Oracle® Enterprise Manager Ops Center

User's Guide

11*g* Release 1 Update 3 (11.1.3.0.0)

E18415-04

November 2011

ORACLE®

Oracle Enterprise Manager Ops Center User's Guide 11*g* Release 1 Update 3 (11.1.3.0.0)

E18415-04

Contents

3 Message Center

4 Oracle Services

5 Job Management

6 Hardware Monitoring

7 OS Monitoring

8 Reports

9 Problem Management

10 Virtual Pools

11 Oracle VM Server for SPARC

12 Oracle Solaris Zones

13 Cluster Monitoring

14 Plans

15 Enterprise Manager Ops Center and Grid Control

Preface

The *Oracle Enterprise Manager Ops Center User's Guide* describes the concepts and procedures for the routine use of the product software to maintain the data center's infractructure.

The descriptions in this document assume that the product software has been installed and configured completely, and that product components such as libraries, profiles, and deployment plans have been established.

Audience

This document is intended for system administrators who are reponsible for monitoring the status of systems in a data center and for managing any provisioning that is required.

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 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 documents in the Oracle Enterprise Manager Ops Center 11g Release 1.0 documentation set:

Related Documents

For more information, see the following documents in the Oracle Enterprise Manager Ops Center 11*g* documentation set:

- *Oracle Enterprise Manager Ops Center Concepts Guide*
- *Oracle Enterprise Manager Ops Center Advanced User's Guide*
- *Oracle Enterprise Manager Ops Center Provision and Update Guide*
- *Oracle Enterprise Manager Ops Center Administration Guide*

- *Oracle Enterprise Manager Ops Center Reference Guide*

- *Oracle Enterprise Manager Ops Center Installation Guide for Oracle Solaris Operating System*

- *Oracle Enterprise Manager Ops Center Installation Guide for Linux Operating Systems*

- *Oracle Enterprise Manager Ops Center Site Preparation Guide*

- *Oracle Enterprise Manager System Monitoring Plug-In for Oracle Enterprise Manager Ops Center Guide*

- *Oracle Enterprise Manager Ops Center Release Notes*

- *Oracle Enterprise Manager Ops Center Readme*

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.
`monospace`	Monospace type indicates commands, file names, and directories within a paragraph, and code in examples.

1

Using the Browser User Interface

Enterprise Manager Ops Center is an integrated solution for managing all aspects of the data center. From a single console, you can discover, provision, patch, manage, and monitor the physical and virtual assets in one or more of your data centers. The operations you perform cover a wide range of Enterprise Manager Ops Center's capabilities:

- You bring assets into Enterprise Manager Ops Center's management , using several types of discovery, and organize them.

- You monitor all the assets, both physical and virtual.

- You communicate status for all problems and take preventative or corrective actions.

- You create reports of assets' performance and utilization.

To perform these tasks, you need a comprehensive knowledge of the types of assets and product features. This chapter describes the product's browser interface:

Overview of the User Interface

The user interface has a rich set of features that are designed to let you observe and control your data center's assets. These features present the range of information from a high-level overview of your data center to the low-level details of a specific asset. This information is presented in the center pane of the UI and is the result of your selections in the Navigation pane and in the Actions pane.

Figure 1–1   Enterprise Manager Ops Center User Interface.
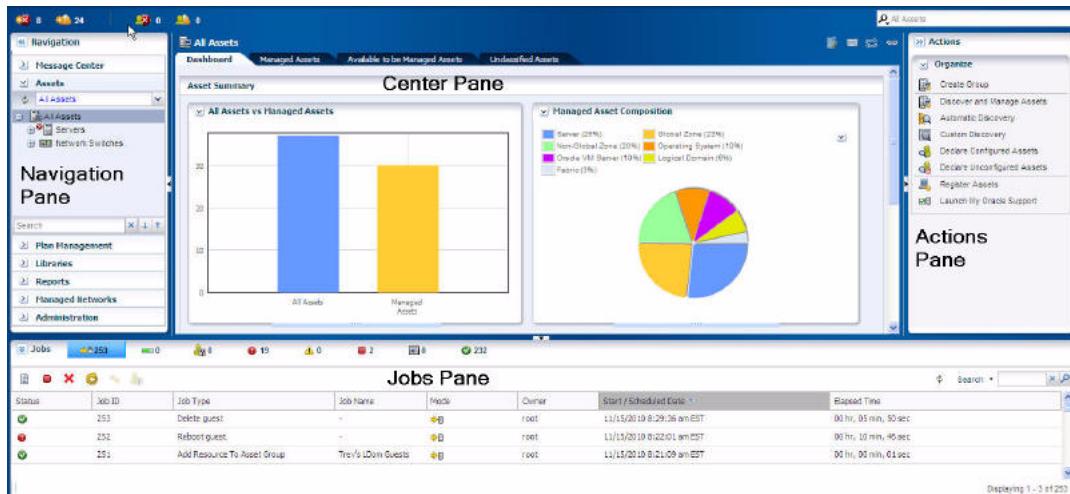
Navigation Pane

The Navigation pane gives access to the following sections:

- Assets – List of all assets that are registered in Enterprise Manager Ops Center.

- Managed Networks – Contains details about the virtual networks defined for Enterprise Manager Ops Center.

- Libraries – Stores all images, policies, and profiles.

- Reports – List of reports that you can create, such as update and firmware reports.

- Administration – Contains administrative functions, such as user administration, logs, and the status of service, including the version and upgrades available for the Agent.

Click on the right-arrow next to the section title to open the section and see its available resources and options. You can only open one section at a time.

Figure 1–2   Opening a section of the Navigation pane
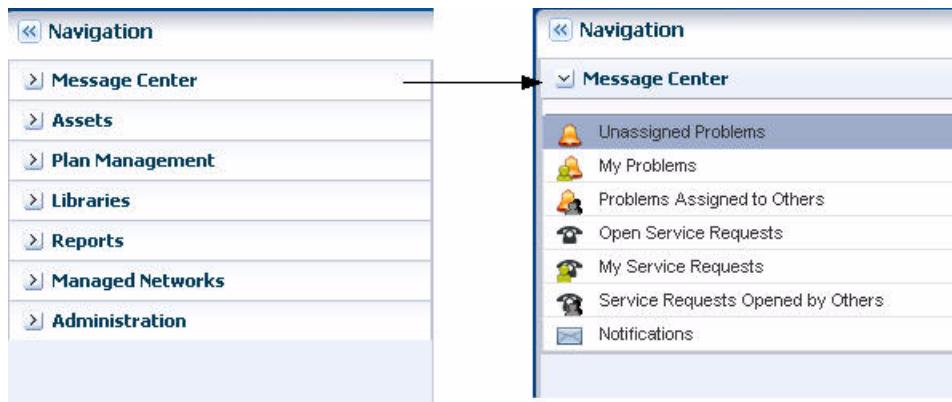
The Navigation pane's Asset section lists every asset that is managed by Enterprise Manager Ops Center, grouped by its type and by other criteria that you provide. In

even the smallest data center, refine the list of assets to show only the assets of interest. Choose one of the views in these categories:

- Standard Views to organize the assets by type.

- Resource Management View to list assets that share resources and are treated as nodes.

- User-Defined Groups to organize assets according to your site's definitions. When a group is created, this customized list becomes the new default view. A list can show hierarchy of assets so when you see a + symbol in the list, click on the symbol to display the subordinate or individual assets.

