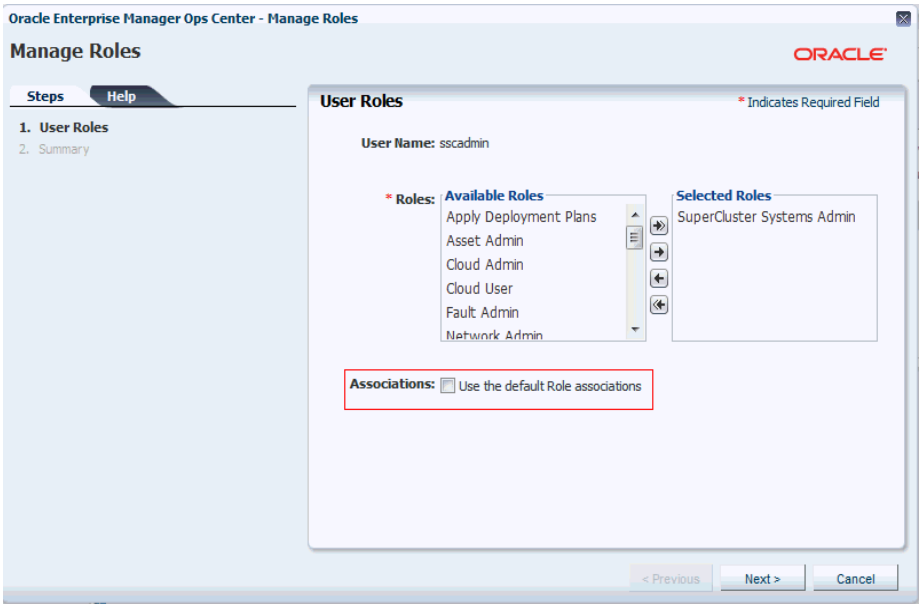

6. In the Specify Asset Privileges screen, expand **All Assets,** then expand **Oracle Engineered Systems**.

Figure 23–5 Specify Asset Privileges

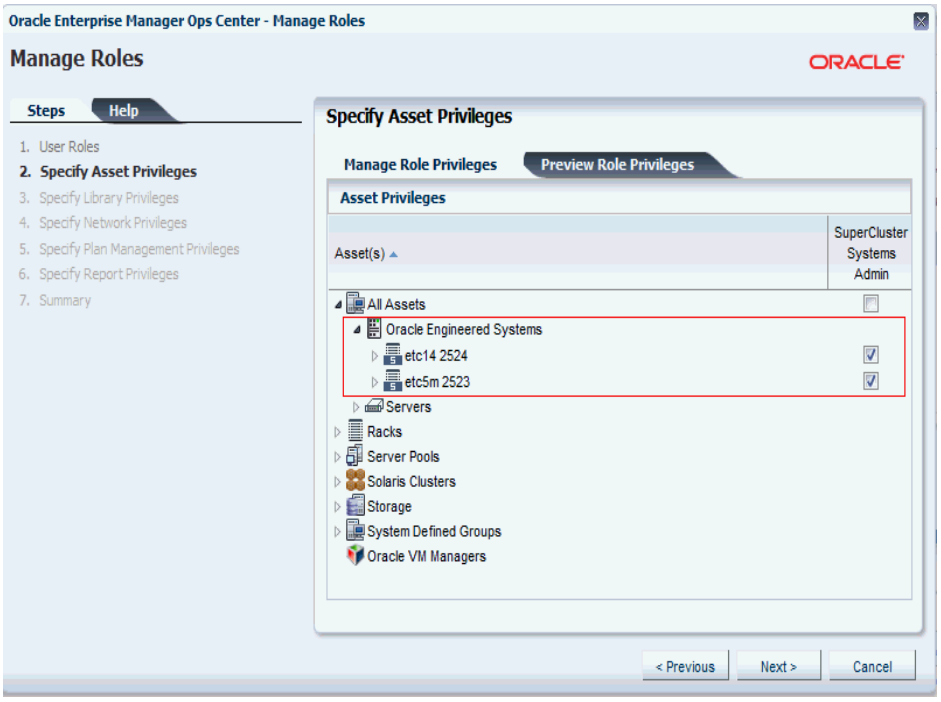

7. Select the Oracle SuperCluster assets that you want the SuperCluster Systems Admin to manage.
8. Click **Next** on the following screens.
9. Review the Summary, then click **Finish**.
   Only the Oracle SuperCluster assets that have been selected are displayed in the Navigation pane under All Assets.

> **Note:** The SuperCluster Systems Admin role prevents you from performing unsupported actions by mistake.

SuperCluster Systems Admin Role Permissions

As a SuperCluster Systems Admin, you can perform most operations available in Oracle Enterprise Manager Ops Center. However, some operations are disabled because they could compromise pre-installed Oracle Engineered Systems' hardware or software infrastructure.

The following actions are disabled for SuperCluster Systems Admin user:

- PDOM management (creating LDOMs on M-series servers)
- LDOM management (Connect Network and Storage Management with the exception of Starting / Shutting down / Rebooting LDOM)
- Install / Update OS on LDOMs and Global zones
- Install Server
- Rack management
- Chassis management
- Network Switch management

- Network infrastructure operations (Attach network to global zone, IPMP operations)
- Firmware updates
- EC Management (downloads, upgrade, EC Proxy / Agent management, storage library management)
- OVM management
- Asset discovery
- Create logical domain
- Switch management (edit action is permitted)
- Remove SuperCluster

Note: Oracle Solaris cloud management operations and virtual datacenter (vDC) functionality cannot be performed by Supercluster Systems Admin user. To perform these actions, create Cloud Admin and Cloud User roles.

Cloud Admin

The Cloud Administrator's responsibilities include setting up of infrastructure and resource allocation so that cloud users can deploy their application onto authorized accounts. They also manage the cloud users accessing the accounts and their authorization.

Prerequisites: The user must be trained on Oracle Enterprise Manager Ops Center, installation and configuration, and the continual maintenance of the product.

Cloud User

Cloud users create virtual servers and deploy applications. Cloud users are restricted to virtual datacenter infrastructure activities and are presented only with the required options on the Oracle Enterprise Manager Ops Center UI.

Prerequisites: The user must be familiar with the use of Oracle Enterprise Manager Ops Center, and also its hardware management and OS management in general.

Oracle Engineered Systems Management

This section describes in detail about Oracle Engineered Systems Management using which you can enable a single or centralized Oracle Enterprise Manager Ops Center installation managing multiple Oracle SuperCluster systems and other assets in the datacenter directly.

Using Oracle Enterprise Manager Ops Center, you can discover Engineered Systems from a single datacenter instance and perform complete management and monitoring of multiple engineered systems.

Figure 23–6 is a pictorial representation of datacenter management of engineered systems using Oracle Enterprise Manager Ops Center.

Figure 23–6 Oracle Engineered Systems Management

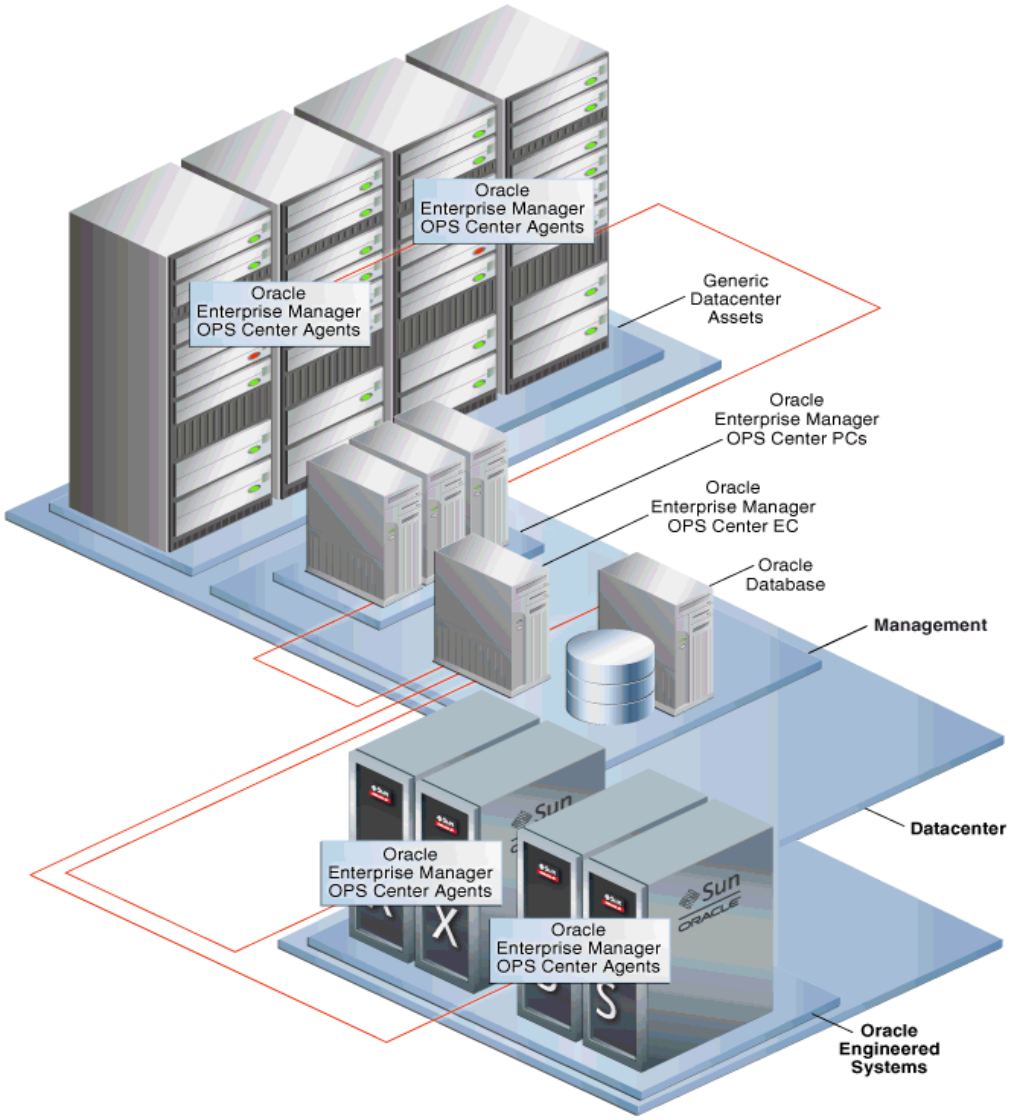

Configuring Oracle Engineered Systems

This section describes the following scenarios:

- Overlapping IB Networks Enabled
- Overlapping IB Networks not Enabled

Overlapping IB Networks Enabled

Overlapping IB networks are not supported by default, but starting with Oracle Enterprise Manager Ops Center 12c Release 2 (12.2.2.0.0), the same can be enabled. Overlapping network systems must be discovered only after the feature is enabled.

To enable overlapping IB networks, perform the following steps:

1. In the Navigation pane, under Administration, click **Enterprise Controller**.
2. In the center pane, click **Configuration**.

3. In the Configuration Management section, select Network/Fabric Manager from the Subsystem drop-down list.
4. Set the value of **oem.oc.networkmgmt.ib.overlapping.enabled** property file to true.

> **Note:** Restart the Enterprise Controller for the changes to take effect.

Overlapping IB Networks not Enabled

One or more Oracle Engineered Systems can be discovered and managed by a single Oracle Enterprise Manager Ops Center instance based on the following conditions:

- None of Oracle Engineered System instances have overlapping private networks connected through IPoIB, that is, networks that have the same CIDR (Classless Inter-Domain Routing) or networks that are sub-blocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping.
- None of the Oracle Engineered System instances or generic datacenter assets have overlapping management or client access networks connected through Ethernet, that is, networks that have the same CIDR or networks that are sub-blocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping. As an exception, you can use the same CIDR (not sub-block) for multiple systems. For example, you can use 192.0.2.1/22 as a CIDR for Ethernet network on one or more engineered systems and/or generic datacenter assets.
- None of the Oracle Engineered System instances have overlapping public networks connected through EoIB, that is, networks that have the same CIDR or networks that are sub-blocks of the same CIDR. For example, 192.0.2.1/21 and 192.0.2.1/24 are overlapping. As an exception, you can use the same CIDR (not sub-block) for multiple systems. For example, you can use 192.2.0.0/22 as a CIDR for public EoIB network on multiple engineered systems.
- None of the networks configured in Oracle Enterprise Manager Ops Center overlaps with any network, that is, overlapping networks are not supported by Oracle Enterprise Manager Ops Center.

> **Note:** To manage two or more engineered systems that have overlapping networks or any networks already present in Oracle Enterprise Manager Ops Center, reconfigure one of the conflicting systems before it is discovered and managed by the same Oracle Enterprise Manager Ops Center.

Limitations

Do not create server pools using private networks attached to members from two or more SuperCluster systems (racks). To create server pools with members from two or more SuperCluster systems, use public networks. Use private networks only in server pools with members belonging to the same SuperCluster system.

Example Oracle SuperCluster Network Configurations

The following are example Oracle SuperCluster network configurations that you can use when configuring the network to discover and manage Oracle SuperCluster systems. Status OK indicates a valid configuration and status Fail indicates an invalid configuration.

Table 23–1 Example SuperCluster Network Configuration-1

	1 GbE	10 GbE	IB
SuperCluster1	192.0.251.0/21	192.4.251.0/24	192.168.30.0/24
SuperCluster2	192.0.251.0/21	192.4.251.0/24	192.168.31.0/24
Status	OK	OK	OK

Status:

OK - SuperCluster1-1GbE and SuperCluster2-1GbE share the same network.

OK - SuperCluster1-10GbE and SuperCluster2-10GbE share the same network.

OK - SuperCluster1-IB does not overlap with SuperCluster2-IB.

Table 23–2 Example SuperCluster Network Configuration-2

	1 GbE	10 GbE	IB
SuperCluster1	192.0.251.0/21	192.0.250.0/24	192.168.30.0/24 - IB fabric connected with SuperCluster2
SuperCluster2	192.6.0.0/21	192.0.250.0/24	192.168.30.0/24 - IB fabric connected with SuperCluster1
Status	OK	OK	OK

Status:

OK - SuperCluster1-1GbE and SuperCluster2-1GbE represent different non-overlapping networks.

OK - SuperCluster1-10GbE and SuperCluster2-10GbE share the same network.

OK - SuperCluster1-IB and SuperCluster2-IB represent the same network as they are interconnected.

Table 23–3 Example SuperCluster Network Configuration-3

	1 GbE	10 GbE	IB
SuperCluster1	192.0.2.1/21	192.0.251.0/21	192.168.30.0/24
SuperCluster2	192.0.0.128/25	192.0.7.0/24	192.168.30.0/24
Status	FAIL	OK	FAIL

Status:

FAIL - SuperCluster1-1GbE and SuperCluster2-1GbE define overlapping networks.

OK - SuperCluster1-10GbE and SuperCluster2-10GbE represent different non-overlapping networks.

FAIL - SuperCluster1-1GbE and SuperCluster2-10GbE define overlapping networks.

FAIL - SuperCluster1-IB and SuperCluster2-IB do not define unique private networks (racks are not interconnected).

Oracle Solaris 11 Software Library Setup

When you discover an Oracle SuperCluster system, default install agents work only if the Oracle Solaris 11 Software Update Library is correctly setup, because Oracle Solaris uses Oracle Solaris 11 Software Update Library.

If Oracle Enterprise Manager Ops Center was just installed, initialize the Oracle Solaris 11 Software Update Library before the discovery is started as it will fail to install agents. Ensure that the Oracle Solaris 11 Software Update Library contains the correct Oracle Solaris packages that your Enterprise Controller and Proxy Controllers use and also the SRUs that are used on the Oracle Engineered System.

To manage generic assets, you also need the correct Oracle Solaris packages for each generic Solaris 11 managed OS. Typically, you must create a full copy of the Oracle Solaris 11 support repository.

Starting with Oracle Enterprise Manager Ops Center 12c Release 2 (12.2.2.0.0), Oracle SuperCluster Solaris 11 OS uses the default HMP packages delivered with Oracle SuperCluster / QFSDP, instead of packages normally delivered by Ops Center installation.

It is recommended to install agents in all domains.

Deploy Proxy Controller

Deploy the Proxy Controller only if you do not have a suitable Proxy Controller in Oracle Enterprise Manager Ops Center that can discover Oracle Engineered Systems.

Perform the following steps to deploy the Proxy Controller on Oracle Enterprise Manager Ops Center.

1. In the Navigation pane, click **Administration**.
2. In the Actions pane, click **Deploy Proxy**.
3. Select **Remote Proxies**, then click **Next**.
4. Enter the Proxy Hostname/IP, SSH User, SSH Password, Privileged Role, and Privileged Password in the respective fields.
5. Click **Next**. The Remote Proxy Controller is deployed. This might take a few minutes.
6. Review the Summary, then click **Finish**.

The Remote Proxy Controller is deployed on Oracle Enterprise Manager Ops Center. You can now perform discovery of the Oracle Engineered Systems.

Prepare Setup for Oracle Engineered Systems Discovery

The setup is based on the following options:

- Network is identified by the Enterprise Controller
- Network is not identified by the Enterprise Controller

> **Note:** Ensure you have a Proxy Controller deployed that can access the network. To deploy a Proxy Controller, see Deploy Proxy Controller.

Network is Identified by the Enterprise Controller

If the Enterprise Controller host identifies the management network of the Oracle SuperCluster (CIDR must be the same), ensure that the network is assigned to the Proxy Controller. If it is not assigned, assign the network to the Proxy Controller.

To assign the network to the Proxy Controller, perform the following steps:

1. In the Navigation pane, select **Administration**.

2. Select a Proxy Controller.
3. In the Actions pane, click **Associate Networks**.

Network is not Identified by the Enterprise Controller

If the Enterprise Controller host does not identify the network (Oracle Engineered System management network must be routable from it), create a fabric definition and a network for the fabric.

Create Fabric Definition

1. In the Navigation pane, under Networks, select **Fabrics** from the drop-down list.
2. In the Actions pane, click **Define Ethernet Fabric**.
3. In the Fabric Name field enter a name for the fabric.
4. (Optional) Enter the description.
5. Click **Next**.
6. Enter the VLAN ID Ranges, then click **Next**.
7. Select the networks to be associated with the fabric, then click **Next**.
8. Review the Summary, the click **Finish**.

   A new fabric is created.

Create Network for the Fabric

After the fabric is created, you must create a network for the new fabric.