If you know that Enterprise Manager Ops Center's information is changing and you do not see an asset, click the Refresh symbol in the Navigation pane. For example, if Enterprise Manager Ops Center discovers a new asset, the Asset list in the Navigation pane might not be current.

Each asset and its status is represented by an icon. These icons are an aid to help you distinguish one type of asset from another type and to see the condition of each. A quick scan of the asset hierarchy shows you which servers are running a Linux OS and which ones are running an Oracle Solaris OS or finds any unconfigured assets.

The Navigation pane shows not only all the assets, but all the other elements managed by Enterprise Manager Ops Center: plans and profiles, networks, libraries of images and data, reports, and administrative functions. Click on the title bar's right arrow to display a hierarchy of these items.

Your selections in the Navigation pane change the display in the center pane. The center pane is where you do most of your work. To increase the size of the center pane, close the Navigation pane. In the Navigation pane's title bar, click the left double-arrow to collapse the pane. When you want to select another asset, expand the pane again by clicking on the right double-arrow.

Center Pane

The initial display, in which no asset is selected, gives an overview of the assets managed by Enterprise Manager Ops Center. Each subpane represents the types of assets in a different type of chart. To identify the asset or asset type, hover the mouse of that portion of the chart.

Figure 1–3   Initial View in Center Pane
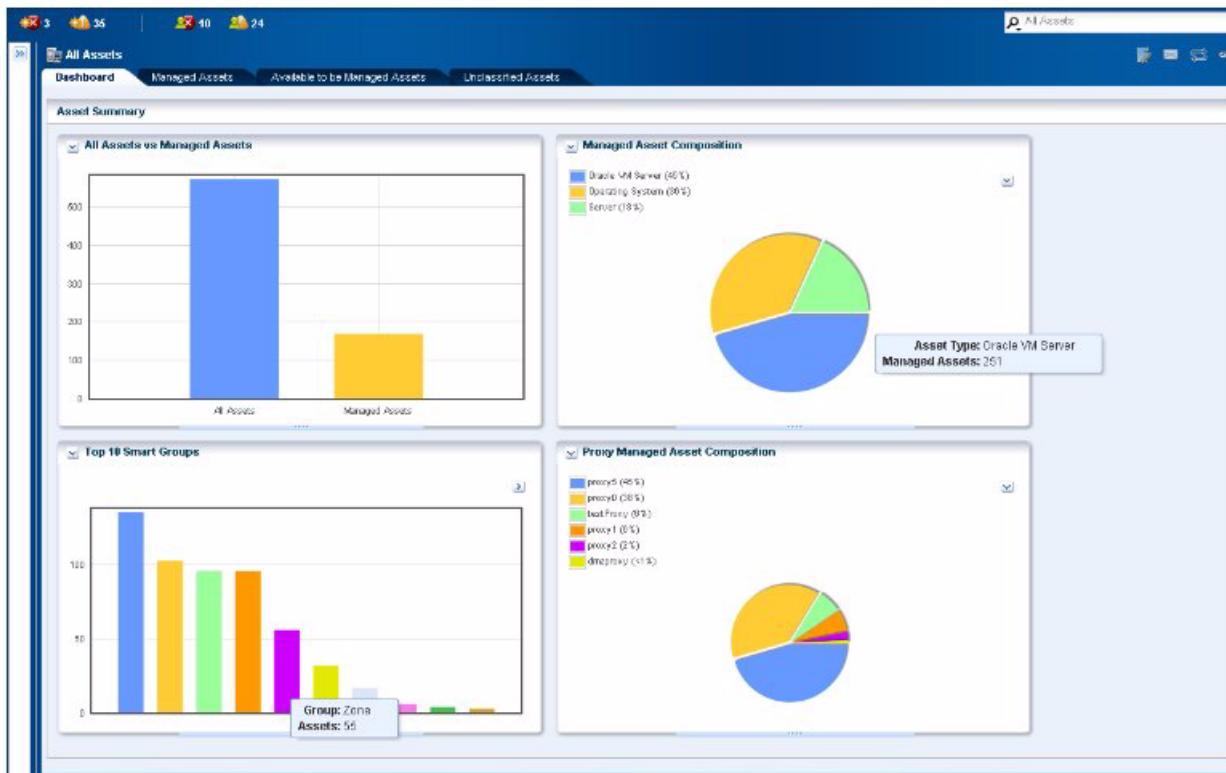
Any selection in the Navigation pane changes the display in the center pane. Depending on the complexity of your selection, allow time for the information, especially in the Membership Graph , to finish loading.

For most types of assets, the center pane shows a dashboard of standard information about the selected asset, a membership graph that shows its relationships, and a status summary. Use the scroll bars to bring all the information into view. For some types of assets, you need to click on the right double-arrow button to see all the tabs.

Figure 1–4 Center Pane for a server
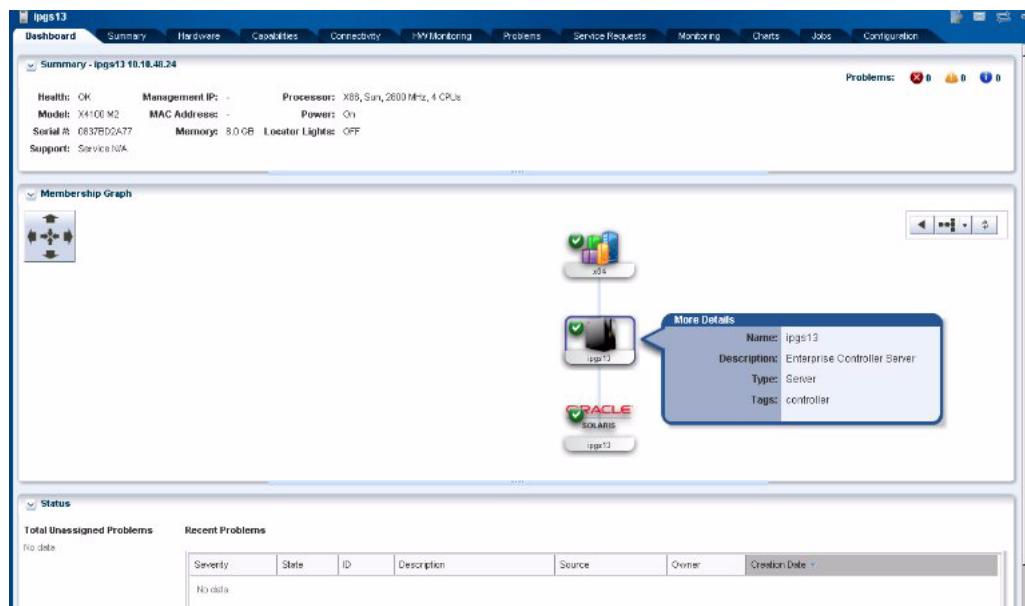
Some types of assets have additional sub-panes such as Services and File Systems. If any of this information is not of interest, you can collapse these sub-panes to hide the display and to prevent the information from refreshing. Click on the down-arrow in the corner of the sub-pane to hide the sub-pane. Click on the resulting right-arrow to reveal the sub-pane.

At any time, when you know that information is changing and you do not see the new information, click on the Refresh symbol in the upper-right corner of the center pane.

Membership Graph

The membership graph is the default display in the center pane and shows the relationships of the selected asset. Depending on what you select in the Navigation pane and on your data center's organization, the membership graph can have a single object or many objects. To control the view, you have these options:

- Use the scroll bars to move the display, bringing assets in large graphs into view

- Use view controls to move the membership graph left and right, and up and down. These view controls are two sets of arrows: four inward arrows and four outward arrows. Click on one of the outward arrows to move the membership graph. Click on one of the inward arrows to center the graph in the center pane.

- Change the orientation of the graph. The hierarchy can be presented from the top down or from the bottom up. If the hierarchy is flat and wide, a more convenient display might be to change to a horizontal display so that the top is now on the left or right. Click the left arrow to restore the display.

- Drill into the graph. If you choose assets or a group at the top of a hierarchy, the membership graph consolidates the display of the assets so that the graph is not unwieldy. Click in the graph to show the actual members.

Figure 1–5   Changing Orientation of Membership Graph
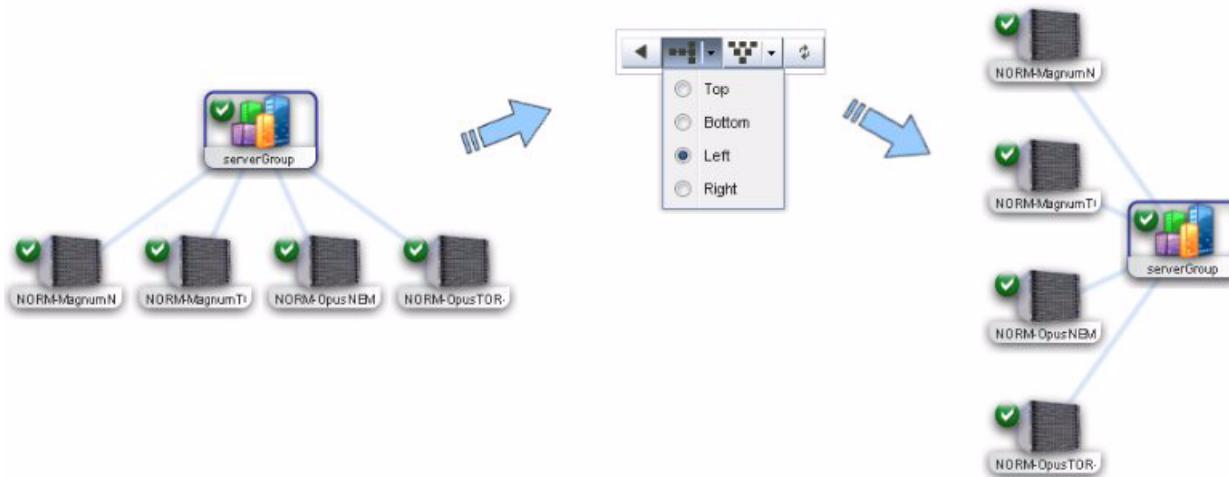
In addition to showing the relationship among members, the lines that connect two members show the status of the relationship. A blue line is working connection and a red line is faulted or disconnection.

Actions and Jobs

You use the Action pane to manipulate assets and to issue commands. Your selections in the Navigation pane or center pane change the display in the Actions pane. Depending on the asset you select, the set of available actions change. An action is not available for these reasons:

- You have not selected an asset that can perform that action. For example, you intend to perform an action on a server's operating system so you select the server. The action is not available until you select the operating system asset that is listed under the server.

- Your user account's role does not allow the action.

At any time, if you are not issuing commands, you can increase the size of the center pane by hiding the text on the Actions pane. In the Action pane's title bar, click the right double-arrow to collapse the pane. The icons for the available actions are displayed but no text. To select another action, click on its icon or expand the pane by clicking on the left double-arrow button to be able to select the names of the actions.

Every action creates a job. Many jobs can run at the same time and some jobs take longer than others to complete. The Job pane is at the bottom of the user interface and its default display is a count of all jobs, according to each one's status. To see the status type, hold or hover the mouse over the job icon. Some jobs have many steps and to follow the progress of a particular job, display the job details. Click on the Jobs title bar to expand the pane, click on the job, and then select the View Job Details icon.

Figure 1–6   Jobs Pane

Some actions creates the job immediately but, in many cases, you need to provide more information to specify the job. These actions launch a wizard to guide you through constructing the job. The last step of every wizard is a Summary page so that you can review your specifications before you create the job.

Complex operations or operations that you want to repeat many times are accomplished through plans. The actions are defined in the profiles that make up the steps of the plan. The Apply Deployment Plan wizard prompts you to identify an existing plan and then select the targets of the plan. Targets are the assets that you want the plan to act on. In the wizard, you select assets from the list of appropriate assets and click the Add to Target List button to include these assets in the Target list. You can select multiple assets or select them one at a time.

Figure 1–7   Wizard with Target Selection
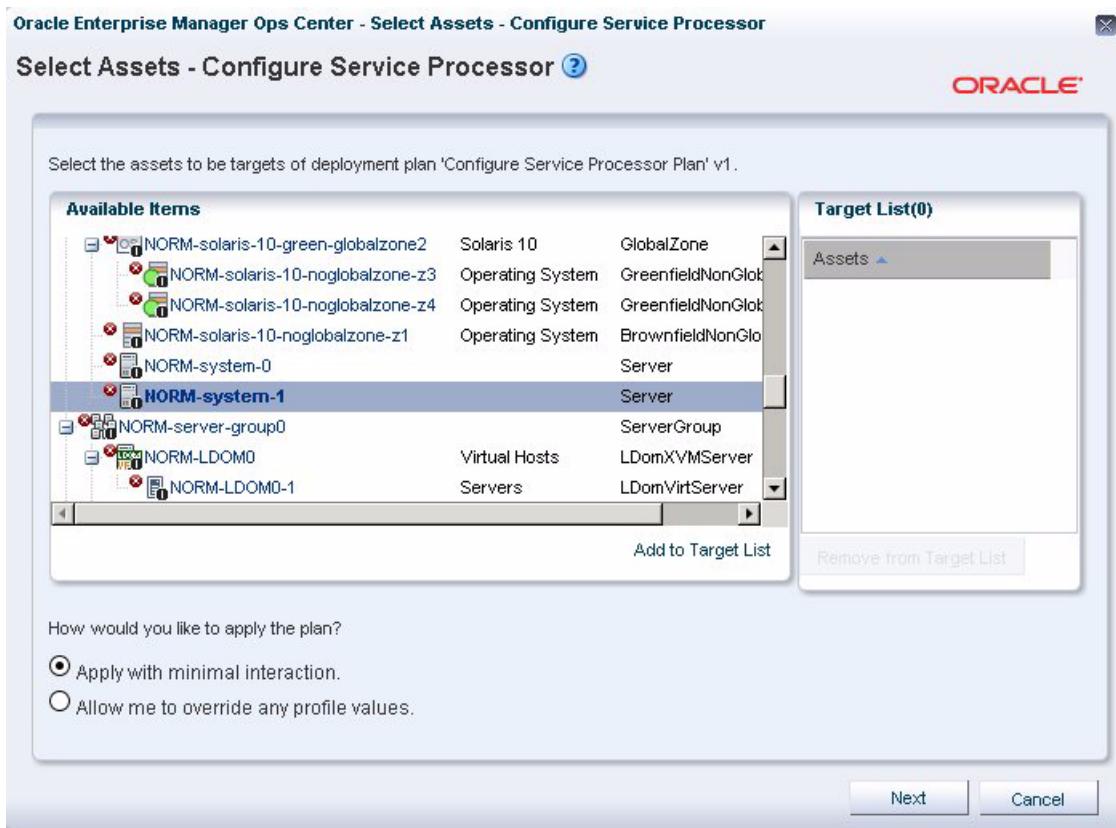
Tabs and Dashboards

The default display in the center pane is the Dashboard or a Summary, depending on the asset type.