1. In the Navigation pane, under Networks, select **Networks** from the drop-down list.
2. In the Actions pane, click **Define Network**.
3. In the Network IP field, enter the IP address (in CIDR format) of the network that represents the management network of the Oracle Engineered System you want to manage.
4. Enter the **Gateway IP address**.
5. In the Network Name field, enter a name for the network.
6. Click **Next**.
7. Assign the newly created fabric to the Proxy Controller, then click **Next**.

   The setup is now ready for Oracle Engineered System discovery.

Ports for Oracle SuperCluster

The proxy Controller for an Oracle SuperCluster engineered system does not have unique ports or protocols. The following table summarizes the set of ports and their protocols used by Oracle SuperCluster.

Table 23–4 Required Ports and Protocols for Oracle SuperCluster

Communication Direction	Protocol and Port	Purpose
Proxy Controller to Exadata's ILOM Service Processors	SSH, TCP: Port 22 IPMI, TCP, UDP: Port 623	Proxy Controller discovers, manages, and monitors the service processor of Exadata.

Table 23–4 (Cont.) Required Ports and Protocols for Oracle SuperCluster

Communication Direction	Protocol and Port	Purpose
Proxy Controller to Exadata cells	SSH, TCP: Port 22	Proxy Controller discovers, manages, and monitors the compute nodes.
Proxy Controller to Oracle ZFS Storage Appliance	SSH, TCP: Port 22 IPMI, TCP, UDP: Port 623	Proxy Controller discovers, manages, and monitors the service processor of the storage appliance.
Proxy Controller to Oracle ZFS Storage Appliance	SSH: Port 215	Proxy Controller discovers the projects of the storage appliance: ■ iSCSI volumes. ■ NFS shares
Proxy Controller to Cisco switch	SSH version 2: Port 22 SNMP: Port 161	Proxy Controller discovers and manages the switch
Proxy Controller to InfiniBand switch	SSH: Port 22 IPMI: Port 623	Proxy Controller discovers and manages the switch.

Discovering Oracle Engineered Systems

Oracle SuperCluster can be discovered only by trained Oracle staff. Oracle SuperCluster discovery is supported as a free service during Oracle SuperCluster installation only. You must request Oracle SuperCluster discovery before the Oracle SuperCluster installation.

Note: After Oracle SuperCluster installation, Oracle SuperCluster discovery is offered only as a paid service.

Asset Protection

Some actions in Oracle Enterprise Manager Ops Center might break an Oracle SuperCluster engineered system configuration. To protect the engineered system from inadvertent changes in the configuration and the resulting loss of service, high-risk actions are disabled for all users. For example, modifying logical domains or network configuration might break the engineered system configuration leading to loss of service.

The following actions are disabled for all users, including Ops Center Admin:

Control Domains

- Attach Network
- Create Logical Domains
- Cancel Delayed Reconfiguration

Logical Domains

- Delete logical domain
- Migrate logical domain
- Connect network

- Add storage
- Move metadata
- Edit attributes of logical domain (with exception of name and description)
- Enable/Disable automatic recovery
- Network/Network Connectivity operations
- Network/Link Aggregation operations
- Network/IPMP group operations
- Network/Bandwidth Flow operations

Global Zones

- Network/Network Connectivity operations
- Network/IPMP group operations
- Network/Bandwidth Flow operations

Common Operations

- Install Software
- Deploy/Update Software
- Update Firmware
- Update BIOS

Disable Asset Protection

The Ops center Administrator role has special privilege to turn off the Asset protection feature in engineered systems.

> **Note:** Use this privilege for emergency purpose only.

Enable the *oes.asset.protection.override.user* parameter in the **/opt/sun/n1gc/lib/XVM_SATELLITE.properties** file to allow all actions, including actions disabled by asset protection.

The Oracle Enterprise Manager OpsCenter Enterprise Controller must be restarted for changes in the XVM_SATELLITE.properties file to take effect.

Adding Multiple Racks

You can add multiple racks to a single Engineered System. To add multiple racks, you need different configuration files in respect to assets that you want to mount on the additional racks. After the configuration files are uploaded and validated, they are merged into a single profile, thus representing a single Engineered System.

Perform the following steps to add multiple racks to the Engineered System asset:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**, then click the discovery profile that you used to discover the Engineered System.

Figure 23–7 Select Engineered System Discovery Profile

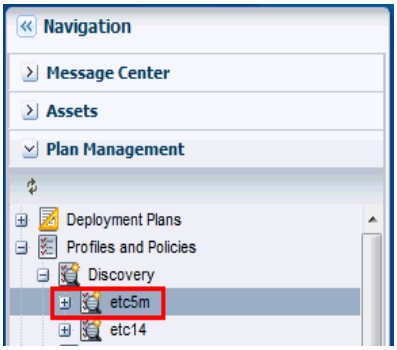

3. In the Actions pane, click **Edit Profile**.
4. Click **Next**.
5. Click **Next**. The Configuration Files wizard is displayed.
6. Click **Browse** to select the configuration file for a new rack that you want to add to the Engineered System.
7. Click **Upload** to upload the configuration file.
   You can also embed discovery information of multiple racks into a single configuration file.
8. Click **Next**, then click **Finish** to create a discovery profile.
9. Run the discovery profile with the new configuration file. See Discovering Oracle Engineered Systems.

   The additional rack is discovered and displayed under the respective Engineered System.

Figure 23–8 Multiple Racks Under a Single Engineered System

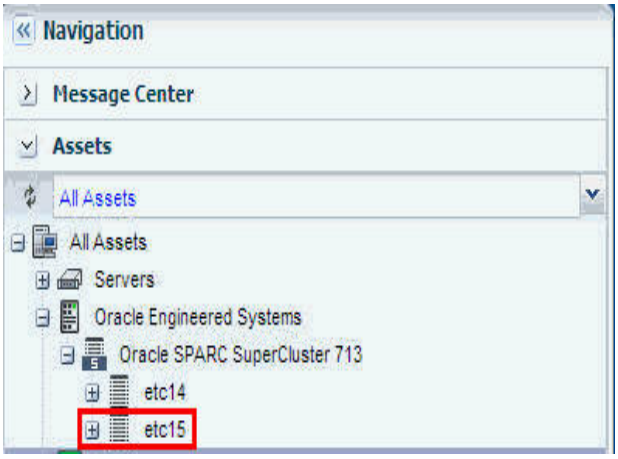

Adding Oracle SuperCluster M6-32 Server

You can add additional Oracle SuperCluster M6-32 servers to the Oracle Engineered System using Oracle Enterprise Manager Ops Center.

Perform the following steps to add an Oracle SuperCluster M6-32 server to the Engineered System asset:

1. Create a discovery profile.
2. Upload the configuration file to discover the Oracle SuperCluster M6-32 server.
3. Click **Finish** to create the discovery profile.

Add the Oracle SuperCluster M6-32 Server

After the discovery profile is successfully created, run the discovery job and add the Oracle SuperCluster M6-32 server to the Oracle Engineered System asset.

Perform the following steps to add the Oracle SuperCluster M6-32 server:

1. In the Navigation pane, click **All Assets**.
2. In the Actions pane, click **Add Assets**.
3. Select **Add and manage various types of assets via discovery probes**, then click **Next**.
4. Select the discovery profile that you created to add the Oracle SuperCluster M6-32 server, then click **Add Now**.
   The Oracle SuperCluster M6-32 server is discovered and displayed under Assets in the Navigation pane.

Figure 23–9 Oracle SuperCluster M6-32 Server

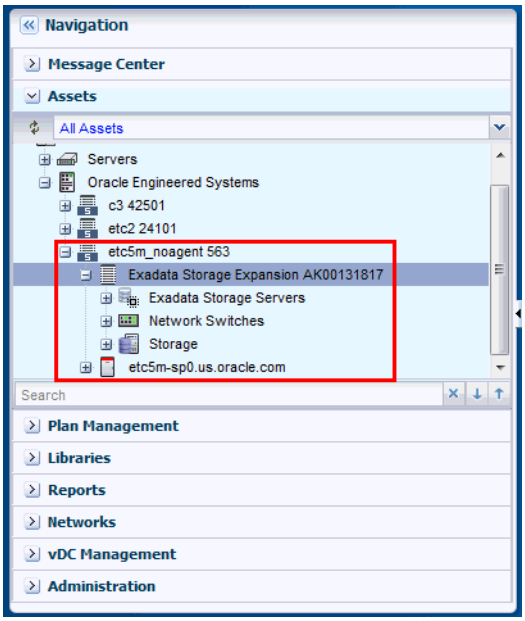

Adding Additional Hardware to the Rack

You can add additional hardware to the existing Engineered System rack. To add additional hardware, ensure you update the configuration file that was used originally for the discovery of the Oracle Engineered System asset with metadata information of the new hardware that you want to add.

Perform the following steps to add additional hardware:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
   The discovery profiles are listed.

3. Select the discovery profile for which you want to upload the updated configuration file.
4. Click the **Edit Profile** icon.

Figure 23–10 Select Discovery Profile

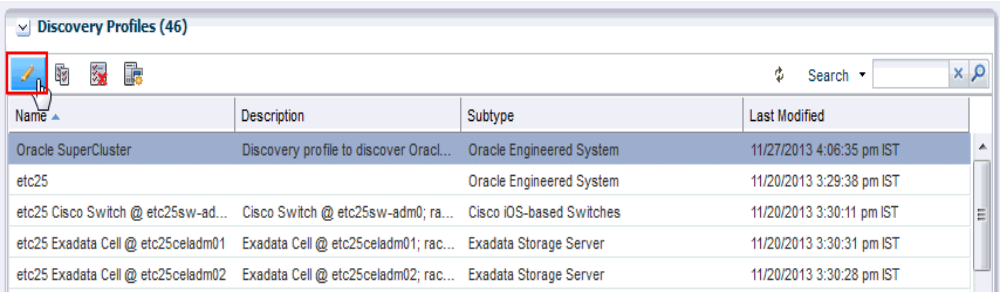

5. Click **Next** in the Identity Profile screen.
6. Click **Next** in the Tags screen.
7. In the Configuration Files(s) screen, click the **Browse** button to select the updated configuration file, then click **Upload**.
8. After the configuration file is uploaded successfully, click **Next**.
9. Review the Summary, then click **Finish** to create the discovery profile.

Add the Additional Hardware

After the discovery profile is successfully created, run the discovery job and add the additional hardware to the same rack that is mounted on the Oracle Engineered System asset.

Perform the following steps to add the additional hardware to the Oracle Engineered System rack.

1. In the Navigation pane, click **All Assets**.
2. In the Actions pane, click **Add Assets**.
3. Select **Add and manage various types of assets via discovery probes**, then click **Next**.
4. Select the discovery profile that you updated to add the additional hardware, then click **Add Now**.
   The additional hardware is discovered and displayed under All Assets in the Oracle Engineered Systems pool in the Navigation pane.

Viewing the Oracle Engineered Systems

You can view all types of engineered systems in the Navigation pane under All Assets. The engineered system assets are grouped by racks on which they are mounted. Assets mounted on a rack are grouped by gear type.

1. In the Navigation pane, under Assets, select **All Assets**.
2. Expand **Oracle Engineered Systems**.

   The Engineered System assets are displayed.
   In Figure 23–11, Engineered System assets (Oracle SPARC SuperCluster T4-4 and Oracle SuperCluster M6-32) are grouped by racks on which they are mounted.

Figure 23–11 Engineered Systems View

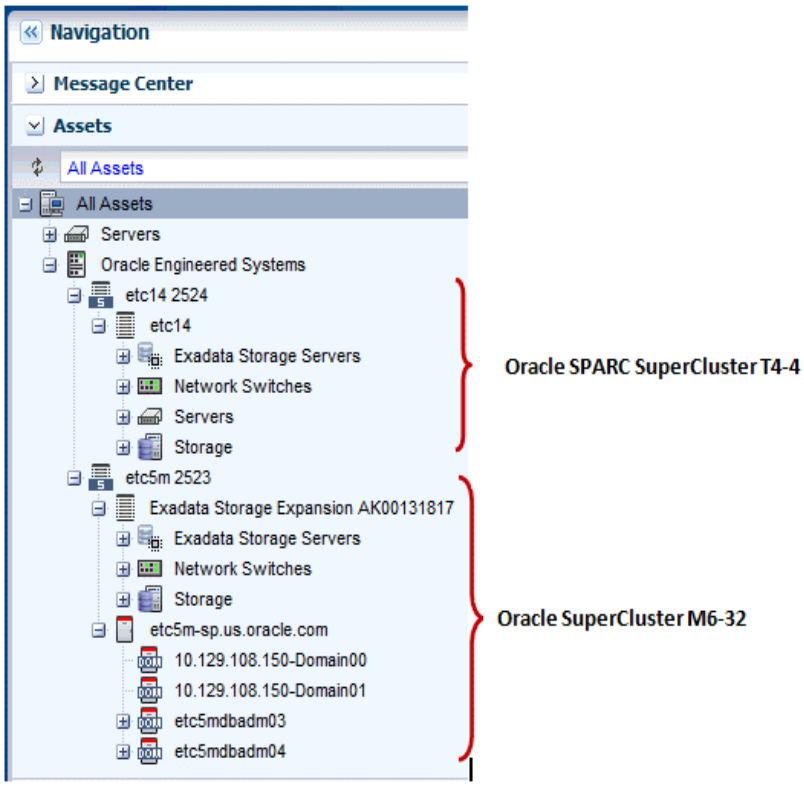

Photorealistic View

The photorealistic view represents the physical rack where you can view the placement of the rack and its components. The front and rear views of the rack are displayed in this view. All slots and the respective assets are displayed.

Each asset in the rack is represented by an image. The health status of assets such as OK, Warning, or Critical is displayed in the form of colored lights as seen on the physical rack itself. OK status is identified by the green color. Warning and critical status are identified by the yellow color. Hover the mouse over the slots in the rack and view details about the assets such as slot number, asset name and description, type of asset, model number of the asset, and its health status.

A question mark in any of the slots indicates a discrepancy from the baseline check. See Baseline Check for more information. For example, see Figure 23–12.

Figure 23–12 Baseline Check

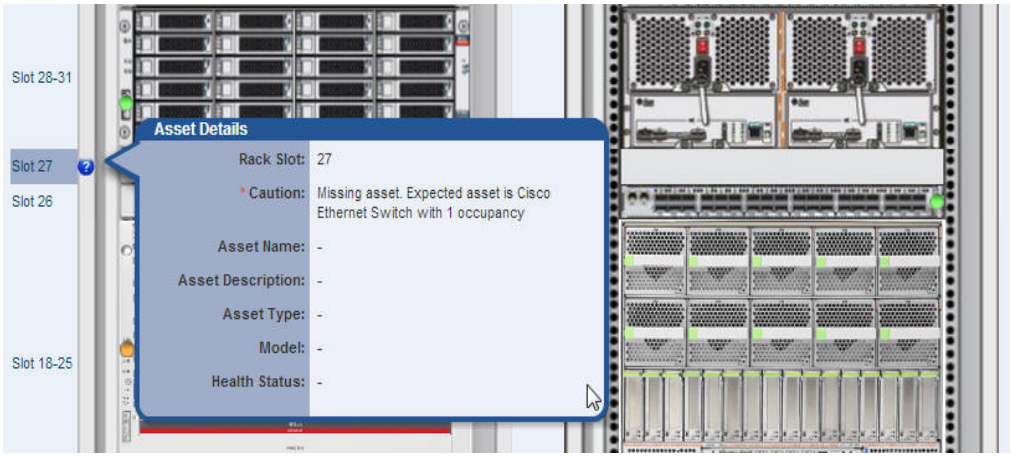

In this document, photorealistic view of the Oracle SuperCluster T5-8 rack is described.

To see the photorealistic view of the rack, perform the following steps:

1. In the Navigation pane, under Assets, select an asset from Oracle Engineered Systems.
2. Select a rack that you want to view.
3. In the center pane, click the **Details** tab.

   A photorealistic view of the Oracle Engineered System rack is displayed.

Figure 23–13 Photorealistic view of Oracle SuperCluster System

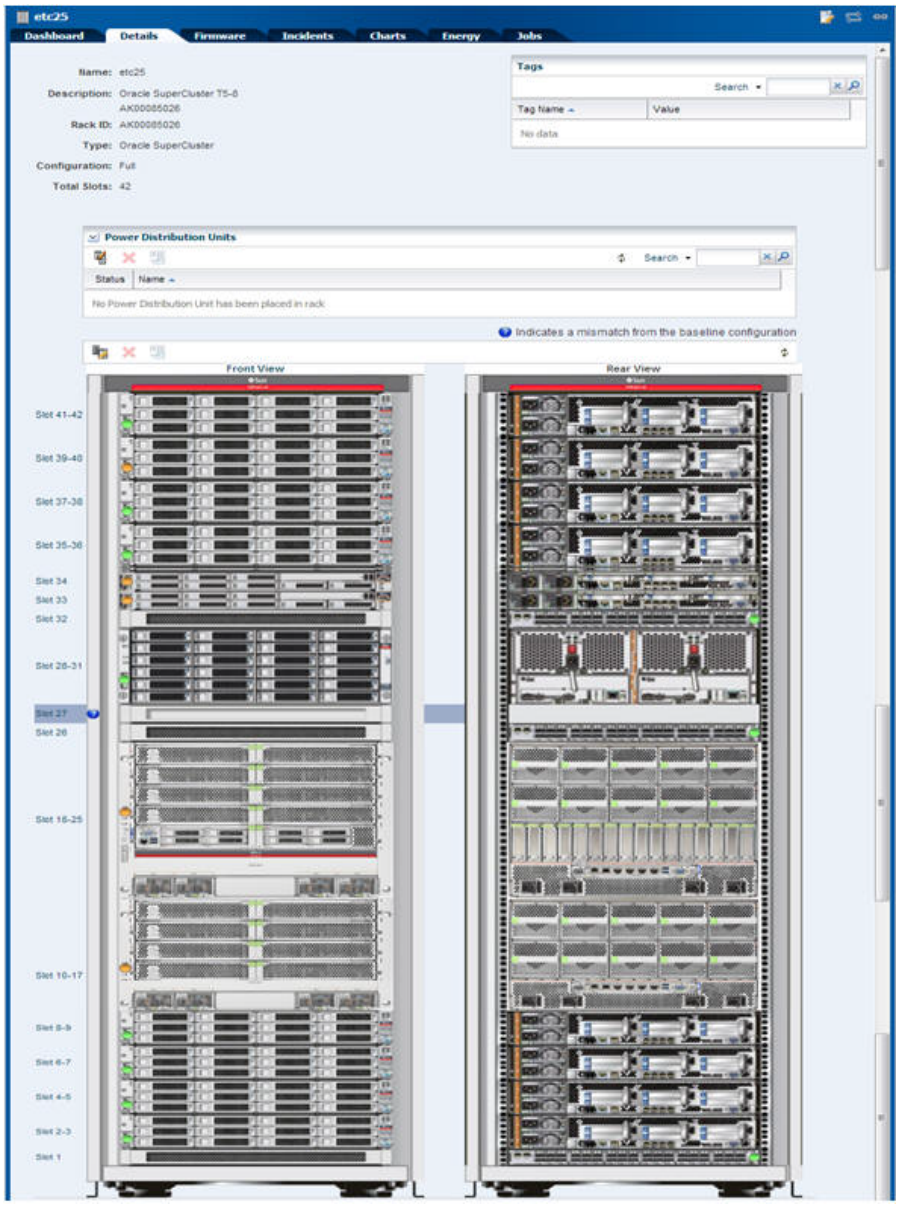

Creating Oracle Engineered Systems Report

Reports provide information about assets, such as job history, firmware, operating system updates, and incidents. Reports are created in Interactive, PDF, and CSV formats. You can generate and view the report for multiple Oracle Engineered Systems.

Perform the following steps to create a report for Oracle Engineered System.

1. In the Navigation pane, click **Reports**.

Figure 23–14 Create Reports

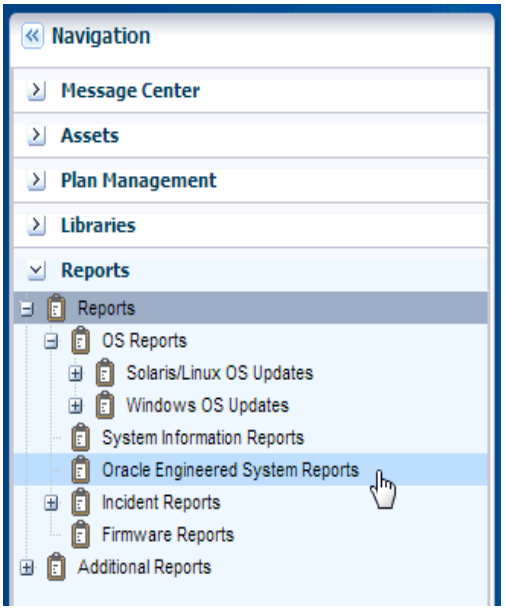

2. Click **Oracle Engineered System Reports**.
3. In the Actions pane, click **Create Oracle Engineered System Report**.
4. In the Define Report Parameters wizard, enter a name and description for the report.
   The Schedule and Output Format are checked by default.

Figure 23–15 Define Report Parameters

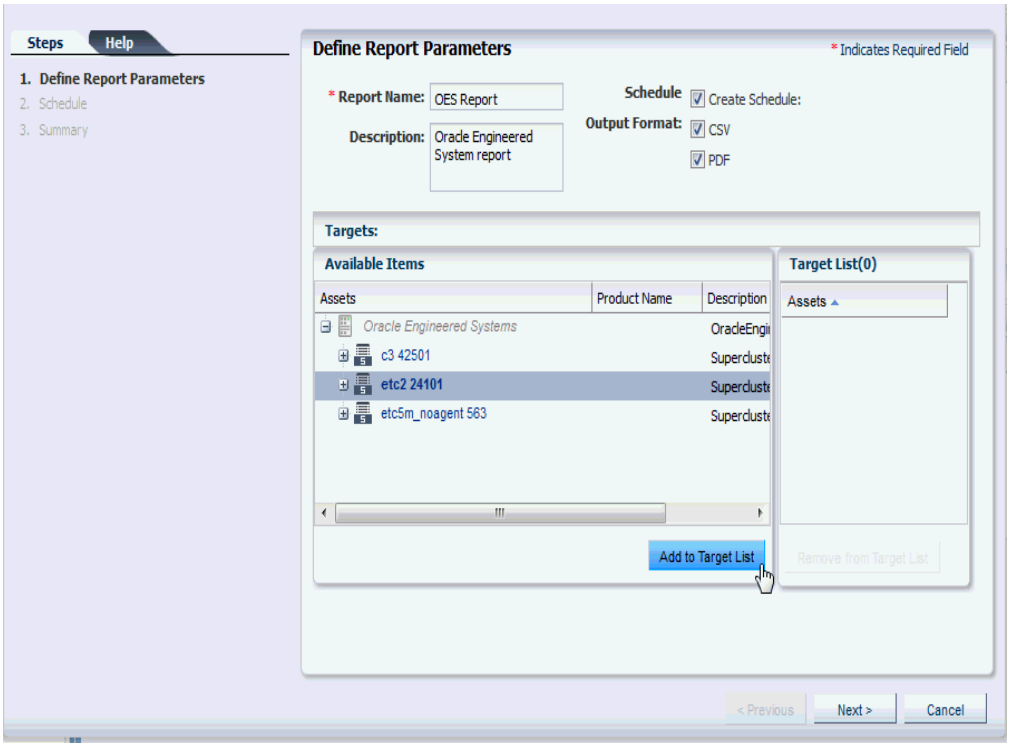

- Select Create Schedule to run the report later or on a recurring schedule.

- Select the output formats of the result that will be generated for the report.

5. In the Targets section, select the asset for which you want to run the report and click **Add to Target List**.
6. Click **Next**. The Schedule wizard is displayed.
7. Select a schedule for the report. You can schedule the report to run on the following instances:
   - Now: Runs the report immediately.
   - At a later date/time: Select a date and time to generate the report.
   - On a Recurring Schedule: Select the month and day when you want to generate the report. Select the Start Time, End Time, and Number of Hours between runs. This is to set the number of times the report is generated between the specified start and end time. For example, if you set the start time at 6.00 a.m, end time at 12.00 a.m, and the number of hours between runs as 2, then the report is run at 6.00 a.m, 8.00 a.m, 10.00 a.m, and 12.00 a.m.
8. Click **Next**. The Summary wizard is displayed.
9. Verify the report parameters and click one of the options as required:
   - **Save Template and Close**: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
   - **Run and Close:** Runs the report and closes the wizard window.

Figure 23–16 Report Created

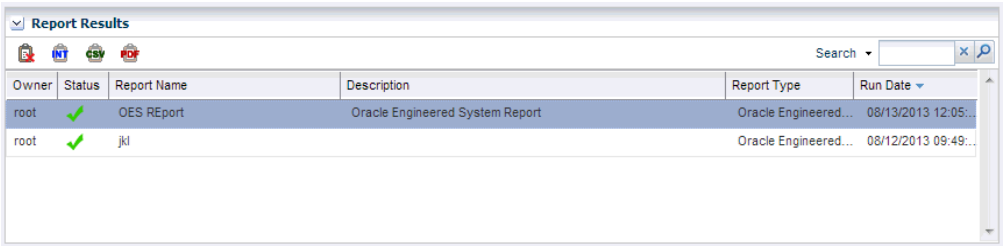

View the Engineered System Report

Using reports, you can view the rack setup for each of the rack within the system, including the asset details related to the rack.

To view an Engineered System Report, perform the following steps.

1. In the Navigation pane, under Reports, select **Oracle Engineered System Reports**.
2. Select a report and click the format in which you want to view the report.

Figure 23–17 View Report

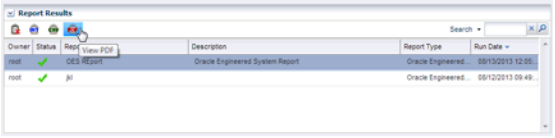

The report is displayed in the selected format.

Oracle SuperCluster

Oracle SuperCluster is an Oracle Engineered System that integrates SPARC compute nodes, a Sun ZFS Storage Appliance, InfiniBand switches, PDUs, and Exadata Storage Servers into a multi-rack system.

Oracle SuperCluster is supported in the following configurations:

- **Oracle SPARC SuperCluster T4-4**: The Oracle SPARC SuperCluster T4-4 hardware includes two SPARC T4-4 servers, three Sun Datacenter InfiniBand Switch 36 switches as InfiniBand backplane, a Cisco 4948 48-port 1Gb Ethernet Switch for external connectivity (SSH version 2 only supported), and a Sun Rack II. Oracle SPARC SuperCluster T4-4 is available in two configurations, half rack and full rack. The half rack includes two SPARC T4-4 servers (with four processor modules per server) and three Exadata Storage Servers. The full rack includes four SPARC T4-4 servers (with four processor modules per server) and six Exadata Storage Servers. In addition, for additional storage capacity, connecting to an Oracle Exadata Storage Expansion Rack is supported.
- **Oracle SuperCluster T5-8**: The Oracle SuperCluster T5-8 hardware includes two SPARC T5-8 servers, three Sun Datacenter InfiniBand Switch 36 switches as InfiniBand backplane, a Cisco 4948 48-port 1Gb Ethernet Switch for external connectivity (SSH version 2 only supported), and a Sun Rack II. Oracle SuperCluster T5-8 is available in two configurations, half rack and full rack. The half rack includes two SPARC T5-8 servers (with four processor modules per server) and four Exadata Storage Servers. The full rack includes two SPARC T5-8 servers (with eight processor modules per server) and eight Exadata Storage Servers. In addition, for additional storage capacity, connecting to an Oracle Exadata Storage Expansion Rack is supported.
- **Oracle SuperCluster M6-32**: Oracle SuperCluster M6-32 delivers 16, 24, or 32 M6 processors with up to 32 TB of memory in a dedicated SPARC M6–32 rack. Compute capacity can be assigned flexibly, depending on customer requirements. Layered Optimized Virtualization allows resources to be configured hierarchically in physical domains (PDoms), logical domains (LDoms), and Oracle Solaris Zones. The SPARC M6-32 server delivers mainframe-class availability; two chassis may also be configured for extreme redundancy. An external storage rack provides nine Exadata Storage Servers, a ZFS Storage Appliance, three Sun Data center InfiniBand Switch 36 switches, and a Cisco 4948 48-port 1Gb Ethernet Switch (SSH version 2 only supported). Additional Oracle Exadata Storage Expansion Racks can be added as required.

Note: Oracle Sun ZFS Storage 7320, Oracle Sun ZFS Storage 7420, and Oracle ZFS Storage ZS3-ES appliances provide a two-node cluster configuration. To discover the storage appliance, the administrative interfaces of both nodes must be private so that each node has a different static IP address. To verify that the appliance's nodes are using private administrative interfaces, you must use the appliance's user interface. For steps on how to determine if the interface is private, see Oracle ZFS Storage Appliance.

Viewing the Oracle SuperCluster System

You can view the Oracle SuperCluster system virtually using the tabs on the center pane, namely Dashboard, Details, Networks, and Incidents. You can also perform actions by clicking the respective actions on the Actions pane. The actions available on the Actions pane depend on the selected asset in the Navigation pane. The actions are

also role based, not all users are allowed to perform all actions. For more information on Roles, see Understanding User Roles.

> **Note:** In this document, it is assumed that Oracle SuperCluster is configured and discovered in Oracle Enterprise Manager Ops Center.

To view the Oracle SuperCluster system, perform the following steps:

1. Log in to Oracle Enterprise Manager Ops Center using the Oracle SuperCluster Systems Admin role.
2. In the Navigation pane, under Assets, expand **Oracle Engineered Systems**.

Figure 23–18 Navigation to SuperCluster System

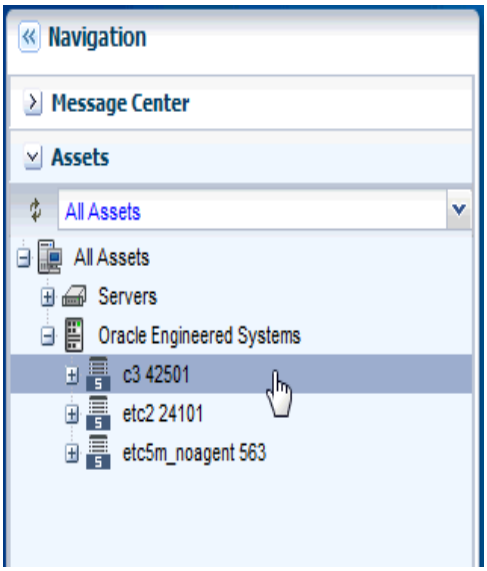

3. Select the Oracle SuperCluster system.

   The Dashboard, Details, Networks, and Incidents tabs are displayed in the center pane.

Figure 23–19 SuperCluster System Center Pane

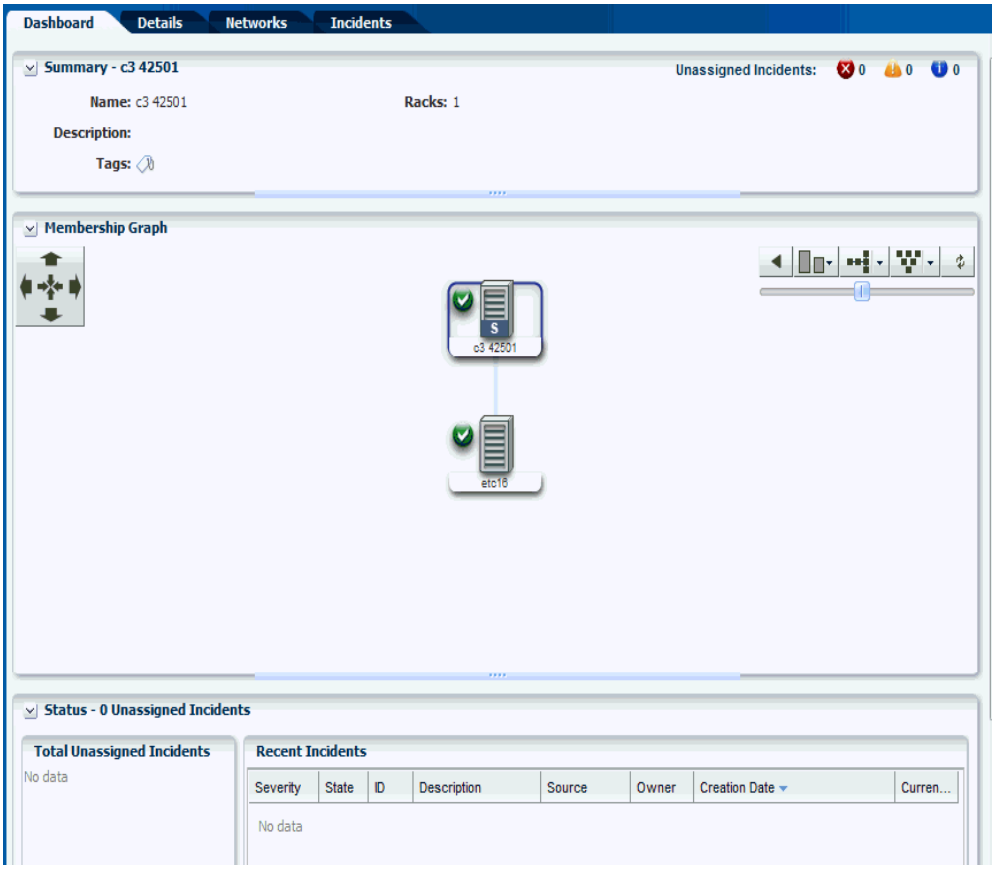

4. Click the required tab on the center pane to view more details.
   - **Dashboard Tab** displays the summary of the SuperCluster system, its membership graph, and the status of the system with details of all incidents.
   - **Details Tab** displays the name and description of the Oracle SuperCluster system. It also displays the number of racks, compute nodes, storage nodes, Exadata storage servers switches, PDUs, and fabrics present in the system.
   - **Networks Tab** displays the Infrastructure Networks table and Network Connectivity table.
   - **Incidents Tab** displays all the unresolved incidents and alerts reported in the Oracle SuperCluster system.

Dashboard Tab

The dashboard includes three sections namely, Summary, Membership Graph, and Status.

- Summary Pane
- Membership Graph Pane
- Status Pane

Summary Pane

The Summary section displays the name of the Oracle SuperCluster system, description or the name of the system identifier, number of racks that are part of the

system, and the number of unassigned incidents. The Unassigned Incidents icon in the summary pane includes incidents resulting from hardware faults. The three icons depict Critical Incidents, Warnings, and Information Incidents respectively.

The incidents are from all the assets that belong to a single Oracle SuperCluster Engineered System that includes multiple racks, or even systems with common IB backbone. If there are two systems, click each of them to display only incidents for assets that belong to that particular engineered system.

Membership Graph Pane

The Membership Graph pane displays the Oracle SuperCluster system as a hierarchy of its components showing the relationship between the Oracle SuperCluster system and fabrics that are grouped in the rack. You can instantly navigate to any asset by double-clicking on the asset in the membership graph.

Using the controls on the top right of the graphic pane, you can change the view of the graph to either a horizontal or a vertical orientation. You can also refresh the view by clicking the Refresh icon. You can also change the graph depths or size of the images.

Status Pane

The Status Pane displays the total unassigned incidents in chart format and also the recent incidents encountered by the Oracle SuperCluster system. In the Recent Incidents section, the following details are displayed:

- Severity displays the severity of the incident.
- State displays the state of the incident.
- ID displays the incident ID.
- Description displays the incident description.
- Source displays the source of the incident.
- Owner identifies the admin that is assigned to the incident and responsible for taking corrective action on the incident.
- Creation Date displays the date on which the incident was created.
- Current Status displays a current status icon of the incident, such as Critical / Warning / Info / Cleared.

Details Tab

In the Details Tab, the following details are displayed:

- Name of the Oracle SuperCluster System
- Description
- Master Subnet Manager
- Number of Racks
- Number of Compute Nodes
- Number of Storage Nodes
- Number of Exadata Storage Servers
- Number of Switches
- Number of PDUs
- Number of Fabrics

In the Tags pane, the tag names and their values are displayed. You can search for a particular tag using the Search feature.

Networks Tab

The Networks Tab displays the Infrastructure Networks Table and Network Connectivity Table.

Infrastructure Networks Table

The Infrastructure Networks Table displays the infrastructure networks that are defined and used inside the Oracle SuperCluster system for communication between the Oracle SuperCluster control components.

- Network Name specifies the name of the managed network.
- Network CIDR specifies the network.
- Partition Key specifies the IB network partition key.
- IP Range specifies the minimum and maximum boundaries of the IP addresses assigned to the network.
- Roles specifies the role of the network.