Dashboards are a high-level overview of an individual asset or a group of assets. The information displayed is determined by the asset or group. Every group and managed asset has a dashboard that provides a summary, a membership graph, and the monitoring status. Additional information might be displayed, depending on the type of group or asset.

The Summary tab provides information on all members of a group and charts to identify the most active and least active members. To learn more information about the asset or details about information in the default display, click on one of the tabs. These tabs vary, depending on the asset type, and are explained in other sections of this guide. If there are too many tabs to fit into the center pane, click on the right double-arrrow button to bring the remaining tabs into view. Hold or hover the mouse over the value in the table to see its definition.

Charts are included in the Dashboard, Summary, and the Charts tab. These charts give you access to the information that Enterprise Manager Ops Center is collecting about all the assets all the time. Hold or hover the mouse over a portion of the chart to see the value it represents. On the Dashboard, you can hide the legend for a chart and you can hide the chart itself. You hide the legend if it is taking too much room and you want to see more detail in the chart. You hide the entire chart if it is not of interest at the time and you want to focus on another part of the center pane. Click on the down-arrow button to hide and click the right-arrow button to show. On the Charts tab, several charts are displayed. When you right-click on one of the charts, a submenu of the following actions is displayed so that you can change the chart:

- Select to pick a point on the line and display its values.

- Zoom to concentrate on a portion of the chart.

- Move to change the location of the chart.

- 100% to restore the original chart.

- Select All

- Property to display a table of all the values in the chart You can control the scope of the chart by increasing or decreasing the time period of the data collection.

Group Dashboards

For a group of assets, the Dashboard tab displays a summary of the members of the group, how the members relate to each other and other managed assets in Enterprise Manager Ops Center, and a summary of known problems and their severity levels. You can drill down for more group details by clicking the Membership tab and the Problems tab.

A group dashboard contains the following information:

- Group summary, including the type of group, the location, description, number of members, and a high-level summary of unassigned issues.

- Membership graph, which provides a graphical representation of the group members and their relationships, and the monitoring status of each member.

- Status, including a pie chart representation of all unassigned problems for the group, and a table of recent problems.

■ Asset Summary contents are determined by the asset type. The asset summary usually includes the group composition and the largest consumers of resources. For example, the summary might include the top five users of CPU, memory, network, and power.

Figure 1–8 shows the summary and membership graph sections of a dashboard. The summary includes group details and known problems. In this case, two warnings and three informational problems were detected for the group. The Membership Graph uses the problem status icons to indicate the member status. The Status section displays the unassigned problems, by severity, as a percentage of the total number of unassigned problems. It includes a pie chart representation of all unassigned problems for the group, and a table of recent problems. The information in this section includes the group composition and a graphical representation of the resources consumed by the top five assets. The resources displayed in these tables is determined by the type of asset. In this figure, the summary displays the top five users of CPU, memory, and network and the percentage used.

Figure 1–8   Asset Summary
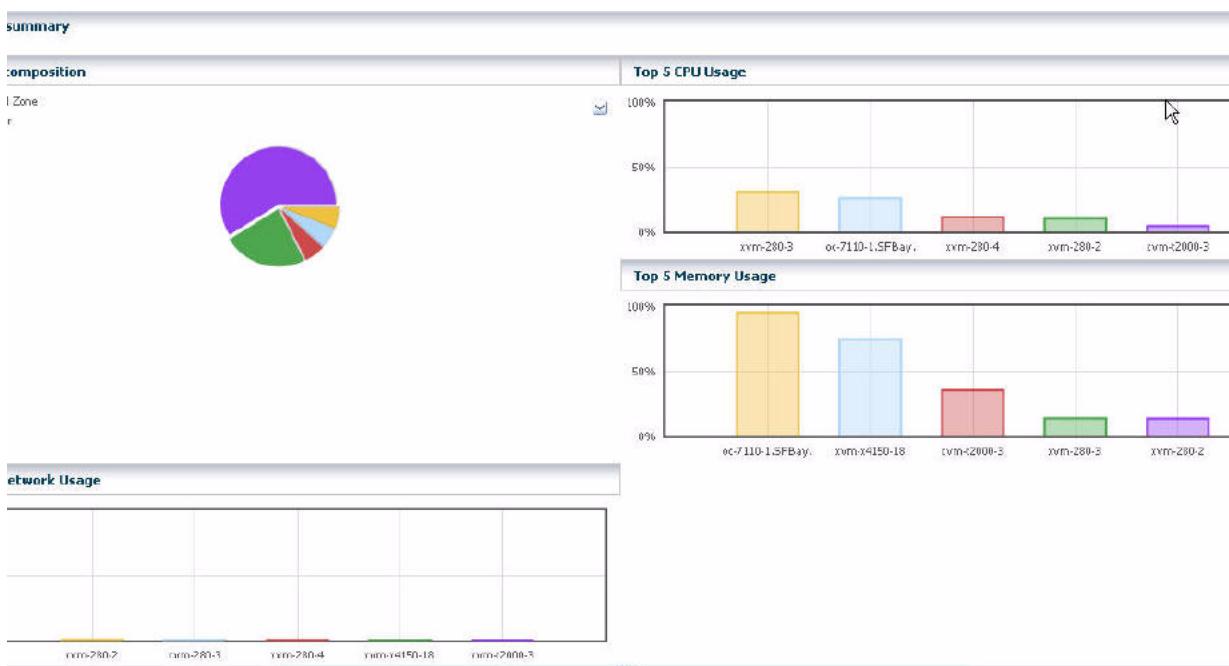
To view a higher level of detail about the group members, click the Membership tab. To view a higher level of detail about problems, click the Problems tab.

Asset Dashboards

For an asset, the Dashboard tab displays a summary of the managed asset, the relationships for the asset, and a summary of known problems and problem severity levels. You can drill down for more details by clicking the tabs in the center pane.

An asset dashboard contains the following information:

■ Asset summary, including the type of group, the location, description, and number of members

■ Membership graph, which provides a graphical representation of the relationships for the asset

- Status, including a pie chart representation of all unassigned problems and a table of recent problems.

- Compliance reports, which display the latest firmware and OS compliance reports that have been run on the asset.

Figure 1–9 shows the summary and membership graph sections of a dashboard for an operating system. The summary includes OS details and a summary of the unassigned problems. Some tabs are asset-specific and others are standard across asset types. The Dashboard, Summary, Problems, Monitoring, and Jobs, are standard.