Network Connectivity Table

The Network Connectivity Table displays the infrastructure networks with IPs assigned to the individual hardware components of the Oracle SuperCluster system.

- Network Name
- Asset Type

Incidents Tab

The Incidents Tab displays all the unresolved incidents and alerts reported in the Oracle SuperCluster System.

Viewing the Oracle SuperCluster System Rack

This section describes how to view the Oracle SuperCluster System rack, visualization of the rack physical layout, aggregated rack components, energy data, and other rack details. As an administrator, you can drill down to an asset contained in the rack (server, storage node, switch, Exadata Cells) or even further.

To view the Oracle SuperCluster System rack, perform the following steps:

1. In the Navigation pane, under Assets, expand **Oracle Engineered Systems**.

Figure 23–20 Navigation to SuperCluster System Rack

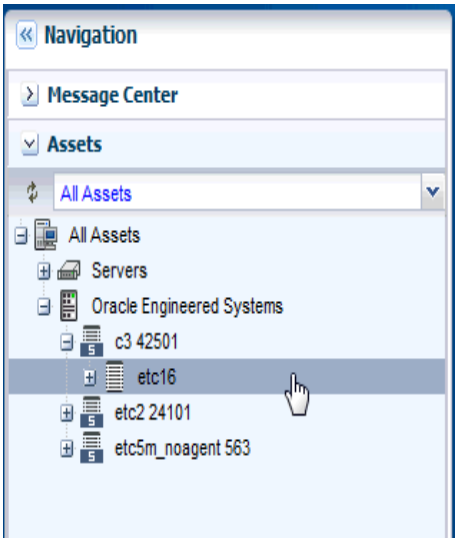

2. Select a SuperCluster system, then select the rack that you want to view.

   The Dashboard, Details, Firmware, Incidents, Charts, Energy, and Jobs tabs are displayed in the center pane.

Figure 23–21 SuperCluster System Rack Center Pane

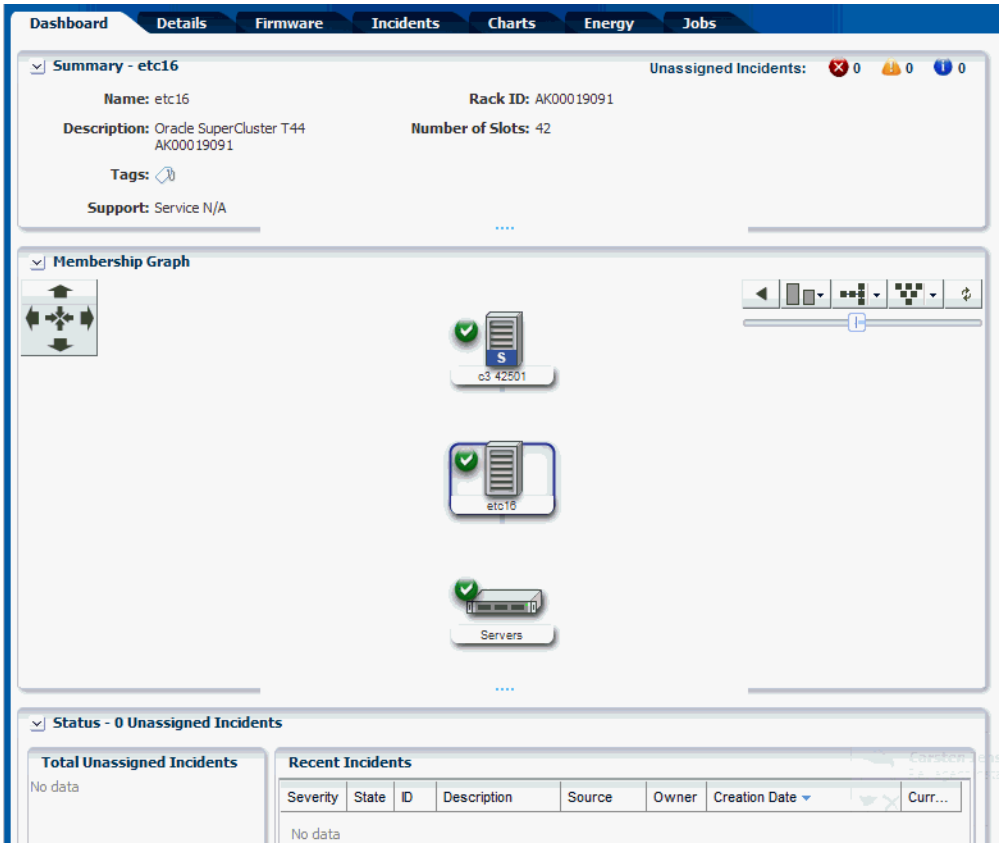

- **Dashboard Tab** displays the summary, membership graph, and status of the rack.

- **Details Tab** displays the rack info with a photorealistic view of the rack.
- **Firmware Tab** displays details of the Compute Nodes, Switches, Storage Appliances, Exadata Storage Servers, and Power Distribution Units in the SuperCluster system
- **Incidents Tab** displays the unresolved incidents and alerts for the selected SuperCluster system.
- **Charts Tab** displays the chart of the aggregate power usage of the selected SuperCluster system. It provides more ways to display the power utilization data. You can change the graphed data to a bar chart or an area chart. You can also export the data for either the current view or all available data to a file in either CSV or XML format.
- **Energy Tab** reports details of the Energy Performance such as Aggregate Power Consumption, Top and Bottom Consumers of Power and CPU Resources, and Average Power Consumption and CPU Utilization of the selected SuperCluster system rack.
- **Jobs Tab** displays the current and historical jobs.

3. Click any of the tabs to view more details of the rack.

Viewing the Exadata Storage Server

The Exadata Storage Server provides storage for the Oracle SuperCluster system. It is a sub-type of the Linux server installed to be a database node.

The Exadata Storage Servers are grouped together in the Navigation pane. In the same way that you view the Oracle SuperCluster system and its rack, you can also view the Exadata Storage Servers.

To view the Exadata Storage Server, perform the following:

1. In the Navigation pane, under Assets, expand **Oracle Engineered Systems**.
2. Select the Oracle SuperCluster asset.
3. Expand **Exadata Storage Servers**, and select an Exadata Storage Server. The server details are displayed on the center pane. Select the respective tabs on the center pane to view or perform actions.

Figure 23–22 View of Exadata Storage Servers

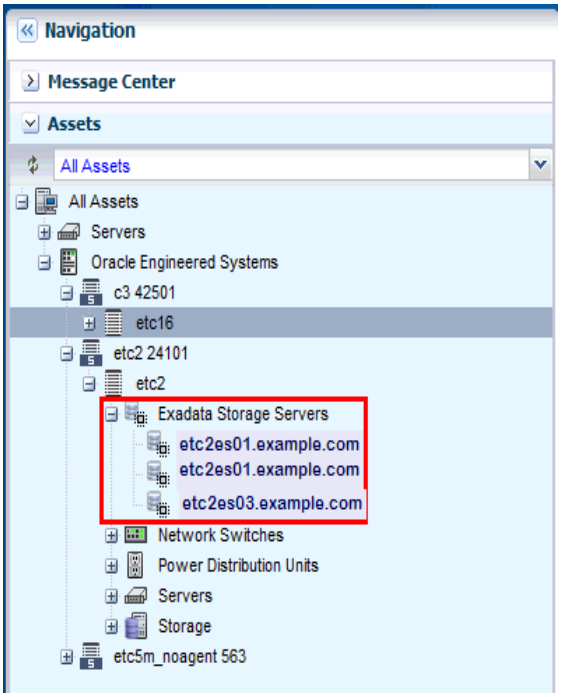

Refer To `Hardware Monitoring` section in the *Hardware Management Guide* for information on the different views and tasks that can be performed on the server cells.

Discovering an Oracle SuperCluster Component

Manual discovery is necessary if a component has been replaced or some properties of a component has changed autonomously. The component can either be a compute node, storage, or switch.

To manually discover an Oracle SuperCluster component, perform the following:

1. In the Navigation pane, select **Plan Management**.
2. Select Profiles and Policies, then select **Discovery**.

   The available discovery profiles are listed in the center pane.
3. Select the required stored discovery profile and edit the profile and metadata if required.
4. Click **Finish**. The profile is updated.
   Run the discovery using the updated profile to discover the Oracle SuperCluster component. See Discovering Oracle Engineered Systems.

Embedded Engineered Systems Management

Using Oracle Enterprise Manager Ops Center, you can discover and manage Oracle Engineered Systems. You can view and access multiple Engineered Systems from a single datacenter through Oracle Enterprise Manager Ops Center. You can also view the Engineered System's assets, incidents reported from the Engineered Systems, and Service Requests from the same UI. You can generate reports for all datacenter assets, including Engineered Systems.

The discovery and management process is performed using one of the following methods:

- Finding Assets
- Adding Assets Using a Discovery Profile

Finding Assets

The Find Assets action is disabled by default. Refer Service Tag Discovery for more information.

You can find assets using their service tags. Service tags are XML files containing product information. Many Oracle systems come equipped with service tags. You can discover hardware assets equipped with service tags using the Find Assets Wizard. This method lets you discover large number of assets quickly.

The Find Assets Wizard searches IP ranges or networks for service tags, then uses credentials that you specify to discover and manage the discovered assets.

> **Note:** Products without service tags cannot be discovered using this method. By default, service tags are enabled only on Oracle Exalogic Elastic Cloud Engineered Systems.

To find assets, perform the following steps:

1. In the Navigation pane, under Assets, select **All Assets**.
2. In the Actions pane, click **Find Assets**. The Find Assets Wizard opens.

Figure 23–23 Find Assets

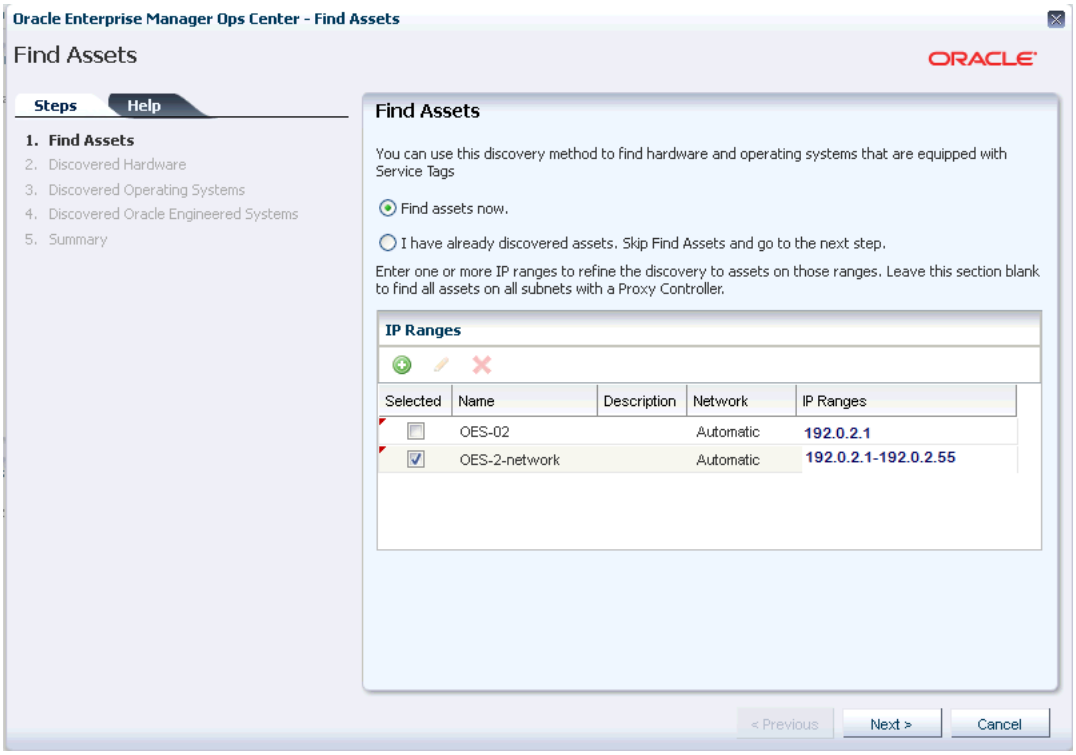

3. Select **Find Assets Now** and click **Next.**

   When the initial discovery is complete, select the assets in each category (hardware, operating systems, and engineered systems) that you want to manage and provide necessary credentials.

   Note: When adding a network range, enter the values without any space between the IP addresses. Use the format *start_ip-end_ ip*.

4. Click **Finish**. The manage asset job is launched.

See Discovering and Managing Assets for more information.

Adding Embedded Engineered System Assets

Embedded Oracle Engineered Systems is managed by a dedicated Oracle Enterprise Manager Ops Center running inside the engineered system fabric. You can add and manage a single or multiple Embedded Engineered System assets using Oracle Enterprise Manager Ops Center. To discover or add an asset, create a discovery profile. After you create the discovery profile, you can use the discovery profile to add the assets.

Creating a Discovery Profile

To create a discovery profile, perform the following steps:

1. In the Navigation pane, select **Plan Management**.
2. Under Profiles and Policies, select **Discovery**.
3. In the Actions pane, click **Create Profile**.
4. Enter a Name and Description for the discovery profile (description is optional).
5. In the Asset Type, expand **Oracle Engineered Systems**, then select **Embedded Oracle Engineered System**.

Figure 23–24 Create Discovery Profile

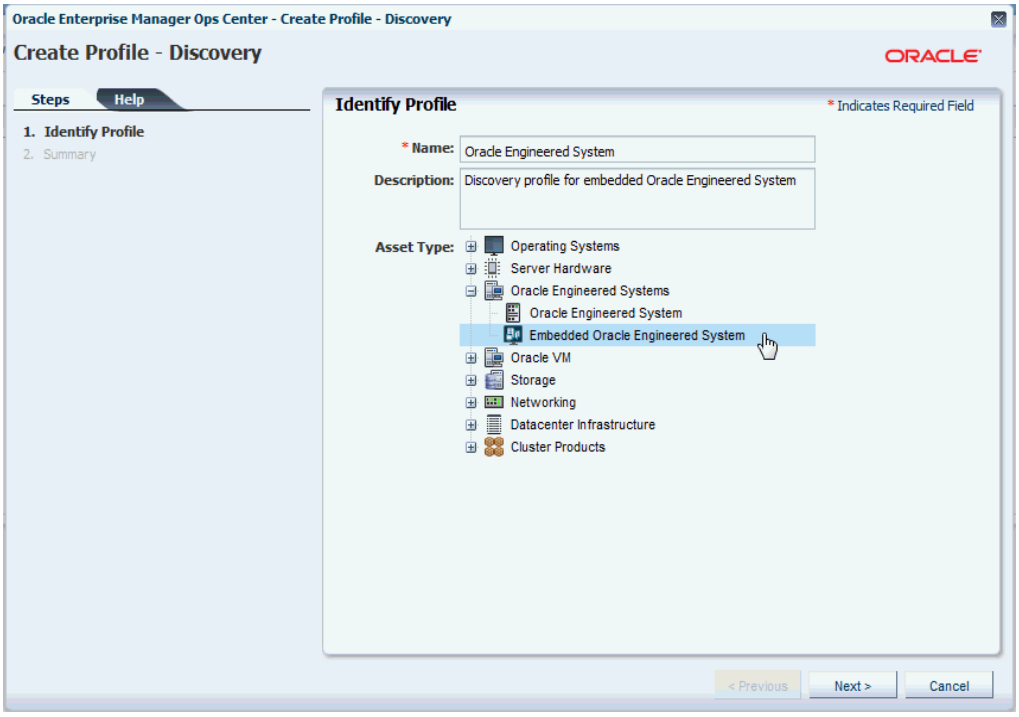

6. Click **Next**.
7. Enter Tags as necessary, then click **Next**.
8. (Optional) Enter the IP range. To use the existing IP range, select it.
   - Host names or IPs: Enter a comma-separated list of host names or IP addresses. To target an IP range, use the IP Ranges field instead.
   - Network: The managed network associated with the host names or IP addresses to route the discovery to the correct Proxy Controller. Select Automatic to route the job to the most appropriate Proxy Controller. The IP address of a target must resolve to only one known network for automatic routing to succeed.
   - IP Ranges: Click the plus icon to add one or more IP ranges. Enter a name, description (optional), network, and IP range for each. To target specific host names or IP addresses, you must enter them when you run the discovery profile.
9. Enter credentials for the JMX protocol.

   The Java Management Extension (JMX) protocol is exclusively used to discover engineered system assets.
10. Review the summary and click **Finish** to create the discovery profile.
    Refer Asset Management for additional information on discovering and adding assets in Oracle Enterprise Manager Ops Center.

Adding Assets Using a Discovery Profile

After you have created a discovery profile, you run a discovery job and manage Engineered System's assets. This lets you discover assets using a pre-existing set of protocols and credentials, and manage assets consistently.

To declare unconfigured assets for service processor configuration, perform the following steps:

1. In the Navigation pane, under Assets, select **All Assets**.
2. In the Actions pane, click **Add Assets**. The Add Assets Wizard opens.
3. Select the option, **Add and manage various types of assets via discovery probes**. The Discovery profile window opens.

Figure 23–25 Add Assets

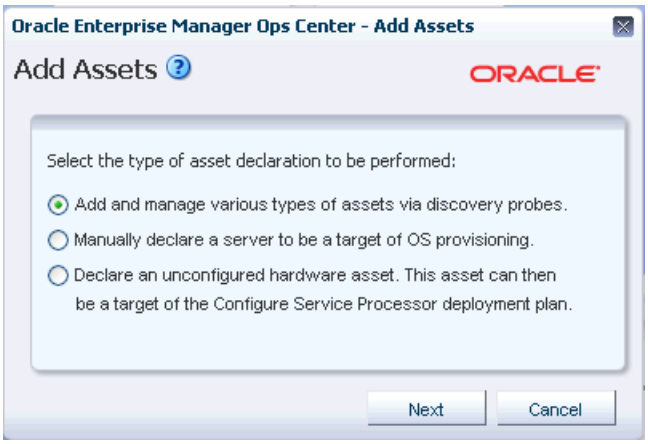

4. Select a discovery profile that you created for discovering the embedded engineered system asset.

Figure 23–26 Select Discovery Profile

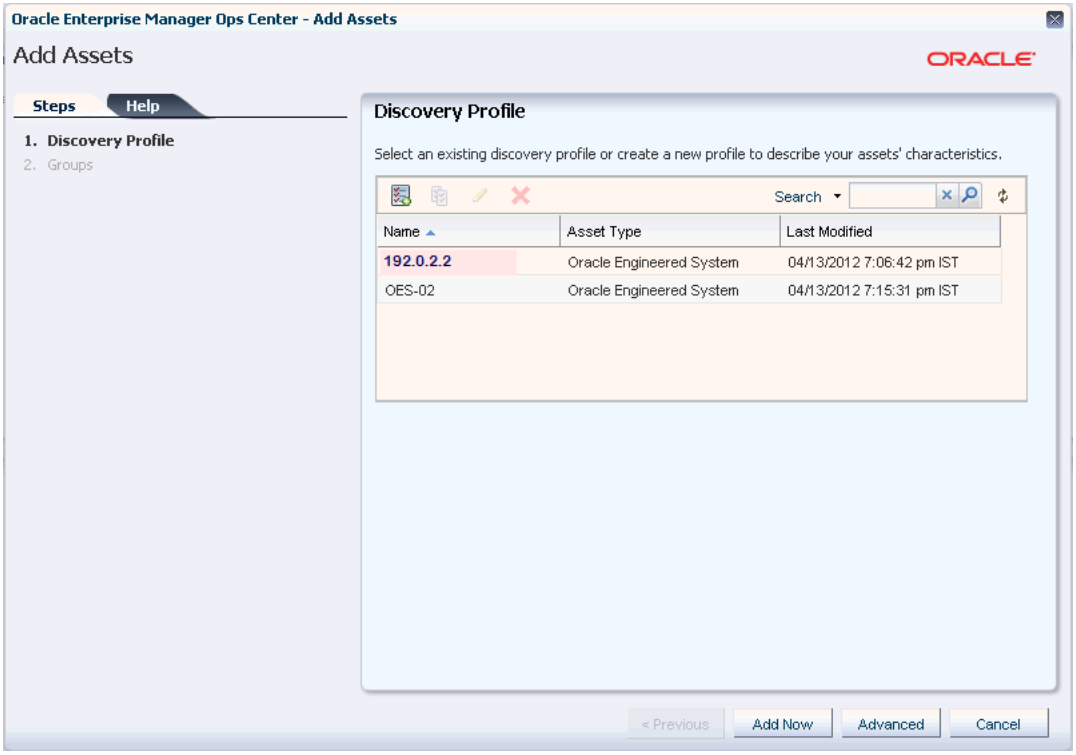

5. Add or edit the IP addresses, host names, and credentials for the targets.
6. Click **Add Now**. The Add Assets job is run. This might take a few minutes.

Viewing the Embedded Engineered Systems

You can view the Embedded Engineered System assets in the Datacenter Enterprise Controller BUI.

To view the Engineered System assets in the Datacenter Enterprise Controller BUI, perform the following steps:

1. In the Navigation pane, under Assets, expand **Embedded Oracle Engineered Systems**.
2. Select the Engineered System that you want to view. The Dashboard, Details, Rack Info, Incidents, Service Request, and Monitoring tabs are displayed in the center pane.
   - **Dashboard Tab**: Displays the summary of the selected Engineered System, its membership graph, and the status of the system.
   - **Details Tab**: Displays the name of the selected Engineered System, description, master subnet manager address, number of racks, compute nodes, storage nodes, switches, and PDUs in the system.
   - **Rack Info Tab**: Displays a photorealistic view of the selected Engineered System.
   - **Incidents Tab**: Displays all incidents reported from the selected Engineered System.
   - **Service Request Tab**: Displays the service request for the selected Engineered System.
   - **Monitoring Tab**: Displays the alert monitoring rules and service dependencies of the selected Engineered System.

To perform active management actions on any Oracle Engineered System asset, click **Launch Web Interface** in the Actions pane to launch the web interface of the selected Oracle Engineered System.

Viewing Relayed Incidents

Relayed Incidents are datacenter incidents that are reported from Oracle Engineered Systems attached to Datacenter Enterprise Controller. You can only view and delete the relayed incidents. No actions can be performed on these incidents.

When you close the relayed incidents from the Datacenter Enterprise Controller, the incidents are closed in the respective Engineered System as well.

To view all datacenter incidents, perform the following:

In the Navigation pane, under Message Center, select **Relayed Incidents**.

All the relayed incidents are displayed.

See Incidents for more information.

Viewing Relayed Service Requests

You can view the relayed service requests related to the Oracle Engineered Systems. You can also file service requests.

To view the relayed service requests, perform the following:

In the Navigation pane, under Message Center, select **Relayed Service Requests**.

All the relayed service requests are displayed. Details such as summary of the request, SR number, serial number, severity, status, contact person, and last updated date are displayed.

See `Oracle Services and Service Requests` for more information.

Baseline Check

Baseline Check is a feature in Oracle Enterprise Manager Ops Center Engineered Systems where the factory default configuration of eighth, quarter, half, and full rack configurations are considered as a normal or baseline setup. When the assets are discovered and associated with the rack, this setup is compared with the default factory configuration. If any discrepancy is found, the baseline check displays a question mark on the slot in the rack where the local configuration differs from the factory default configuration.

Hover the cursor over the question mark to view details of the warning, for example, missing asset, asset mismatch, or unknown asset. Based on the warning, you can decide to adjust and move the assets in the rack.

Related Resources for Engineered Systems

For instructions on performing general Oracle Enterprise Manager Ops Center actions, see the following resources.

- See Chapter 2, "Asset Management"
- See Chapter 5, "Software Libraries"
- See Chapter 11, "Hardware"
- See the Deploy How To library at `http://docs.oracle.com/cd/E40871_01/nav/deployhowto.htm`.
- See the Operate How To library at `http://docs.oracle.com/cd/E40871_01/nav/operatehowto.htm`.

Part VI

Reference

Part VI contains the following appendices:

- Appendix A, "Oracle Solaris Cluster"
- Appendix B, "Logs and Directories"
- Appendix C, "JumpStart Enterprise Toolkit"
- Glossary

For more information about modifying the product software, see the *Oracle Enterprise Manager Feature Reference Appendix Guide*.

A

Oracle Solaris Cluster

Use Oracle Enterprise Manager Ops Center to install a new Oracle Solaris Cluster and to maintain existing Oracle Solaris Clusters.

For information about Oracle Solaris Cluster 3.2, see the product documentation at http://docs.oracle.com/cd/E19787-01/index.html

Starting in Release 12.2.2.0.0, you can launch the Oracle Solaris Cluster 4.2 Manager:

1. In the Navigation pane, select **Solaris Cluster** and then select a cluster.
2. In the Actions pane, click **Launch Web Console**.
3. Log into the Oracle Solaris Cluster Manager.

Upgrading an Oracle Solaris Cluster

The cluster update procedures operate on complete clusters. Do not attempt to upgrade only a cluster node.

To upgrade a cluster to a newer version of the Oracle Solaris Cluster software, use a profile that defines the new version. Either select the target cluster and then the deployment plan or select the deployment plan and then the target. In either case, you use the Solaris Cluster Upgrade Job Wizard to specify how the deployment operates.

Before You Begin

Import the cluster profile.

To Upgrade the Cluster

1. Expand Plan Management in the Navigation pane.
2. In the Profiles and Policies section, click Update Profiles.
3. Select **Oracle Solaris Upgrade Job** in the Actions pane. The Oracle Solaris Upgrade Job Wizard starts.
4. Enter a name for the upgrade job.
5. Click display the cluster targets.
6. Choose one or more clusters and click Select/Add to Target List.
7. Select Dual Partition Mode.
8. Select the policy to use if a task in the upgrade cannot be completed.
9. Select the profile that you imported.

10. Review the summary of the job and click Finish. The upgrade job is submitted. To follow the progress of the job, click the View Job Details icon in the jobs pane.

Cluster Profiles

You use cluster profiles in a deployment plan to perform the following operations:

- Install Oracle Solaris Cluster software
- Upgrade Oracle Solaris Cluster software

The profiles and deployment plans rely on pre-action and post-action scripts to suspend cluster operations and save configuration information during the update process.

The following cluster-specific profiles and scripts are located in Enterprise Manager Ops Center's Local Content software library. As the Oracle Solaris Cluster software is revised, new profiles are made available to you through the Java.net site's Ops Center Cluster Profiles project. You download these new profiles and scripts, then import them into the Local Content library. See Obtaining the Cluster Profiles and Scripts for instructions to get these files.

- Profiles for Provisioning Oracle Solaris Cluster Software version 3.2u3
  - SPARC
    - sc-3.2u3-core-sparc
    - sc-3.2u3-manager-sparc
    - sc-3.2u3-agents-sparc
  - x86
    - sc-3.2u3-core-x86
    - sc-3.2u3-manager-x86
    - sc-3.2u3-agents-x86
- Profiles for Upgrading Version 3.2u3 to Version 3.3 or for Provisioning an Oracle Solaris Cluster Software Version 3.3
  - SPARC
    - sc-3.3-core-sparc
    - sc-3.3-manager-sparc
    - sc-3.3-agents-sparc
  - x86
    - sc-3.3-core-x86
    - sc-3.3-manager-x86
    - sc-3.3-agents-x86
- Scripts for pre-action and post-action
  - Post-action script for provisioning either version 3.2 or 3.3: `SolarisCluster-Post.ksh`
  - Post-action script for upgrading version 3.2 to 3.3: `SolarisCluster3.21109-Post`
  - Pre-action script for upgrading version 3.2 to 3.3: `SolarisCluster3.21109-Pre`

Obtaining the Cluster Profiles and Scripts

To obtain profiles and scripts for a new release of Oracle Solaris s Cluster software, download them to the system that is running the Enterprise Controller.

> **Note:** When you download, transfer, upload, or import these files, do not change the file name. Keep the same name throughout the procedure.

To Obtain the Cluster Profiles and Scripts

1. Go to `http://java.net/projects/oc-cluster-profiles`
2. Click the **Downloads** button.
3. Select the folder for the profiles you want:
   - To provision Oracle Solaris Cluster Software version 3.2u3, select the **Version3.2u3** folder.
   - To upgrade version 3.2u3 to version 3.3, select the **Version3.3** folder.
   - To provision Oracle Solaris Cluster Software version 3.3, select the **Version3.3** folder.
4. In the version folder, select the platform you want, either SPARC or x86.

   The platform folder contains three files: Agents, Core, and Manager. You need all three files to complete upgrade or provision operations.
5. Use the browser window's **Save** feature to save each file, keeping its file name.
6. Return to the home page for the project and select the **Action-Scripts** folder.
7. Select the scripts that complete the core profiles:
   - To upgrade version 3.2u3 to version 3.3, select the Pre-Upgrade and Post-Upgrade scripts.
   - To provision version 3.2u3 or version 3.3, select the Post-Provision script.
8. Upload the profiles and scripts, according to Uploading the Cluster Profiles and Scripts.

Uploading the Cluster Profiles and Scripts

The files containing the profiles and scripts must reside on the same system that is running the browser and they must reside in the same directory. This operation uploads all files in the selected directory.

1. Move the profiles and scripts to the appropriate system and directory.
2. Expand **Libraries** in the Navigation pane.
3. Click **Local Content** in the Solaris/Linux OS Updates library.
4. Click **Bulk Upload Packages and Patches** in the Actions pane.
5. Click **Distribution** to select the distribution that applies to these files.
6. Select **Upload from Directory**.
7. Specify the path to the directory or click **Browse** to locate and select it.

   All files in the directory and its subdirectories are uploaded.

8. Click **Submit**.
9. Import the profiles into Enterprise Manager Ops Center, according to Importing Cluster Profiles and Scripts.

Importing Cluster Profiles and Scripts

Before You Begin

- To import profiles, use the browser on the same system or transfer the profiles to the system on which you are running the browser.
- To import scripts, transfer the files to the Enterprise Controller's system.

> **Note:** When you download, transfer, upload, or import these files, do not change the file name. Keep the same name throughout the procedure.

To Import a Cluster Profile

1. Expand **Plan Management** in the Navigation pane.
2. Click **Update Profiles**.
3. Click the **Import Profile** icon.
4. For each profile and script:
   a. Enter the name of the file you downloaded. Do not change the name.
   b. Click **Import**.
5. Expand **Libraries** in the Navigation pane.
6. Select **Solaris/Linux OS Updates**.
7. Select **Local Content**.
8. Click **Upload Local Action**.
9. For each profile:
   a. Enter the name of the profile as you want it to be displayed in the Plan Management section.
   b. Select the distribution of the OS that uses the profile.
   c. Enter the name of the file that contains the profile. Do not change the name.
   d. Click **Upload**.
10. For each script:
    a. Select action type, the matching parent category, and the OS distribution.
    b. Enter the name of the file that contains the script. Do not change the name.
    c. Click **Upload**.
11. For provisioning profiles, edit the core profiles to include the Post-Provision scripts, according to Editing the Core Profile for Provisioning

    To upgrade a cluster, no modification is needed because the scripts are available to the upgrade profile after the import operation.

Editing the Core Profile for Provisioning

1. Expand **Plan Management** in the Navigation pane.
2. Click **Update Profiles**.
3. In the OS Update Profiles table, select the core profile.
4. Click the **Edit Profile** icon.
5. In the Edit OS Update Profile window, navigate from Local to either Post-actions or Pre-action.
6. Select the script and click **Required**. For example, to edit the core profile for provisioning, navigate to Post-actions, select `SolarisCluster-Post.ksh`, and click Required. The file is now listed in the Profile Contents section.
7. Click the **Saved as Named Profile** button.

B

Logs and Directories

Oracle Enterprise Manager Ops Center performs each action as a job. The details of a job show the order of tasks in the job and the managed assets that are targets of the job. You can view the details of a job from either the browser or the command-line interface. Each job is stored until it is deleted explicitly. See Viewing Jobs for instructions.

In addition to the job, log files record events of different types and for different purposes. Some log files are protected by file permissions and require a user with root access to view them. Some log files can be displayed in the product's browser interface, using the following procedure:

1. Click the Enterprise Controller in the **Administration** section of the Navigation pane.
2. Click the **Logs** tab in the center pane.
3. Select a log from the drop-down list:
   - cacao log
   - UI log
   - Proxy log
   - Update error log
   - Update channel download log
   - Update channel error log
4. (Optional) Click **Refresh Log File** to update the display.

Installation

- Log of the most recent installation or uninstallation: `/var/tmp/opscenter/installer.log.latest`
- Log of previous installation or uninstallation operations: `/var/tmp/opscenter/installer.log.xxxx`
- Log of a specific installation:

  `/var/opt/sun/xvm/oracle/app/oraInventory/logs/silentInstall<yyyy-mm-dd-hh-mm-sspm>.log`
- Log of an agent installation: `/var/scn/install/log`

Upgrades

The log of upgrade actions are in these files:

- Enterprise Controller: `/var/opt/sun/xvm/update-saved-state/update_EC_minor_bundle_12.2.`*`n.xxxx`*`/updatelog.txt`
- Co-located Proxy Controller: `/var/opt/sun/xvm/update-saved-state/update_PC_minor_bundle_12.2.`*`n.xxxx`*`/updatelog.txt`
- Remote Proxy Controller: `/var/scn/update-saved-state/update_proxy_bundle_12.2.`*`n.xxxx`*`/updatelog.txt`

If an upgrade fails, the database rolls back and a log of database actions is stored in the following directory: `/var/opt/sun/xvm/update-saved-state/update_EC_minor_bundle_12.2.`*`n.xxxx`*`/dblogs` directory.

Connection, Job, and User Account Activity

Oracle Enterprise Manager Ops Center records events for several purposes such as for performance, security, and diagnosing problems. The various log files contain some of the same events and different levels of detail.

Performance and Security

The audit log files record the following types of events:

- Adding and deleting a user account
- Changing the roles for a user account
- Login activity and information about the connection
- Starting and ending jobs

The files are located on the Enterprise Controller in the following location:

- On Oracle Solaris: `/var/cacao/instances/oem-ec/logs/audit-logs.*`
- On Linux: `/var/opt/sun/cacao2/instances/oem-ec/logs/audit-logs.*`

Each audit log file has a maximum size of 10 Mb. When this limit is reached, the file is closed and a new file is created with an incremented file extension. The maximum number of audit log files is 15, accumulating 150 Mb of logged activity. When `audit-logs.14` is closed, the next audit file is `audit-log.0`, overwriting the original `audit-log.0` file. Figure B–1 shows the series of log files.

Figure B–1 Contents of Log Directory on Oracle Solaris 11

```
root@ocbrm-ipgs15:/var/cacao/instances/oem-ec/logs# ls -l
total 64146
-rw-r--r--   1 root     sys        39173 Apr 29 12:06 audit-logs.0
-rw-r--r--   1 root     sys            0 Apr 23 16:02 audit-logs.0.
-rw-r--r--   1 root     sys      2456675 Apr 29 13:41 cacao.0
-rw-r--r--   1 root     sys            0 Apr 23 16:01 cacao.0.lck
-rw-r--r--   1 root     sys     10000142 Apr 29 00:21 cacao.1
-rw-r--r--   1 root     sys     10000082 Apr 26 22:46 cacao.2
-rw-r--r--   1 root     sys     10000092 Apr 24 23:59 cacao.3
```

The entries in the audit log file have the following syntax:

`datetime action connect_info additional_info`

where

action
LOGIN
DISCONNECT If a connection expires, the disconnection is not logged.
JOB_START
JOB_END
USER_ADD
USER_DELETE
ROLES_ASSIGN
SCHEDULED_JOB_STARTED
REMOTE_INFO Indicates a connection through the browser interface and includes the IP address and port of the `http` client making the connection, as in the following example:

```
REMOTE_INFO rmi://127.0.0.1 yogi 52, Remote Info: User yogi connected from
10.157.134.249:57391 / JMX Session: com.sun.cacao.sessionrmi://127.0.0.1:9
com.sun.cacao.useryogi
```

Some changes to the domain model are also recorded:

- Refresh
- Set
- Unregister

connect_info
Unique identifier for the connection, depending on the type of connection:

- Connections through the browser interface or the command line interface: `rmi://`*`ip_address username connection_id`*
- Connections through the API: `jmxmp://`*`ip_address:port username connection_id`*

additional_info

- For job actions, the additional information is the job ID, which consists of the Enterprise Controller's name and the job number as listed in the Job pane.
- For user actions, the additional information is the username.

Example B–1 shows the contents of an audit log for the following operations:

- User root logs in at 3:06.
- User root creates a new user, stanfield.
- User root gives the OPS_CENTER_ADMIN privilege to user stanfield.
- User root logs out.
- User stanfield logs in at 3:12.
- User stanfield starts a DHCP configuration job.
- Job is completed.
- User stanfield logs out.