Figure 1–9   Summary and Membership Graph of an OS Asset
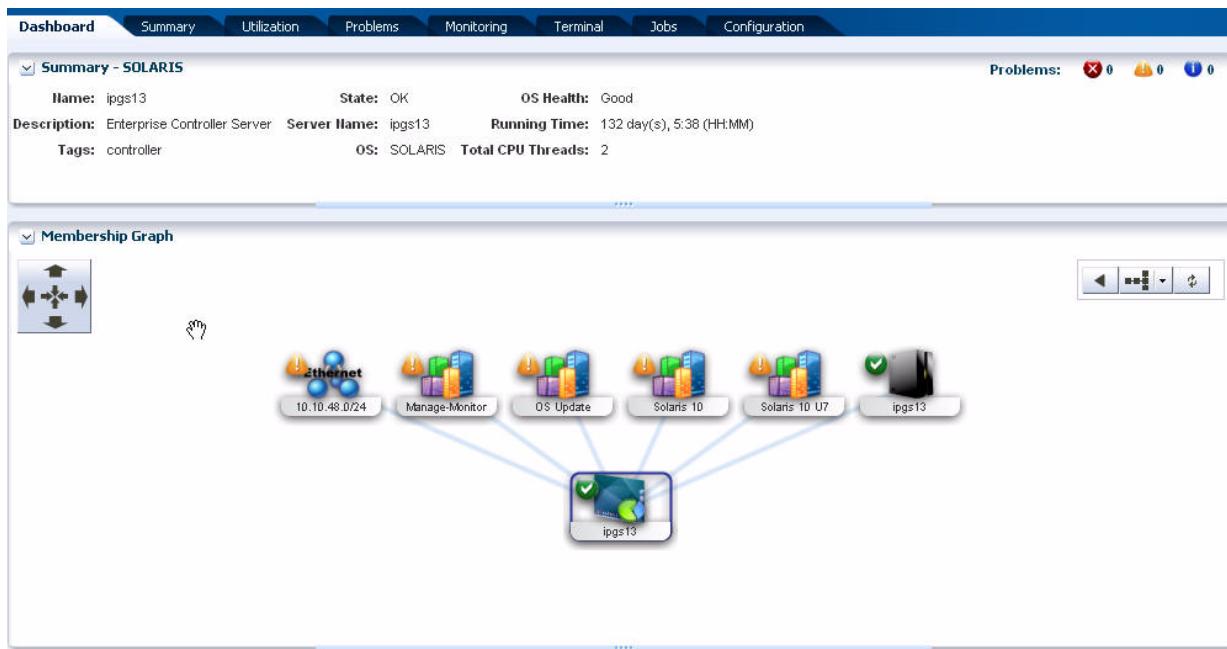
You can drill down for more details by clicking one of the tabs in the center pane.

Viewing a Dashboard

1. Click Assets in the Navigation pane.

2. If you have user defined groups, the filter defaults to All User Defined Groups. To change the view, select an option, such as All Assets, from the drop-down menu.

3. Select a group or an individual asset. The Dashboard displays in the center pane.

4. (Optional) Click a tab to view greater detail. The Summary tab provides a high-level view of the asset's attributes. More details are available on the subsequent tabs.

Status of Your Session

Some icons are always displayed to show you the current status. In the title bar, icons report on the session. Some icons such as Help, current account, and Logout are familiar to users of any application.

Figure 1–10   Icons in the Title Bar

The others show you the mode of the session:

- Connected to or disconnected from Internet access

- Connected to or disconnected from Knowledge Base Service

- Connected to or disconnected from Enterprise Manager Ops Center Hosted Interface and Service Request

Below the title bar, a set of icons summarizes the Problems status. The number next to an icond indicates the number of assets reporting that condition. Hold or hover the mouse cursor over the icon to see information that identifies the most recent problem. These icons keep you updated while you perform other operations. To see information about the problems, click on the icon to change the display: the Message Center opens in the Navigation pane and the center pane shows the tables for that problem category.

Figure 1–11   Problem Icons

Searching in Enterprise Manager Ops Center

The user interface provides several search functions:

- In the Navigation pane, you can search for the asset that matches the string you enter. The first result is highlighted in the asset tree for your selection. The up and down arrows show you the next matching asset.

- When All Assets is selected in the center pane, search fields are displayed in the Managed Assets, Available to be Managed Assets, and Unclassified Assets tabs. These search fields search only within the specified category.

- The global search field, in the upper-right corner of the user interface, can search for any Managed Network, Storage Library, or Asset.

- In the Jobs pane, you can search for the job that matches the string you enter.

Checking the Version Number of the Product Software

You might need to report the version number of the software. This version number is also the version of the Enterprise Controller. At any time, click on Administration in the Action pane. The center pane's default display is the Configuration tab for the Enterprise Controller and it displays the complete version number of the software, including the build number.

2

Group Management

Groups are administrative structures that organize assets so that you can locate an asset quickly or perform operations on all assets of the same type. Groups can contain any number of assets, and assets can be members of more than one group. This chapter describes how to use groups to manage your assets.

Smart Groups

Oracle Enterprise Manager Ops Center creates groups of major asset types automatically. Within these Smart Groups, a subgroup is created for each type of asset that Oracle Enterprise Manager Ops Center is managing. Use Smart Groups to locate and view assets of a specific type. You can also act on Smart Groups, such as changing monitoring thresholds and updating discovery credentials.

Oracle Enterprise Manager Ops Center creates these Smart Groups:

- Operating Systems – Contains all operating systems with subgroups for each type of operating system such as Oracle Solaris OS, Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux. The subgroups are further organized by version, such as Oracle Solaris 9 and Oracle Solaris 10 software.

- Servers

- Storage

- Systems – Contains all hardware that can receive firmware updates and OS provisioning with subgroups for each type of hardware (such as SPARC and x86). The subgroups contain are further organized by subtypes such as x86 32-bit and x64.

- Chassis – Contains all hardware that can receive firmware updates but not OS provisioning with subgroups for each type of hardware such as Sun Blade 6000 Series and Sun Blade 8000 Series.

- Network Switches

- Services – Contains all assets organized by the actions for which they can be targets. It contains subgroups for each of the major actions in the software: OS Update, Manage-Monitor, OS Provisioning, FW Provisioning, and Virtualization. Each of these subgroups contains the assets that can be targeted with that action, and includes related actions in the Actions pane. For example, the OS Update group contains all operating systems that can be updated, and includes New Update OS Job in its Actions pane.

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.

Viewing Smart Groups

Smart groups are automatically generated groups that organize your assets by type.

To View a Smart Group

1. Select the Assets in the Navigation pane.

2. Select a Smart Group from the dropdown list at the top of the pane.

3. (Optional) Select a subgroup in the Navigation pane. The smart group is displayed.

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.

To View Group Data

1. Expand Assets in the Navigation pane.

2. Select All User-Defined Groups.

3. Select a group. The group dashboard is displayed. This includes:

   ■ 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.

Starting Jobs on a Group

You can perform operations on all the assets in a homogeneous group by choosing the group as the target for asset-specific actions. For example, you can provision an OS on many systems simultaneously.

The selected action is applied only to compatible assets. For example, if you select the Locator Light On action for a group of servers, the job targets only servers that have locator lights.

To Start a Job on a Group

1. Select All User-Defined Groups from the Assets list in the Navigation Pane.

2. Select the group. Asset-specific actions are displayed in the Actions pane.

3. Click an action. The action is applied to all compatible assets within the group.

Creating a Group

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

To Create a Group

1. Expand All Assets in the Assets section of the Navigation pane.

2. Click Create Group in the Actions pane. The Create New Group/Subgroup window is displayed.

3. Enter group information, then click Next. Group information includes the following:

   ■ **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 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** – Uset this option to specify any existing group as a child of this group.

   ■ **Preview Group Before Creation** – Select this option to see the assets and subgroups that will be added to the new group.

4. If you selected Configure Group Rules, the Configure Asset Type Rules step is displayed. Configure one or more rules, then click Next. Configuring a rule includes the following:

   ■ **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. 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.