Example B–1 Example of an Audit Log

```
5/23/14 3:06 PM LOGIN rmi://127.0.0.1 root 13
5/23/14 3:06 PM REMOTE_INFO rmi://127.0.0.1 root 13, Remote Info: User root
connected from 192.0.2.1:45338 / JMX Session:
com.sun.cacao.session^Armi://127.0.0.1:2 com.sun.cacao.user^Aroot
```

```
5/23/14 3:12 PM USER_ADD rmi://127.0.0.1 root 13, Remote Info: User root connected
from 192.0.2.1:45338 / JMX Session: com.sun.cacao.session^Armi://127.0.0.1:2
com.sun.cacao.user^Aroot Add user stanfield: SUCCESS
5/23/14 3:12 PM ROLES ASSIGN rmi://127.0.0.1 root 13 Roles [OPS_CENTER_ADMIN]
granted to user stanfield
5/23/14 3:12 PM DISCONNECT rmi://127.0.0.1 root 13
5/23/14 3:12 PM LOGIN rmi://127.0.0.1 stanfield 18
5/23/14 3:12 PM REMOTE_INFO rmi://127.0.0.1 stanfield 18, Remote Info: User
stanfield connected from 192.0.2.1:45351 / JMX Session:
com.sun.cacao.session^Armi://127.0.0.1:3 com.sun.cacao.user^Astanfield
5/23/14 3:13 PM JOB_STARTED rmi://127.0.0.1 stanfield 18
sm4170m2-11-n172.27.immediate - DHCP Server Configuration on sm4170m2-11-n172
5/23/14 3:13 PM JOB_END  Job sm4170m2-11-n172.27 Completed with Status: SUCCESS
5/23/14 3:13 PM DISCONNECT rmi://127.0.0.1 stanfield 18
```

Diagnosing Problems

The following log files contain detailed information about the same events as the audit log files except for login information. They include the interactions between components of the product software.

- On Oracle Solaris: `/var/cacao/instances/oem-ec/audits/`
- On Linux: `/var/opt/sun/cacao/instances/oem-ec/audits/`

The following log files are specialized for specific events:

- Messages from operating system such as Info and Warning: `/var/adm/messages*`
- Login and connection information: `/var/opt/sun/xvm/logs/audit-logs*`
- Events in the user interface component: `/var/opt/sun/xvm/logs/emoc.log`
- Events between controllers and agents:
  - On an Oracle Solaris Enterprise Controller: `/var/cacao/instances/oem-ec/logs/cacao.`*n*
  - On a Linux Enterprise Controller: `/var/opt/sun/cacao/instances/oem-ec/logs/cacao.`*n*
  - On each Oracle Solaris Proxy Controller: `/var/cacao/instances/scn-proxy/logs/cacao.`*n*
  - On each Linux Proxy Controller: `/var/opt/sun/cacao/instances/scn-proxy/logs/cacao.`*n*
  - On each Oracle Solaris agent: `/var/cacao/instances/scn-agent/logs/cacao.`*n*
  - On each Oracle Linux agent: `/var/opt/sun/cacao/instances/scn-agent/logs/cacao.`*n*

High Availability

In a High Availability configuration, each Enterprise Controller is a Clusterware node. The Clusterware resource activity is logged each time the active Enterprise Controller's resource action script's `check()` function is executed. The default interval is 60 seconds.

On Oracle Solaris: `/var/opt/sun/xvm/ha/EnterpriseController.log`

Software Update Component

The Software Update component has its own server. The following files record activity for this server:

- Audit Log
  - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/audit.log`
  - On Linux: `/usr/local/uce/server/logs/audit.log`
- Errors
  - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/error.log`
  - On Linux: `/usr/local/uce/server/logs/error.log`
  - Download jobs: `/opt/SUNWuce/server/logs/SERVICE_CHANNEL/error.log`
- Job Log
  - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/job.log`
  - On Linux: `/usr/local/uce/server/logs/job.log`

Agents

- Agent log:
  - On Oracle Solaris: `/var/cacao/instances/scn-agent/logs/cacao.`*n*
  - On Oracle Linux: `/var/opt/sun/cacao/instances/scn-agent/logs/cacao.`*n*
- Agent update log files: `/var/scn/update-agent/logs` directory after an update
- Other agent log: `/var/opt/sun/xvm/logs`

Local Database

- On the Enterprise Controller:
  - For installation events:

    `/var/opt/sun/xvm/oracle/cfgtoollogs/dbca/OCDB/*`

    `/var/tmp/opscenter/installer.log.latest`
  - For operational events, reported by the `ecadm sqlplus` utility:

`/var/opt/sun/xvm/oracle/diag/rdbms/ocdb/OCDB/alert/log.xml.*`
`/var/opt/sun/xvm/oracle/diag/rdbms/ocdb/OCDB/trace/alert_OCDB.log.*`
`/var/opt/sun/xvm/oracle/diag/tnslsnr/`*hostname*`/oclistener/alert/log.xml.*`
`/var/opt/sun/xvm/oracle/diag/tnslsnr/`*hostname*`/oclistener/trace/listener.log.*`

  - For schema changes:

    `/var/opt/sun/xvm/log/satadmsqlplus.log`

    `/var/opt/sun/xvm/logs/alter_oracle_schema.out`

    `/var/opt/sun/xvm/logs/alter_oracle_storage.out`
  - For backup, restore, and migrate operations:

    `/var/opt/sun/xvm/logs/sat-backup-`*date-time*`.log`

    `/var/opt/sun/xvm/logs/sat-restore-`*date-time*`.log`

/var/opt/sun/xvm/logs/migrate.log

  - For data files: /var/opt/sun/xvm/oracle/oradata/OCDB
  - For redo log files: /var/opt/sun/xvm/oracle/oradata/OCDB
- On the Proxy Controller: /var/opt/sun/xvm/proxydb/*
- On each agent: /var/opt/sun/xvm/agentdb/*

Controlling the Number of Common Agent Container Log Files

The Common Agent Container cacao is a common Java container for JDMX/JMX management and handles the interactions between controllers and agents. All events are recorded in the cacao log files. Any event above the level of INFO is also logged in the syslog. You can view the contents of the current log file using the UI or by viewing the contents of the following files:

- On an Oracle Solaris Enterprise Controller: /var/cacao/instances/oem-ec/logs/cacao.*n*
- On a Linux Enterprise Controller: /var/opt/sun/cacao/instances/oem-ec/logs/cacao.*n*
- On each Oracle Solaris Proxy Controller: /var/cacao/instances/scn-proxy/logs/cacao.*n*
- On each Linux Proxy Controller: /var/opt/sun/cacao/instances/scn-proxy/logs/cacao.*n*
- On each Oracle Solaris agent: /var/cacao/instances/scn-agent/logs/cacao.*n*
- On each Oracle Linux agent: /var/opt/sun/cacao/instances/scn-agent/logs/cacao.*n*

The maximum file size is 1 MB. When the limit is reached, the current log file is closed and a new one created. The default number of log files is three. You can change the number of log files that are retained, using the Common Agent Container's management utility, cacaoadm.

To view the current number of log files maintained for the Enterprise Controller, issue the following command on the system where the Enterprise software is running:

```
# cacaoadm get-param log-file-count -i oem-ec
log-file-count=3
```

To view the number of log files maintained for a Proxy Controller, issue the following command on the system were the proxy controller software is running:

```
# cacaoadm get-param log-file-count -i scn-proxy
log-file-count=3
```

To view the number of log files maintained for an agent use the following command on the system were the agent is running:

```
# cacaoadm get-param log-file-count -i scn-agent
log-file-count=3
```

To Change the Number of Log Files for an Enterprise Controller or Proxy Controller

To change the number of log files on the Enterprise Controller, a Proxy Controller, or both:

1. Verify that there are no active jobs.
2. Stop the Common Agent Container service on the Enterprise Controller.
   - On Oracle Solaris:

     ```
     # /opt/SUNWxvmoc/bin/satadm stop -w -v
     ```

   - On Linux:

     ```
     # /opt/sun/xvmoc/bin/satadm stop -w -v
     ```

3. Stop the Common Agent Container service on a Proxy Controller:
   - On Oracle Solaris:

     ```
     # /opt/SUNWxvmoc/bin/proxyadm stop -w -v
     ```

   - On Linux:

     ```
     # /opt/sun/xvmoc/bin/proxyadm stop -w -v
     ```

4. Specify the maximum number of log files to be retained in addition to the current log file. In the following example, the count of 10 specifies that nine log files of previous events are retained in addition to the log file for current events. On the Enterprise Controller:

   ```
   # cacaoadm set-param log-file-count=10 -i default
   ```

   On a Proxy Controller:

   ```
   # cacaoadm set-param log-file-count=10 -i scn-proxy
   ```

5. Start the Enterprise Controller.
   - On Oracle Solaris:

     ```
     # /opt/SUNWxvmoc/bin/satadm start -w -v
     ```

   - On Linux:

     ```
     # /opt/sun/xvmoc/bin/satadm start -w -v
     ```

6. Verify that the Enterprise Controller has been restarted completely before attempting other operations. For example, if you have stopped both the Enterprise Controller and a Proxy Controller, wait for the Enterprise Controller to restart before restarting each Proxy Controller.

   On an Oracle Solaris Proxy Controller:

   ```
   # /opt/SUNWxvmoc/bin/proxyadm start -w -v
   ```

   On a Linux Proxy Controller:

   ```
   # /opt/sun/xvmoc/bin/proxyadm start -w -v
   ```

7. Verify that all controllers have restarted completely before attempting other operations.

To Change the Number of Log Files for an Agent

To change the number of log files on an agent:

1. Verify that there are no active jobs.
2. Stop the Common Agent Container service on a the agent:
   - On Oracle Solaris:

     ```
     # /opt/SUNWxvmoc/bin/agentadm stop -v
     ```

   - On Linux:

     ```
     # /opt/sun/xvmoc/bin/agentadm stop -v
     ```

3. Specify the maximum number of log files to be retained in addition to the current log file. In the following example, the count of 10 specifies that nine log files of previous events are retained in addition to the log file for current events.

   ```
   # cacaoadm set-param log-file-count=10 -i scn-agent
   ```

4. Start the agent:
   - On Oracle Solaris:

     ```
     # /opt/SUNWxvmoc/bin/agentadm start -v
     ```

   - On Linux:

     ```
     # /opt/sun/xvmoc/bin/agentadm start -v
     ```

5. Verify that all controllers have restarted completely before attempting other operations.

C

JumpStart Enterprise Toolkit

Use JumpStart Enterprise Toolkit (JET) to extend the JumpStart installation functionality provided within the Oracle Solaris 9 and 10 operating systems.

JET is a framework designed to simplify and extend the JumpStart installation capabilities for provisioning. the Oracle Solaris 9 or 10 operating system. JET provides a set of helper scripts to simplify the use of Jumpstart for the installation of Solaris 10 and earlier on both SPARC and x86 servers. Oracle Solaris 11 does not use JET, instead, it uses the automated installer (AI).

The SUNWjet and JetFLASH packages are installed on the Proxy Controller during installation. See the JET page on the Oracle Technology Network at http://www.oracle.com/technetwork/systems/jet-toolkit/jet-toolkit-1614844.html for more information on JET, including additional packages available for download and a link to the user documentation.

Note: JET must run on a Proxy Controller that is running on an Oracle Solaris 9 or 10 operating system.

JumpStart Enterprise Toolkit Configuration File Location

The JET module parameters are available for use in OS profiles. See the `module.conf` configuration files that are associated with JET modules for information about parameters for specific JET modules. The configuration files are located in the `/opt/SUNWjet/Products` directory on the Proxy Controller. For example, the configuration files for the custom module is located in the following directory on the Proxy Controller: `/opt/SUNWjet/Products/custom/custom.conf`. You can review the parameters for these modules by looking at the `sample.template` file in the `/opt/SUNWjet/Templates` directory on a Proxy Controller.

SUNWjet Parameters

The main JET framework is supplied in a single SVR4 package called SUNWjet. This package contains everything necessary to do a standard Oracle Solaris installation using either bootp or dhcp.

When you specify JET parameters with an OS profile, the following parameters from the base_config JET module are automatically updated within the OS profile and must not be modified:

- `base_config_ClientArch`
- `base_config_ClientEther`

- base_config_client_allocation
- base_config_sysidcfg_network_interface
- base_config_sysidcfg_ip_address
- base_config_sysidcfg_netmask
- base_config_sysidcfg_nameservice
- base_config_sysidcfg_system_locale
- base_config_sysidcfg_terminal
- base_config_sysidcfg_timeserve
- base_config_sysidcfg_timezone
- base_config_sysidcfg_root_password
- base_config_sysidcfg_security_policy
- base_config_sysidcfg_protocol_ipv6

The following list describes the parameters that are associated with the base_config JET module. These parameters provide basic operating system configuration information. Values for many of these parameters use the term targetableComponent to represent the target system.

- base_config_client_allocation: The mechanism used to build this client. By default, the options listed in /opt/SUNWjet/etc/jumpstart.conf are used. Leave the value blank unless you need to do something different from the default for this specific client. If you are provisioning the Oracle Solaris 10 1/06 x86 release, set the value of this variable to GRUB to enable GRUB-based booting and installation.
- base_config_ClientArch: Kernel architecture, such as sun4u or x86. By default, this is set to the kernel architecture of the targetable component.

  Default Value: [targetableComponent:kernel_arch]
- base_config_ClientEther: Ethernet MAC address. By default, this is set to the Ethernet MAC address of the targetable component.

  Default Value: [targetableComponent:ethernet_mac_address]
- base_config_ClientOS: Version of the OS to be provisioned.

  Example: Solaris9_u7_sparc
- base_config_dedicated_dump_device: If set, the dumpadm utility configures the partition as a Dedicated Dump Device. See dumpadm(1M) for supported Operating Environments.
- base_config_defaultrouter: Value to use for /etc/defaultrouter.
- base_config_disable_sysid_probe: If set, skip the sysid step on the first reboot. This can significantly increase provisioning efficiency on systems that have many unused network adapters.

  Default Value: yes
- base_config_dns_disableforbuild: Delay DNS configuration until later. If DNS is not available in the build environment, set this variable to yes.
- base_config_dns_domain: DNS domain entry for the /etc/resolv.conf file.
- base_config_dns_nameservers: Space-separated list of IP addresses to use for DNS name server entries in the /etc/resolv.conf file.

- base_config_dns_searchpath : List of entries to go in the DNS search line in /etc/resolv.conf file.
- base_config_dumpadm_minfree: Set a limit so that crash dumps do not fill up the dump file system. See the dumpadm(1M) -m option for possible values.

  Example: 20000k
- base_config_enable_altbreak: If set, enable alternate break sequence.
- base_config_enable_rootftp: If set to any value, enable root FTP access.
- base_config_enable_rootlogin: If set to any value, enable network root login from telnet, rsh, and ssh.
- base_config_enable_savecore: If set to any value, enable save core for Solaris 2.6 systems.

  Default Value: yes
- base_config_grub_append: For Oracle Solaris 10 1/06 x86 systems, specifies additional options or arguments to pass to the GRUB bootloader.
- base_config_ipmp_networkifs: Space-separated list of interfaces to be defined under IPMP control. For each interface listed, define sets of variables to provide the netgroup, mode, test1, test2, netmask, host name, log-ip, host name2, and log-ip2 for the interface.

  Example: qfe0_qfe4!database-net l 10.0.0.1 10.0.0.2 24 oracle-db 10.0.0.3 apache 10.0.0.4
- base_config_networkifs: Space-separated list of additional network interfaces to be defined. For logical interfaces, use underscores (_) rather than colons (:). Use the format c*n*t*n*d*n*. For each interface listed, define sets of variables to provide the netname, netmask, host name, and IP address for the interface.

  Example: le1!netB 255.255.255.0 myhost-netB 192.168.1.0
- base_config_nfs_mounts: Space-separated list of remote NFS mount points. Use ? to separate the mount source from the mount target, as shown in the example.

  Example: fs?1.1.1.1:/fs
- base_config_nfsv4_domain: Set up the NFSv4 domain to prevent being prompted at first reboot. If not set, look first for the entry in base_config_dns_domain, and second for the domain value in /etc/default/nfs.
- base_config_noautoshutdown: If set to any value, disable power management.

  Default Value: pm_disabled
- base_config_nodename: Value to use for /etc/nodename if not the default host name.
- base_config_notrouter: If set to y, then disable IPv4 forwarding and create the /etc/notrouter file.
- base_config_ntp_servers: Space-separated list of names or IP addresses for the NTP servers. The first server has a prefer tag. This section places lines of the form: server [prefer] into the /etc/inet/ntp.conf file. For additional NTP control, use the custom module to deploy your own custom ntp.conf file.
- base_config_patchdir: Path to the patches. If blank, use information from the jumpstart.conf file and the IP address of the JET server. If your patch files are not stored on the JET server, then provide an NFS-style path to the location of the patches.

- base_config_poweroff_afterbuild: If set, shut down the system after the build completes.
- base_config_productdir: Path to the products. If blank, use information from the jumpstart.conf file and the IP address of the JET server. If your package files are not stored on the JET server, then provide an NFS-style path to the location of the packages.
- base_config_products: JET modules to provision.
- base_config_profile: Create a custom JumpStart profile. By default, if you leave this variable blank, the OS provisioning plug-in creates the /opt/SUNWjet/Clients/hostname/profile based on the other base_config_profile variables. Alternatively, you can create your own custom JumpStart profile. To use the profile that you created manually, set the base_config_profile variable to the name of the created profile. By default, the OS provisioning plug-in looks for the profile in the /opt/SUNWjet/Clients/hostname directory. To direct the plug-in to a profile in another directory, provide an absolute path name in the base_config_ profile variable.

  Note: If you are provisioning Oracle Solaris OS on x86 target hosts, you must create a custom JumpStart profile that deletes any existing partitions and point to that profile in the base_config_profile variable.

- base_config_profile_add_clusters: Space-separated list of cluster packages to add.
- base_config_profile_add_geos: Comma-separated list of geographical regions to add.

  Example: N_Europe, C_Europe
- base_config_profile_add_locales: Comma-separated list of locales to add.

  Example: fr_FR, ja_JP.UTF-8
- base_config_profile_add_packages: Space-separated list of packages to add.
- base_config_profile_additional_disks: A list of disks to use and configure in addition to the boot disk. Use the format *cntndn*. For each disk listed, define sets of variables for each slice to identify the mount point and the size.
- base_config_profile_cluster: Oracle Solaris software group package.
  - Default Value: SUNWCreq
  - Example: SUNWCreqSUNWCuserSUNWCprogSUNWCallSUNWCXallSUNWCrnet
- base_config_profile_del_clusters: Space-separated list of cluster packages to remove.

  Example: SUNWCpm SUNWCpmx SUNWCdial SUNWCdialx
- base_config_profile_del_geos: Comma-separated list of geographical regions to delete.
- base_config_profile_del_locales : Comma-separated list of locales to delete.
- base_config_profile_del_packages: Space-separated list of packages to remove. To prevent interactive installations on Solaris x86 headless target hosts, set this value to SUNWxwssu SUNWxwscf.

- base_config_profile_dontuse: A comma-separated list of disks that must not be used. Use the format c*n*t*n*d*n*. This variable applies only if base_config_profile_usedisk is not set.
- base_config_profile_root: Root space (free, or size in Megabytes)

  Default Value: free.
- base_config_profile_s3_mtpt: Mount path to slice 3.

  Note: If you are using VxVM and you want your boot disk to look like the mirror, then leave slices 3 and 4 empty.

- base_config_profile_s3_size: Size of slice 3 (in Megabytes).
- base_config_profile_s4_mtpt: Mount path of slice 4.
- base_config_profile_s4_size – Size of slice 4 (in Megabytes).
- base_config_profile_s5_mtpt: Mount path of slice 5.

  Default Value: /var
- base_config_profile_s5_size: Size of slice 5 (in Megabytes).
- base_config_profile_s6_mtpt: Mount path of slice 6.

  Default Value: /usr
- base_config_profile_s6_size: Size of slice 6 (in Megabytes).
- base_config_profile_s7_mtpt: Mount path of slice 7.

  Default Value: /opt

  Note: If you are using Oracle Solaris Volume Manager (SVM), the default behavior is to use slice 7 as a location for metastate databases. If you are using the SVM default configuration, do not use slice 7 for data.

- base_config_profile_s7_size: Size of slice 7 (in Megabytes).
- base_config_profile_swap: Swap space (in Megabytes).

  Default Value: 256
- base_config_profile_usedisk: Defines the boot disk onto which the OS will be loaded. Use the format c*n*t*n*d*n* or the keyword rootdisk. If the value is rootdisk, then the current boot disk is used.

  Default Value: rootdisk
- base_config_shutup_sendmail: If set, create an alias host name to disable sendmail.

  Default Value: yes
- base_config_sysidcfg_default_route: Router IP address to use during JumpStart for Solaris 9 or later environments. If blank, JumpStart uses value from the defaultrouter_base_config variable. If that is also blank, or for another net interface, JumpStart sysidcfg gets a router IP from the JET server.

- base_config_sysidcfg_ip_address: IP address to use at initial boot. By default, this is set to the IP address of the targetable component.

  Default Value: [targetableComponent:ethernet_ip_address]

- base_config_sysidcfg_nameservice: Name service to configure at initial boot.

  Default Value: NIS

- base_config_sysidcfg_netmask: Netmask to use at initial boot. By default, this is set to the netmask of the targetable component.

  Default Value: [targetableComponent:ethernet_netmask]

- base_config_sysidcfg_network_interface: Network interface to use at initial boot.

  Default Value: NONE

- base_config_sysidcfg_protocol_ipv6: Whether to use IPv6 protocol at initial boot.

  Default Value: no

- base_config_sysidcfg_root_password: Encrypted root password.

- base_config_sysidcfg_security_policy: Kerberos security policy to use at initial boot.

  Default Value: NONE

- base_config_sysidcfg_system_locale: System locale to use at initial boot.

  Example: n_US.ISO8859-1

- base_config_sysidcfg_terminal: Terminal emulator to set at initial boot.

  Default Value: vt100

- base_config_sysidcfg_timeserver: Where to get system time for initial boot. If blank, system time comes from the JET server. Alternatively, you can set this variable to localhost to get the system time from the hardware clock on the client.

- base_config_sysidcfg_timezone: System time zone to use for initial boot.

  Example: US/Pacific

- base_config_sysidcfg_x86_kdmfile: For Solaris x86 systems, specifies the name of a keyboard, display, and mouse configuration file to append to the sysidcfg file.

  Default Value: /sysidcfg-addon-file

- base_config_ufs_logging_filesys: For Oracle Solaris 7 and later systems, a space-separated list of mount points to use for logging. To enable logging on all UFS file systems, use the keyword all. Oracle Solaris 9 09/04 enables logging by default. To disable logging on a specific file system, add a hyphen in front of the mount point. To disable logging on all file systems, use the keyword none.

  Default Value: all

  Note: You cannot mix keywords and mount points. You can specify the root file system (/), although the root file system is included as part of the all and none keywords.

- base_config_update_terminal: If set, put the sysidcfg terminal type into inittab.

Default Value: yes

- base_config_x86_confflags: For Oracle Solaris 9 x86 systems, specifies arguments to be used with the confflags attribute of the add_install_client command.

  Example: -f -P /boot/solaris/dca

- base_config_x86_console: For x86 systems, set the console to the correct tty port if you are not going to connect a keyboard and monitor to the client. Setting this variable enables you to perform installs through the serial port. For b1600, v20z, and v40z systems, use ttya. For lx50, v60x, and v65x systems, use ttyb.
- base_config_x86_disable_acpi: For x86 systems, any value disables ACPI. Disabling ACPI might make the installation process proceed better due to how the interrupts are handled.
- base_config_x86_disable_kdmconfig: For Oracle Solaris x86 systems, disables the kdmconfig interactive utility for configuring the keyboard, display, and mouse of the target host. If you are installing an Oracle Solaris OS with the GRUB bootloader, set this variable value to yes.
- base_config_x86_nowin: For x86 systems, prevents Oracle Solaris from trying to run Windows during the install.

  Default Value: yes

- base_config_x86_safetoreboot: For x86 systems, controls whether the system automatically reboots. If your PXE boot is a one-time option, and the next reboot attempts to boot from disk, you must set this option to yes.

Downloading Additional JET Packages

Additional JET packages are available. To download JET packages and the JET user guide, go to http://www.oracle.com/technetwork/systems/jet-toolkit/index.html.

Glossary

account

An account entitles designated cloud users the right to use computing, network, and storage resources of vDC. The account provides the required capabilities to manage these resources. Account defines the amount of vCPU, memory and storage resources that can be used from the available vDC resources.

actions pane

The Actions pane is used to start jobs based on the current selection in the Navigation pane. Selections in the Navigation pane or center pane change the display of operations in the Actions pane. The Actions pane is subdivided into four sections – Operate, Organize, Deploy, and Update.

active

Reflects the state of system and indicates whether monitoring is actually being performed. The active state is not editable. When a rule is not enabled, monitoring is not active. The status is displayed on the Alert Monitoring Rules page, which is accessed from the Monitoring tab. Text in the Active field indicates whether the parameter is active.

activate

Changes an inactive Oracle Solaris boot environment to the new default boot environment on reboot.

Agent Controller

The Agent Controller software communicates with the Enterprise Controller and is installed automatically when an asset is discovered to make the asset a managed asset. You can choose to manage resources remotely with proxy resources without putting an agent on the system. Some features of the product don't work without the agent, but discovery manages the assets without putting an agent on them.

agentless

A system that is managed with Oracle Enterprise Manager Ops Center without the Agent Controller software being installed.

alert monitor

Monitors the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

Alternate Boot Environment

An alternate boot environment, or ABE, is an inactive Oracle Solaris boot environment.

annotations

Annotations are scripts or comments that you can associate with a incident. Annotations can be automated operations to solve a incident, a suggested action, or a comment. You can associate an annotation with a specific incident. Annotations can be added to the Incidents Knowledge Base.

assemblies

Assemblies are kind of infrastructure templates that contain a configuration of multiple virtual machines with their virtual disks and the inter connectivity between them. Assemblies can be created as a set of .ovf (Open Virtualization Format) and .img (disk image) files, or may all be contained in a single .ova (Open Virtualization Format Archive) file.

assets

Assets are physical or virtual piece of hardware, storage device, or operating system that you can manage with Oracle Enterprise Manager Ops Center.

audit log

An audit log file stores details about user log ins, changes to user accounts, and job details. It shows the activity on the Enterprise Controller and the Proxy Controller.

Auto-Balancing Policy

An auto-balancing policy determines if, and how, a server pool is automatically load balanced. By default, automatic balancing is not selected. When you designate the server pool for automatic balancing, the software reviews the load on the virtualization hosts for the interval and day that you request. The software then migrates the guests, as needed, to balance the load. You can require administrator approval before the guests are moved. Also see placement policy and policy.

bandwidth flow

Bandwidth flow is the speed of a connection, or the amount of data that flows from a site's server out to the viewer at any given time.

Baseline

A dated collection of Oracle Solaris patches, patch metadata, and tools. Oracle releases Solaris baselines on a monthly basis. You can use the black lists and the white lists to modify a baseline and create a custom patch set.

baseline check

Baseline check is a feature of Oracle Enterprise Manager Ops Center Engineered Systems where the factory setup of eighth, quarter, half, and full rack configurations are considered as a normal or ideal setup. When the assets are discovered and associated with the rack, this setup is compared with the normal factory setup.

black list

A list of Oracle Solaris operating system patch IDs that you never want to apply to an asset. The black list is used when you are using a baseline to update an Oracle Solaris operating system.

See also **white list**.

block storage

A block storage library consists of LUNs (Logical Unit Number). Each LUN is a slice of a storage volume, which is storage space provided by a collection of disks.

Boolean Control Parameter

A monitoring rule that uses a true-false check.

Boot Environment

A collection of mandatory file systems (disk slices and mount points) that are critical to the operation of the Oracle Solaris operating system. These disk slices can be on the same disk or distributed across multiple disks.

branded zone

Zones that are capable of emulating user environments from operating systems other than Oracle Solaris 10. Zones supports different versions of Oracle Solaris operating system in the zones for running applications.

category

For Oracle Enterprise Manager Ops Center's Local Content, a category is the type of software that is uploaded to Oracle Enterprise Manager Ops Center for use at a site. The parent category is one of the types defined in Oracle Enterprise Manager Ops Center. The local category is a category defined for the site, for example a script for a quarterly inventory.

channel

An operating system distribution, such as Oracle Solaris 10 5/09 on x86 platform or Oracle Linux 5.5. A channel is also called a distribution.

Cloud

A cloud is a set of physical resources that can be divided and allocated to multiple users who can in turn create and use virtual resources as needed without impact to or awareness of the other users' resources. A cloud is implemented as a pool of servers sharing the same virtualization type, storage, networks and fabrics.

cluster heartbeat

Cluster heartbeat is used to verify if the Oracle VM Servers in a clustered server pool are up and running. The heartbeat function has a network component, where a TCP/IP communication channel is created with each Oracle VM Server. Each Oracle VM Server sends regular keep-active packets and these packets are used to determine if each Oracle VM Server is active.

connected mode

This is the default connection mode for Oracle Enterprise Manager Ops Center. With this mode, patch data is regularly downloaded from Knowledge Base through an Internet connection.

Control Domain

A domain that is created when Oracle VM Server for SPARC software is installed. The control domain contains the software packages for Oracle VM Server, including the domains manager application and the domains manager daemon (ldmd) process required for managing the logical domains. The interface to the hypervisor is through the domains manager. The control domain enables you to create, and manage logical domains and allocate virtual resources to the domains.

critical file system

File systems that are required by the Oracle Solaris operating system. When you use Solaris Live Upgrade, these file systems are separate mount points in the vfstab file of

the active and inactive boot environments. Example file systems are `root`, `/usr`, `/var`, and `/opt`. These file systems are copied from the source to the inactive boot environment.

Dashboard

Displays a high-level overview of an asset or a group of assets on the user interface. The information of the selected asset or group is displayed in the Center Pane.

Deployment Plans

Defines the sequence of steps that must be carried out on an asset to deploy. Deployment plans also include the specification or profile that each step should apply, and the resources that are required to apply it such as network addresses, host names and so on. Customized deployment plan enables you to perform hardware, firmware and operating system provisioning activities in a repeatable fashion.

disconnected mode

This is the alternate connection mode for Oracle Enterprise Manager Ops Center. Instead of relying on an Internet connection for updates, patch data is acquired using the harvester script and moved to the Enterprise Controller.

discovery

This is the method for adding assets to Oracle Enterprise Manager Ops Center. Assets can be discovered using a variety of protocols, by their service tags, or by declaring hardware so that it can be configured and provisioned with an operating system.

distribution

For an operating system, a distribution is a specialized version of the operating system.

Domain Name Service (DNS)

DNS is a network protocol that issues IP addresses within a specified range to devices on the network.

Dynamic System Domains

In M-Series servers, you can partition the available hardware resources into smaller logical systems called as dynamic system domains. Dynamic System Domains run their own copies of the operating system and offer a very high level of isolation from other domains in the system because the partitioning occurs at the hardware level.

Dynamic Storage Library

When the block storage library uses LUNs constructed from a storage array that is a managed asset, the block storage library is dynamic. You can add storage capacity as needed by adding LUNs supplied by the storage array.

When the block storage library relies on a storage array that is not a managed asset, the block storage library is static. Because Oracle Enterprise Manager Ops Center has less information about the storage array, you cannot increase the number of LUNs in the storage library.

enabled

A monitoring rule that is enabled is actively monitoring a parameter. By default, all rules are enabled. Users can disable and enable parameters on a per asset or group basis. The status is displayed on the Alert Monitoring Rules page, which is accessed

from the Monitoring tab. Text in the Enabled field indicates whether the parameter is enabled.

Enterprise Controller

This is the central server for Oracle Enterprise Manager Ops Center software. The Enterprise Controller hosts the user interface and communicates with the Knowledge Base. Enterprise Controller stores management information, such as firmware and operating system images, plans, profiles, and policies and also stores the asset data and site customizations. All operations, or jobs, are initiated from the Enterprise Controller.

Enumerated Control Parameter

A monitoring rule that uses a series of values.

Exclusive IP Mode

A dedicated network interface is allocated to the zone. You can choose the network interface when you assign the network to a zone.

Expression Parameter

A monitoring rule that uses an instruction to execute something that returns a value.

/etc Directory

The directory that contains critical system configuration files and maintenance commands.

/etc/netboot Directory

The directory on a WANboot server that contains the client configuration information and security data that are required for a WANboot installation.

/export File System

A file system on an operating system server that is shared with other systems on a network. For example, the /export file system can contain the `root` (/) file system and swap space for diskless clients and the home directories for users on the network. Diskless clients rely on the /export file system on an operating system server to boot and run.

Fabrics

Fabrics are network topologies where network nodes connect with each other through one or more network switches. A true fabric provides a direct connection between any two ports, and supports single step/lookup-based processing. Regardless of its various components, a fabric appears on the outside as a single, logical device with a single, consistent state.

The term is popular in telecommunication, Fibre Channel storage area networks, and other high-speed networks, including InfiniBand.

Filesystem Storage

A software or storage library that relies on a file system on the Enterprise Controller's system or a shared file system on an NFS server that the Enterprise Controller mounts.

global zone

In Oracle Solaris Zones, the global zone is both the default zone for the system and the zone used for system-wide administrative control. The global zone is the only zone from which a non-global zone can be configured, installed, managed, or uninstalled.

Administration of the system infrastructure, such as physical devices, routing, or dynamic reconfiguration (DR), is only possible in the global zone. Appropriately privileged processes running in the global zone can access objects associated with other zones.

group

A group is a user-defined set of assets. Assets can be added to groups based on asset attributes such as type or location. A group can include other groups. Assets can be manually added in addition to the rules based addition using attributes. Any type of asset that can be in a group can be added manually to any user-defined asset group.

guest

Guest refers to a virtual machine that is configured and installed in a virtualization host. For example, the logical domains in an Oracle VM Server host are referred to as guests in a server pool.

Guest Domain

A guest domain is a non-I/O domain that consumes virtual device services that are provided by one or more service domains. A guest domain does not have any physical I/O devices, but only has virtual I/O devices, such as virtual disks and virtual network interfaces.

GUID

Globally Unique Identifier. A pseudo-random 128-bit number that is computed by Windows to identify any component in the computer that requires a unique number. In Oracle Enterprise Manager Ops Center, GUIDs are used to identify LUNs.

Hardware Virtualization (HVM)

Hardware virtualization is a technology that is used to create multiple virtual systems on a single piece of physical hardware. When you create a hardware virtualized (HVM) guest, you must supply an ISO file in a repository to create the virtual machine.

Hardware Virtualized with Paravirtualized Drivers (PVHVM)

PVHVM is identical to HVM, but has additional paravirtualized drivers for improved performance of the virtual machine. PVHVM improves the performance level of Microsoft Windows running in guests.

host name

The name by which a system is known to other systems on a network. This name must be unique among all the systems within a particular domain (usually, this means within any single organization). A host name can be any combination of letters, numbers, and dashes (-), but it cannot begin or end with a dash.

hypervisor

A hypervisor is the software that enables multiple virtual machines to be multiplexed on a single physical machine. The hypervisor code runs at a higher privilege level than the supervisor code of its guest operating systems to manage use of the underlying hardware resources by multiple supervisor kernels.

Image Packaging System (IPS)

Image Packaging System is an Oracle Solaris 11 package that contains operating system components and a manifest that provides basic metadata.

incident

An event that triggers an alert when a monitored attribute does not meet the monitoring parameters. A new incident is displayed in the Unassigned Incidents queue in the Message Center. From the Message Center you can view and act on incidents.

Incident Knowledge Base

A custom database of annotations that are associated with known incidents.

InfiniBand

InfiniBand is a switched fabric communications link primarily used in high-performance computing. Its features include quality of service and failover, and it is designed to be scalable. The InfiniBand architecture specification defines a connection between processor nodes and high performance I/O nodes such as storage devices.

InfiniBand transmission rates begin at 2.5 GBps.

I/O Domain

An I/O domain has direct access to a physical I/O device, such as a network card in a PCI EXPRESS (PCIe) controller. An I/O domain can own a PCIe root complex, or it can own a PCIe slot or on-board PCIe device by using the direct I/O (DIO) feature. An I/O domain can share physical I/O devices with other domains in the form of virtual devices when the I/O domain is also used as a service domain.