5. If you selected Configure Subgroups, the Configure Subgroups page is displayed. Drag one or more groups from the Available Groups list to the Selected Groups list, then click Next.

6. If you selected Preview Group Before Creation, the Preview Group page is displayed. Review the list of assets and subgroups that will be added to the new group, then click Next.

7. The summary page is displayed. Click Finish. A Create Group job is launched.

See the *Enterprise Manager Ops Center Reference 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 manuall are not removed.

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

To Edit a Group

1. Expand All User-Defined Groups from the Assets drop-down list in the Navigation pane.

2. Select a group.

3. Click Edit Group in the Actions pane.

4. Edit the group information, then click Next. Group information includes the following:

   - **Group Name**

   - **Description**

   - **Group Location**

   - **Configure Group Rules** – Select this option to add or change the rules for a group. Group rules add any assets to the group that match the attributes and rules.

   - **Configure Subgroups** – Select this option to add or change child groups of this group.

   - **Preview Group Before Creation** – Select this option to view the assets and groups that will be in the group when you commit the changes. This option is available only if you change group rules or subgroups.

5. If you selected Configure Group Rules, the Configure Asset Type Rules step is displayed. Configure one or more rules, then click Next. Configuring a rule includes the following:

   - **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. 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.

6. If you selected Configure Subgroups, the Configure Subgroups page is displayed. Drag one or more groups from the Available Groups list to the Selected Groups list to add them as subgroups, then click Next.

7. If you selected Preview Group Before Creation, the Preview Group page is displayed. Review the list of assets and subgroups that will be added to the new group, then click Next.

8. The summary page is displayed. Click Finish. An Update Asset Group job is launched.

Adding Assets to a Group

You place assets into groups to aid in organizing your operations. 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 Group

1. Expand All Assets in the Assets section of the Navigation pane.

2. Select one or more assets in either of the following ways:

   - Select an asset from the Asset list, then click Add Asset to Group in the Actions pane.

   - Select one or more assets from the Managed Assets list in the center pane, then click the Add Asset to Group icon.

3. In the Add Assets to Group/Subgroup window, select the destination group, then click Add Assets to Group. An Add Resource to Group job is launched.

Adding SPARC Enterprise M-Series Servers to a Group

For the SPARC Enterprise M-Series servers, create a homogeneous group for each type of server and then use the group as a target for actions such as updating firmware on all the servers in the group or provisioning the OS on all the dynamic system domains. The homogenous groups must contain only servers of the same model.

To Add a SPARC Enterprise M-Series Server to a Group

1. Expand All Assets in the Navigation pane.

2. Select an M-Series server.

3. Click Add Asset to Group in the Actions pane.

4. In the Add Assets to Group/Subgroup window, select the destination group, then click Add Asset to Group. An Add Resource to Group job is launched.

5. After the asset is added to the group, assign the Admin role to the group for users that must manage the assets.

Moving Assets to a Group

When you move assets to a new 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 re-added to the group after you move it. It might be possible to edit the group's rule to exclude the asset you have moved.

To Move Assets to a New Group

1. Expand All User-Defined Groups from the Assets list in the Navigation pane.

2. Select the group.

3. Select one or more assets from the list in the center pane.

4. Click the Move Asset to Group icon. The Move Assets to Group/Subgroup window is displayed.

5. Select the destination group.

6. Click Move Assets to Group to start the job.

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 and it is removed manually, the asset is not re-added.

To Remove an Asset from a Group

1. Select All User-Defined Groups from the Assets list in the Navigation Pane.

2. Select the group.

3. Select the asset from the list in the center pane.

4. Click the Remove Asset from Group icon in the center pane.

5. Click OK to confirm you want to remove the asset. A job is launched to remove the asset from the group.

Moving a Group

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

To Move a Group

1. Select All User Defined Groups from the Assets drop-down list in the Navigation pane.

2. Select a group.

3. Click Move Group in the Actions pane. The Move Group window is displayed.

4. Select a destination for the group:

   - Top level

   - Another user-defined group

5. Click Move Group. A job is launched to move the group.

Deleting a Group

Deleting a group removes the group and all of its subgroups from the user-defined group hierarchy. It does not delete any of the assets.

To Delete a Group

1. Expand All User Defined Groups from the Assets drop-down list in the Navigation Pane.

2. Select a group.

3. Click Delete Group in the Actions pane.

4. Click OK to confirm you want to delete the group. A job is launched to delete the group.

3

Message Center

The Message Center collects and presents all problems, alerts, and notifications. Use the Message Center to view and manage problems, notifications, and service requests and to warranty information. This chapter describes the Message Center.

Contents of the Message Center

The Message Center displays and manages alerts and alarms, which are called Problems, it generates notifications, it can view and display service requests, and it can display the warranty information for an asset.

Problems

An alert is generated when a monitored attribute does not meet the monitoring parameters. If a problem does not already exist, a new problem appears in the Unassigned Problems queue in the Message Center. If you have the Manage or Admin role, you can assign problems to other users with the Manage or Admin role for that asset. From the Message Center you can view and take action on problems.

Use Maintenance Mode, as described in Problem Management, to temporarily disable assets from triggering alerts or problems when you know an incident will occur, such as when you plan on performing maintenance, and you do not want problems to appear in the user interface.

Notifications

Notifications enable you to configure the software to send you an e-mail or pager message when a critical or warning problem is detected. Different users can be given different notification profiles, letting you control what severity of messages are sent to each user and by what method.

Service Requests and Warranty Information

You can file and view service requests to My Oracle Support from the Message Center and view warranty information. See Oracle Services for information about using these features.

Viewing Unassigned Problems

New problems appear in the Unassigned Problems queue in the Message Center. A bar chart on the page visually displays the number of new problems by severity. The page also contains a table that categorizes the problems by severity and asset type. Select a row in the table to display all unassigned problems for the selected severity

category and to drill down for more details. Each problem is assigned an automatically generated ID to help track the issue.

The following information is available:

- Severity
- Problem ID
- Problem Description
- Cause
- Number of Suggested Actions
- Start Time
- Duration
- URL

To View Unassigned Problems

1. Click Message Center in the Navigation pane.
2. Click Unassigned Problems in the Navigation pane.

Viewing My Problems

Problems assigned to you appear in the My Problems queue in the Message Center. A bar chart on the page visually displays the number of new problems by severity. The page also contains a table that categorizes the problems by severity and asset type. Select a row in the table to display all problems for the selected severity category and to drill down for more details. Each problem is assigned an automatically generated ID to help track the issue.

To View Your Problems

1. Click Message Center in the Navigation pane.
2. Click My Problems in the Navigation pane.

Viewing Problems Assigned to Others

Problems assigned to other users appear in the Problems Assigned to Others queue in the Message Center. A bar chart on the page visually displays the number of new problems by severity. The page also contains a table that categorizes the problems by severity and asset type. Select a row in the table to display all problems assigned to others for the selected severity category and to drill down for more details. Each problem is assigned an automatically generated ID to help track the issue.

To View Problems Assigned to Others

1. Click Message Center in the Navigation pane.
2. Click Problems Assigned to Others in the Navigation pane.

Placing an Asset in Maintenance Mode

You can place assets in maintenance mode prior to maintenance. Assets in maintenance mode do not generate notifications.