IPMP

IPMP (IP network multipathing) provides physical interface failure detection and transparent network access failover. You can configure one or more physical interfaces into an IP multipathing group, or IPMP group. After configuring IPMP, the system automatically monitors the interfaces in the IPMP group for failure.

JET Templates

JumpStart Enterprise Toolkit provides a framework to simplify and extend the JumpStart functionality provided within the Oracle Solaris operating system. You can use JET to install Oracle Solaris on the SPARC and x86/64 platforms. You create JET templates to customize the operating system configuration options as required.

JMX

Java Management Extensions (JMX) technology provides the tools for building distributed, modular, and dynamic solutions for managing and monitoring devices, applications, and networks. The JMX API defines the notion of MBeans, or manageable objects, which expose attributes and operations in a way that enables remote management applications to access them. The public API in Oracle Enterprise Manager Ops Center can be accessed through JMX-Remoting.

Knowledge Base

The Knowledge Base is the repository for metadata about Oracle Solaris and Linux operating system components. Knowledge base stores information about patch dependencies, patch compatibilities, withdrawn patches, downloads, and deployment rules and also stores URL of operating system vendor download sites and downloads the components at set intervals. The Enterprise Controller must have Internet connection to connect to the Knowledge Base.

least allocated

Least allocated is a parameter in the server pool placement policy. The lowest allocated CPU and memory is the total static resource allocation across all guests on the virtualization host. The other placement policy parameter is relative load.

libraries

A collection of virtual machine images and disk images that are located under the same file system. When a server pool is created, one or more libraries are assigned to the server pool. Server pools can share the same libraries.

link aggregation

Link aggregation is a standard defined in IEEE802.3ad. An aggregated link consists of several interfaces on a system configured as a single, logical unit. Link aggregation increases the speed and high availability of a connection between a server and a switch.

LUN

LUN stands for Logical Unit Number. In storage, a LUN is the number assigned to a SCSI protocol entity, that handles (I/O) operations. A SCSI target provides a LUN for each storage volume.

management

An asset is managed when Oracle Enterprise Manager Ops Center can monitor it and target it with jobs. Operating systems can be managed with or without an Agent Controller, but operating system update functions are only available with an Agent Controller.

manifest

Each Oracle Solaris 11 package has an associated manifest that describes how the package is put together. The package manifest provides basic metadata about the package (such as name, description, version, and category), what files and directories are included, and the package dependencies.

maintenance mode

Disables incidents from displaying in the UI, but does not disable monitoring. This mode is useful when you do not want incidents generated during system maintenance.

membership graph

Shows a graphical relationship between assets and status of the connection. A blue line shows the working connection and a red line represents the faulted or disconnected status. The membership graph is displayed in the Center Pane.

message center

Displays all incidents, alerts, and notifications. Message Center helps you to view and manage incidents, notifications, and service request, and display warranty information.

MTU

MTU stands for Maximum Transmission Unit. MTU is the largest packet size, in bytes, that can be sent over a network.

monitoring policy

A set of monitoring rules that defines alert conditions. Policies are either system-defined, user-defined, or generic. Each monitoring policy contains one or more alert monitors for a specific type of resource. An alert is raised when the state is outside the pre-defined condition.

monitoring rule

Contains monitoring parameters that state the values and boundaries for an asset's activity. The set of rules is called a monitoring policy.

MPxIO

MPxIO provides a multipathing solution for storage devices accessible through multiple physical paths. MPxIO is included as a part of the distribution in Solaris 10 onwards.

NAT

NAT stands for Network Address Translation. NAT is a protocol that enables a network to use many internal-only IP addresses and a few Internet-facing IP addresses.

navigation pane

Navigation pane is an important part of the user interface of Oracle Enterprise Manager Ops Center. navigation pane contains Message Center, Assets, Plan Management, Networks, Libraries, Reports, vDC Management, and Administration. The Assets section of the Navigation pane lists all the asset that are managed by Oracle Enterprise Manager Ops Center, grouped by its type and the required criteria.

network

A network enables guests to communicate with each other or with the external world (that is, the Internet). When a server pool is created, one or more networks is assigned to the server pool. Server pools can share the same networks.

network bonding

Network bonding refers to the combination of network interfaces on one host for redundancy and/or increased throughput. Redundancy is the key factor you use to protect your virtualized environment from loss of service due to failure of a single physical link. This network bonding equals as the Linux network bonding. Using network bonding in Oracle VM might require some switch configuration.

network domain

A system of centralized network administration, in which the permissions that grant access to resources in the network are maintained in one or more servers. Network Domains use a hierarchical structure that enables you to assign permissions to collaborate with different departments in an organization.

A large network may have several domains based on the needs of each set of users.

NIS

NIS stands for Network Information System. NIS is a network naming and administration system for smaller networks. NIS is similar to the Internet's domain name system (DNS) but designed for a smaller network.

non-global zone

A virtualized operating system environment created within a single instance of the Oracle Solaris operating system. One or more applications can run in a non-global zone without interacting with the rest of the system. Non-global zones are also called zones.

non-sparse copy

A clone of the type "non-sparse copy" is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

notifications

An email, pager, or user interface message that is automatically sent by Oracle Enterprise Manager Ops Center when specified conditions are met. You can configure separate notification profiles for different assets and different users. You can configure the software to send notification for specific incidents, or when a critical or warning incident is detected.

Opaque Data

An opaque data is a data type that is incompletely defined in an interface, so that its values can only be manipulated by calling subroutines that have access to the missing information.

/opt

A file system that contains the mount points for third-party and unbundled software.

Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control is a single, integrated solution for managing all aspects of the Oracle Cloud and the applications running on it. Oracle Enterprise Manager Cloud Control couples a potent, top-down monitoring approach to delivering the highest quality of service for applications with a cost-effective automated configuration management, provisioning, and administration solution.

Oracle Engineered System

Oracle Engineered Systems are hardware and software integrated systems that are designed for a specific enterprise purpose. Oracle Engineered System helps in reducing the cost and complexity of the IT infrastructures, and increases the productivity and performance.

Oracle Services

Provides integrated methods of maintaining and displaying current contracts, warranty information, contract dates, and service requests in Oracle Enterprise Manager Ops Center.

Oracle Solaris Clusters

Oracle Solaris Clusters is a high availability software product for Solaris operating system. Oracle Solaris Clusters are used to improve the availability of software services such as databases, file sharing on a network, electronic commerce websites, or other applications. You can now manage Oracle Solaris Clusters as any other asset using Oracle Enterprise Manager Ops Center.

Oracle Solaris Zones

Oracle Solaris Zones is a software partitioning technology used to virtualize operating system services, and provide an isolated and secure environment for running

applications. When you create a non-global zone, you produce an application execution environment in which processes are isolated from all other zones. This isolation prevents processes that run in a zone from monitoring or affecting processes that run in any other zones. See also global zone and non-global zone.

Oracle Solaris 11 Software Update Library

Oracle Solaris 11 Software Update Library repository is located on the Enterprise Controller. This contains the Oracle Solaris 11 packages that you need to install, provision, and update your Oracle Solaris 11 operating system.

Oracle VM Server for SPARC

Oracle VM Server is a virtualization technology that enables the creation of multiple virtual systems by a hypervisor in the firmware layer, interposed between the operating system and the hardware platform. This is designed to abstract the hardware and can expose or hide various system resources, allowing for the creation of resource partitions that can operate as discrete systems, complete with virtual CPU, memory and I/O devices.

Oracle VM Server for SPARC was previously known as Logical Domains, it is a virtualization technology designed to run on CMT based servers.

Oracle VM Server for x86

Oracle VM Server for x86 is a managed virtualization environment or part of such an environment, that is designed to provide a lightweight, secure, server-based platform for running virtual machines. Oracle VM Server for x86 is based upon an updated version of the underlying Xen hypervisor technology, and includes Oracle VM Agent.

Oracle Solaris ZFS

An Oracle Solaris operating system file system that uses storage pools to manage physical storage.

OS Provisioning Profile

Defines the image, provisioning, and installation requirements.

OS Configuration Profile

Defines the OS and network configuration.

Paravirtualization

Paravirtualization enables you to select a location for the mounted ISO file from which you create the virtual machine. Before you create the virtual machine using the paravirtualized method, you must mount the ISO file on an NFS share, or HTTP or FTP server.

parent repositories

Any hosted Oracle repository that Oracle Solaris 11 Software Update Library can use to upload, or sync, content.

photorealistic view

Photorealistic view displays the front and rear views of the rack. All slots and the respective assets are displayed. Positions within the rack are displayed in a 2-dimensional view. All assets in the rack have a specific image. The health status of assets such as OK, Warning, and Critical are displayed in the form of colored buttons.

placement policy

Determines whether the guest is placed on a virtualization host with the lowest relative load or the least allocated. By default, new guests are placed on the server with the lowest load and are automatically started. The placement policy is defined when a server pool is created. Server pools can have different placement policies.

policy

Defines how a job is performed and sets the automation level of the job. A policy file is similar to a response file. If there is a conflict between a profile and policy, the profile overrides the policy.

Private vNet

vNet that is unique to a given account is called Private vNet.

profile

Defines the configuration of components for a specific type of system. By using a profile, you can define what is enabled, and not enabled, to be installed on a system. If there is a conflict between a profile and policy, the profile overrides the policy.

Proxy Controller

Proxy Controllers link the managed assets to the Enterprise Controller and act as proxies for operations that must be located close to the managed assets, such as operating system provisioning. Proxy Controllers distribute the network load and provide for fan-out capabilities to minimize network load. Proxy Controllers perform management operations on assets and report the results to the Enterprise Controller. An Oracle Enterprise Manager Ops Center installation must have at least one functioning Proxy Controller.

relative load

Relative load is a parameter in the server pool placement policy. Lowest relative load is based on the lowest memory and CPU utilization for the virtualization host over the past three weeks. The other placement policy parameter is least allocated.

repository

A repository is a central place that stores an aggregation of data in an organized way, usually in a computer storage. Depending on how the term is used, a repository may be directly accessible to users or may be a place from which specific databases, files, or documents are obtained for further relocation or distribution in a network.

root

The top level of a hierarchy of items. `root` is the one item from which all other items are descended. See `root` directory or `root` (`/`) file system.

root directory

The top-level directory from which all other directories stem.

Root Domain

A root domain has a PCIe root complex assigned to it. This domain owns the PCIe fabric and provides all fabric-related services, such as fabric error handling. A root domain is also an I/O domain, as it owns and has direct access to physical I/O devices.

root file system

The top-level file system from which all other file systems stem. The `root ( / )` file system is the base on which all other file systems are mounted, and is never unmounted. The `root ( / )` file system contains the directories and files critical for system operation, such as the kernel, device drivers, and the programs that are used to boot a system.

RPM

A package manager used by many versions of the Linux operating system.

rule parameters

Define the monitoring parameters. The following types of rule parameters are available: Threshold, Boolean Control, Enumerated Control, and Expression. Some parameters are editable. All active parameters can be disabled.

SAN Storage Library

Storage Attached Network (SAN) storage which is used for providing storage spaces for managed assets in Oracle Enterprise Manager Ops Center. The SAN storage library consists of groups of LUNs.

script

A command file that is associated with one of Oracle Enterprise Manager Ops Center's actions, either before the action occurs (pre-action script), or after the action completes (post-action script).

security group

The organization of users and other domain objects into groups for easy administration of access permissions is known as a security group. A Security Group enables you to specify certain security settings on an instance specific basis. You have the ability to filter traffic based on IP's (a specific address or a subnet), packet types (TCP, UDP or ICMP), and ports (or a range of ports). You can also grant access to an entire security group so that your trusted computers can get access to each other without having to open ports to the public.

server management

Server management is used to manage the physical Oracle VM Servers in a server pool, for example, to update the Oracle VM Agent on the different Oracle VM Servers.

server pool

A server pool is a resource pool of virtualization hosts that share compatible chip architecture, which facilitates actions such as moving guests between virtualization host instances. Members of the server pool have access to the same network and storage library resources. Guests can access the images contained in the server pool's library. Several server pools can share the same network and library storage resources.

server templates

Server templates provide pre-built images for creating vServers. They can be uploaded individually or as part of an Assembly. Server templates can be created from an existing vServer.

service tag

Service tags are XML files that identify assets uniquely. Assets with service tags can be discovered using the Find Assets wizard.

Service Domain

A service domain provides virtual device services to other domains, such as a virtual switch, a virtual console concentrator, and a virtual disk server. You can have more than one service domain, and any domain can be configured as a service domain.

Shared IP Mode

The global zone shares its network interface with one or more zone. You must define the network interface when you assign the network to the global zone.

shared storage

A shared storage library in Oracle Enterprise Manager Ops Center is one that is accessible by the server and operating system. It is not related to Zones on Shared Storage in Oracle Solaris 11.1.

snapshot

Snapshot, a point in time image of a volume is a non-bootable copy of a boot environment that uses much less disk space than a boot environment. You can create a boot environment from a snapshot.

software libraries

A software library can be a local file system on the Enterprise Controller or a mount point on an NFS server. The software library is used to store the operating system images for provisioning, branded images, flars, firmwares, profiles, operating system updates, custom programs and scripts.

sparse copy

A clone of the type "sparse copy" is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size.

static route

Specifies the route taken by the network for external access. You define a default gateway for the network; however, this default gateway may not be reachable to a given subnet. In this case, you must add a static route for this specific subnet.

status pane

The Status pane in the Jobs section describes about the state of the incidents like jobs in progress, jobs failed, jobs partially successful, jobs stopped, jobs schedules, jobs successful and so on.

Support Repository Update (SRU)

Support Repository Update (SRU) is a package of Oracle Solaris 11 operating system updates that releases on a regular basis.

SCCM

Microsoft System Center Configuration Manager (SCCM), is used to update Windows operating systems.

syncing

Syncing is the process of reconfiguring or updating the Oracle Solaris 11 Software Update Library with the Oracle Solaris 11 Image Packaging System (IPS).

synchronizing

Updates an inactive boot environment to match an active boot environment.

system groups

Default asset groups that automatically organize your assets by type in the user interface.

System-defined Rules

Attribute specific monitoring rules that are hard-coded into drivers. You can disable a system-defined rule, but cannot edit, move, or reconfigure these types of rules.

Thin Clone

A thin clone is a clone of a physical disk that takes up only the amount of disk space actually in use; not the full specified disk size.

threshold parameters

A monitoring rule that uses a numeric value above or below a defined level.

time server

The network device that provides accurate time for synchronizing network activity.

unmanaged storage

Unmanaged storage is the storage resource that is unknown to Oracle Enterprise Manager Ops Center. When you add storage to zones using the native CLI or manage existing zone environments, the zone's storage is not identified and termed as unmanaged.

User-defined Network Domain

A network domain provides custom network resources from an Ethernet or InfiniBand fabric to virtualization hosts, server pools, or virtual datacenters so that new networks can be created as needed. A user-defined network domain supplements the Default Network Domain that is always available and cannot be deleted.

User-defined Rules

Monitoring rules that are associated with, and determined by, the type of managed resource. You can apply a user-defined rule to many different attributes.

/usr File System

A file system on a standalone system or server that contains many of the standard UNIX programs.

Sharing the large `/usr` file system with a server rather than maintaining a local copy minimizes the overall disk space that is required to install and run the Solaris software on a system.

/var File System

A file system or directory (on standalone systems) that contains system files that are likely to change or grow over the life of the system. These files include system logs, vi files, mail files, and UUCP files.

vDC

vDC is a collection of physical servers and storage that are placed on a common network. These physical resources are organized into a pool that are accessed by self-service users. This offers an access point through which you can allocate and control the resources inside. This is created during the set up phase.

vNets

vNets are managed networks and their associated logical (L2) fabrics that can be associated with a vDC and its Accounts.

vServer

vServer is an entity that provides the outward interface of a standalone operating system. This may be a Virtual Machine (VM) or a Solaris Container or a similar construct. This consumes CPU and memory resources. This can be a member of one or multiple vNets.

vServer Type

vServer type is a profile for vServer creation that defines size of memory, size of disk and number of vCPUs to be used when creating a new vServer instance, that is used in combination with a Server Template.

VID

VLAN Identifier. Part of the VLAN tag inserted into Ethernet frame that specifies its VLAN.

virtual disk image

A virtual disk image is a representation of a virtual storage device that is associated with a virtual machine. Such storage can represent a virtual hard disk or a virtual CD/DVD.

virtualization host

Oracle VM Server that are managed by Oracle Enterprise Manager Ops Center is referred to as virtualization host. The virtualization host contains a hypervisor and its local resources and network connections.

virtual machine

A virtual machine is a software implementation of a computing environment in which an operating system or program is installed and run.

A virtual machine typically emulates a physical computing environment, requests for CPU, memory, hard disk, network, and other hardware resources that are managed by a virtualization layer which translates these requests to the underlying physical hardware.

virtual machine template

A Virtual Machine Template provides a standardized group of hardware, and software settings that is used repeatedly to create virtual machines configured with those settings.

virtual server image

A virtual server image is the persisted specification and state of a virtual machine. A virtual server is created when you create a guest. The virtual server image contains the general specification of the guest such as CPU, network, memory, and the type of physical storage that is backing the guest. A virtual server image is also referred to as a guest image.

Virtual Local Area Network (VLAN)

VLAN is a group of network resources connected to different network segments that behave as if they were connected to a single network segment. All transmissions from the VLAN are identified by a unique VLAN tag.

volume

A volume is an identifiable unit of data storage that is sometimes physically removable from the computer or storage system. In tape storage systems, a volume may be a tape cartridge. In mainframe storage systems, a volume may be a removable hard disk. Each volume has a system-unique name or number that enables it to be specified by a user.

white list

A list of Oracle Solaris operating system patch IDs that you always want to be applied to a host. The white list is used when you are using a baseline to update an Oracle Solaris operating system.

See also **black list**.

WINS

WINS stands for Windows Internet Naming Service. The WINS server converts NetBIOS names to IP addresses.

WS-Man

Web Services for Management (WS-MAN) is a specification for managing servers, devices, and applications using web services standards. WS - Man provides a common way for systems to access and exchange management information across the entire IT infrastructure. The public API in Oracle Enterprise Manager Ops Center can be accessed through WS-Management.

World Wide Name (WWN)

WWN is a unique identifier in a Fibre Channel or Serial Attached SCSI storage network. Each WWN is an 8-byte number derived from an IEEE OUI and vendor information.

zone

Also called non-global zones, are a virtualized operating system environment created within a single instance of the Oracle Solaris operating system. One or more applications can run in a non-global zone without interacting with the rest of the system.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Y

Z