To Place an Asset in Maintenance Mode

1. Select an asset in the Navigation pane.

2. Click Place in Maintenance Mode in the Actions pane. A confirmation window is displayed.

3. Click Place.

Removing an Asset from Maintenance Mode

You can remove assets in maintenance mode after maintenance. Assets in maintenance mode do not generate notifications. Once removed from maintenance mode, the assets will generate notifications as normal.

To Remove an Asset from Maintenance Mode

1. Select an asset in the Navigation pane.

2. Click Remove From Maintenance Mode in the Actions pane. A confirmation window is displayed.

3. Click Remove.

Viewing Notifications

Ops Center generates notifications for all event operations and when the set thresholds are exceeded.

You can enable or disable the notifications for a server or blade, chassis, or OS. The status of notifications for the asset is displayed in the asset's monitoring window.

To View Notifications

1. Expand Message Center in the Navigation pane.

2. Click Notifications. The center pane displays a table of all notifications.

3. (Optional) Use the search function in the upper right section of the window to locate specific notifications.

Deleting Notifications

You can delete any notification. You can also delete all notifications.

To Delete Specific Notifications

1. Click the notifications icon in the upper right section of the BUI. The notifications window is displayed.

2. Select the notification or notifications you want to delete.

3. Click the Delete icon.

To Delete All Notifications

1. Click the notifications icon in the upper right section of the BUI. The notifications window is displayed.

2. Click the Delete All Notifications icon. A confirmation window is displayed.

3. Click Yes. All notifications are deleted.

4

Oracle Services

Oracle Services provides integrated methods for maintaining and displaying current contracts, warranty information, contract dates, and service requests. This chapter describes Oracle Services for managed assets.

About Contracts, Warranties, and Service Requests

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 problem in Enterprise Manager Ops Center, view service request details, and file a service request.

- Contracts and Warranties

  Maintaining a valid inventory of the asssets in your data center, including contracts and warranties, can be a time-consuming and labor-intensive process. Use 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, Enterprise Manager Ops Center generates an alert .

- Service Request

  Use Enterprise Manager Ops Center You to create new service requests. You can also review your requests and the requests of other users.

  > **Note:**   You cannot display service requests created outside of Oracle Enterprise Manager Ops Center. To see the status of any service requests filed outside of 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 Mangager 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. 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 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 4–1   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 a Problem in the Message Center and the contract details are displayed in orange text in the asset's Summary tab.

> **Note:**   Updating contract and warranty information requires running the product software in Connected mofe. 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. If 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 ni 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 the current and completed service requests.

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.

Figure 4–2   All Open Service Requests

3. To view details of a particular service request, highlight a row, then click the View Service Request icon.

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 a problem 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 Problem

1. Click Message Center in the Navigation pane.

2. Click My Problems or Unassigned Problems

3. Select the problem, then click the Open Service Request icon in the center pane.

Figure 4–3   Open Service Request icon

To File a Service Request From an Asset

1. Click the hardware in the Assets section of the Navigation pane.

2. Click Open Service Request in the Actions pane. .

5

Job Management

All tasks, such as asset discovery or OS provisioning, create jobs. Because each task runs as a job, you do not have to wait for one job to finish before you start the next action. This chapter describes how to work with jobs and work with the results.

About Jobs

To launch a new job, go to the Navigation pane to select an asset and then go to the Actions pane to select an operation. For some tasks, you can start from the Actions pane and then identify the targets. After the job is launched, you can monitor and control it from the Jobs pane.

From the Jobs pane, you monitor the progress of current jobs and can review historical jobs. You can see the status of all jobs and view detailed information about a specific job in its current state. You can view information about jobs, view categories of jobs, and take actions on jobs. The following information is displayed for each job:

- 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 of the job's 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 – A user stopped the job.
    - 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.
- Job ID – A unique identification number.
- Type of job – For example, the Discovery-Custom type identifies a job as a result of a Custom Discovery action.
- Name of job
- Mode of job – Simulated or Actual Run. Some action can be simulated.
- Owner of job – The user who launched the job.
- Date and time – The timestamp of when the job started or is scheduled to start.

■   Elapsed time

You can take the following actions on jobs:

■   View details about the job, including the status of all tasks.

■   Respond to prompts during a job.

■   Interrupt a running job. You can re-run the job later.

■   Remove a job from the history.

■   Restart 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.

Viewing Job Details

You can view detailed information about a job:

■   Job ID

■   Run ID: If a job has been run multiple times, each run of the job will have a separate run ID.

■   Job Type

■   Start Date

■   Elapsed Time

■   Status

■   Creation Date

■   Description

■   Mode (Simulation or Actual Run)

■   Owner

■   Task Execution Order

■   Failure policy

■   Recurring job status and Next Scheduled Time

■   List of tasks

■   Task execution order

■   Task progress/results

■   Event logs: Event logs are generated by tasks as they progress.

To View Job Details

1. Expand the Jobs pane at the bottom of the UI.

2. Select a job.

3. Click the View Job Details icon or double-click the job. The job details of the most recent run ID are displayed.

4. (Optional) Click the down-arrow near the Job Details title to see information about how the job was run.

5. (Optional) Select a different run ID from the list. The job details for the new run ID are displayed.

To see only those jobs for a particular asset, use the asset's Jobs tab to see both running jobs and jobs that have been completed.

To View Jobs for an Asset

1. Select the asset in the Assets section of the Navigation pane. The Summary page of the selected asset is displayed in the center pane.

2. Select the Jobs tab in the center pane. The Jobs page is displayed, showing the Current and Historical jobs for the selected asset, and their status.

Answering Questions

Some jobs require user input to complete. A job that requires a response has the status of Waiting for User Input.

To Answer Questions For a Job

1. Select the job that requires input.

2. Click the Answer Questions icon. The Answer Pending Questions window is displayed.

3. (Optional) Click the page icon next to a component to display the component details.

4. Answer the questions in one of the following ways. Click Reset to set the response for every question to the default value.

   - Select Yes or No for each question using the buttons in the right column.

   - Click Yes to All to set the response for every question to yes.

   - Click No to All to set the response for every question to no.

5. Click Submit. The job is resubmitted 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 tasks that have completed, been interrupted, and not started.

To Stop a Job

1. Expand the Jobs pane at the bottom of the UI.

2. Select the job ID of a running job.

3. Click the Stop Selected Jobs button.

4. Click Stop Job to confirm. The job is stopped.

Re-running a Job

If a job is completed, you can repeat the job with a new Run ID.

If a job is partially completed, failed, or stopped, you can repeat the job on failed or incomplete targets.

To Re-Run a Job

1. Expand the Jobs pane at the bottom of the UI.

2. Select one or more job IDs to re-run.

3. Click the Re-run Selected Job button.

4. Click Run Job to confirm. The job runs again with a new Run ID.

To Re-Run a Job on Failed Targets

1. Expand the Jobs pane at the bottom of the UI.

2. Select the job that you want to re-run.

3. Click View Job Details.

4. Click Rerun Job on Failed Targets.

5. Click Run Job. The job is run again with the same Run ID.

Copying a Job

You can copy a job, using an old job as a template for a new job.

To Copy a Job

1. Expand the Jobs pane at the bottom of the UI.

2. Select the job that you want to copy.

3. Click Copy Job. The New Job wizard is displayed.

4. (Optional) Modify the job information for its new task. Click Next.

5. Select a schedule for the job and click Next.

6. Review the job information 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 stop the job and then delete it.

To Delete a Job

1. Expand the Jobs pane at the bottom of the UI.

2. Select the job ID or IDs that you want to delete.

3. Click the Delete Selected Jobs icon.

4. Click Delete Job to confirm. The job is deleted.

Changing the Maximum Time for a Job

The default time for a job is 180 minutes. Complex jobs or some environments can require more time for a job to finish.

To Change the Maximum Time for a Job

1. Edit the /var/opt/sun/xvm/satellite.properties file.

2. Edit the `osp.default_timeout` property to increase the number of minutes allowed for jobs. The example adds another hour to the expiration time.

```
#################################
    # OS Provisioning
    #################################
    # Default Timeout of OSP Job
    osp.default_timeout=240
```

3. Restart the Enterprise Controller.

6

Hardware Monitoring

All hardware assets are monitored for their status, according to the asset's monitoring profile. The user interface reports information for a selected asset in a series of tabbed displays. The tabs and the type of information is specific for the asset type.

- Health status

- Power state

- Power usage

- Hardware variables and connectivity

Based on your observations, you can control your hardware assets and do the following actions:

- Create charts

- Use locator lights to identify a specific asset

- Power systems on and off

- Reset a server

- Get access to the serial console

Viewing Hardware Details

Enterprise Manager Ops Center reports the information that it can acquire from an individual asset. Hardware information is displayed in increasing detail on the Summary tab, the Hardware tab, and the Monitoring tab.

See the *Enterprise Manager Ops Center Reference Guide* for a list of the asset attributes that can be used in monitoring.

To View Details About Hardware

1. Expand Assets in the the Navigation pane.

2. Select All Assets.

3. (Optional) Select an asset type to filter the assets. The default groups for hardware are Systems, Chassis, and Switches.

4. Select an asset. The Summary tab provides a high-level view of the asset's attributes and, for most assets, displays the firmware version.

5. Select the Hardware tab to see information about that asset's hardware and firmware components. Depending on the asset, you can refine the information to specific components.

6. Select the Monitoring tab, if the asset has one, to view the current state of various hardware variables. For each variable, this tab also shows the values for the warning threshold, the critical threshold, and the non-recoverable threshold.

Server Details

For server hardware, the Summary page displays:

- Name
- Description and Tags
- 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
- Notification

Use the Hardware tab to view information about each component of the system:

- System: Description, type, and version of all firmware installed except for disk firmware. See the Disk tab for firmware version.
- CPU: Name, Model, Architecture, Speed, Manufacturer
- Memory: Name, Type, Size in bytes, Manufacturer, Part number, Serial number
- Network Adaptors: Name or each, MAC Address, Manufacturer, Part number, Serial Number
- Disks: Name of each, Model, Size in bytes, Slot ID, Node ID, Firmware Version, Manufacturer, Root Disk, RAID Disk
- Power Supply: Name, Manufacturer, Part number, Serial number
- Disk Controller: Name, Model Number, Firmware Version, BIOS Version, PCI Address, PCI Version ID
- Disk Expander: Name, Manufacturer, Version, Model Number, Firmware Version, Chassis ID
- Fan Tray: Name, Manufacturer, Part number, Serial number
- Fans: Name, Speed

Chassis Details

For chassis hardware as a group, the Summary page shows:

- The five largest consumers of CPU

- The five largest consumers of memory

- The five largest consumers of the network

For chassis hardware, the Summary page displays:

- Group Name

- Description

- Location

- Type

M-Series Server

The hardware resources in a SPARC Enterprise M-Series Server are divided into one or more logical units, called dynamic system domains. Enterprise Manager Ops Center can monitor each domain, in addition to the server hardware.

For an M-Series server, the Dashboard tab displays:

- Number of Dynamic System Domains it is supporting

- Model

- Product Serial Number

- Description

- Support contract

- XCP Firmware version

- OBP Firmware version

- XSCF Firmware version

- Operator Panel Switch Status: Locked

- Current Alert Status

The Summary tab repeats some of the Dashboard's information and adds details. 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.

- Name
- Model
- Product Serial Number
- Management IP
- MAC Address
- Current Alert Status
- Power
- Locator Light
- Notification

- All firmware versions

- For each domain: Name, Model, Health, Power, Locator Light, Notification

The Hardware tab shows the state of the server or, if a Dynamic System Domain is selected, the state of that domain. The Hardware tab reports the following:

- Model

- Serial Number

- State

- Power

- Locator Light

- Notification

- Operator Panel Switch State At the System level, the hardware report includes:

- The Unallocated Resources table lists all the physical system boards and their status: PSD ID, Assignment Status, Power Status, Connection Status, Diagnostics Status, 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, 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, Operational Status

You change the display to show 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, Value

- Memory: Name, Type, Size in bytes, Serial number, Part number, Status For Sensors: Name, Description, Type, Value

- Board: Name, Serial number, Part number, Memory mirrored, Version, Status

- Power Supply: Name, Serial number, Part Number, Status For Sensors: Name, Description, Type, Value

- Board: Name, XSB Mode, Memory Mirrored, Serial Number, Part Number, Version, Status For Sensors: Name, Description, Type, Value

- IO Unit: Name, Serial Number, Part Number Version, Status For Sensors: Name, Description, Type, Value

- XSCF: Name, Host Name, Serial Number, Part Number, Version, Status

- Fan Tray: Name, Manufacturer, Part number, Serial number

- Fans: Name, Speed For Sensors: Name, Description, Type, Value

Enterprise Manager Ops Center monitors the voltage for the Board and IO Unit components and the speed for the Fan components. Click on the Monitoring tab to see the actual value and the threshold values.

See SPARC Enterprise M-Series Server Support for information requirements.

Sun ZFS Storage Appliance

The Sun ZFS Storage Appliance support both file storage and application use.

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

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

Switch Details

Enterprise Manager Ops Center can manage 10G Ethernet Fabric Switches and Datacenter Infiniband Switches. These types of switches reside in the system or blade system, providing the switch fabric.

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

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),

For more information about supported hardware, see the Supported Systems.

Monitoring Hardware Health

Enterprise Manager Ops Center monitors the sensors in the hardware and displays the following information:

- CPU temperature

- Ambient temperature

- Fan speed in revolutions per minute

- Voltages

- LEDs

States of Hardware Health

If a hardware asset can report a value for a hardware variable, Enterprise Manager Ops Center reports its current state and compares it to the threshold value.

- Good – The hardware asset is working properly.

- Unknown – 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 – 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.

Monitoring Hardware Variables

Enterprise Manager Ops Center monitors hardware assets according to the monitoring profile for that type of asset. The following hardware variables can be monitored:

- Current

- Disk

- Fan

- Power supply

- Temperature

- Voltage

To see the default profile for monitoring hardware, see Monitoring Profiles.

To View Hardware Variables

1. Expand Assets in the Navigation pane.

2. Expand the hardware type.

3. Select the hardware. The Summary page of the hardware is displayed in the center pane.

4. Click the Monitoring tab to view the variables.

5. Select the Variable type. The variables are listed with their Warning, Critical, and Non-recoverable threshold values.

See Editing A Monitoring Rule to change a threshold value.

Monitoring Connectivity

Connectivity is the network interface of the system. You can view information about a hardware asset's Network Interface Card (NIC).

To Monitor Connectivity

1. Expand Assets in the Navigation pane.

2. Expand the hardware type and select the hardware asset. The Dashboard page of the hardware is displayed in the center panel.

3. Click the Connectivity tab. The details about the network interface cards such as name, connection status, MAC address, and the corresponding IP address are displayed.

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. Use the following procedure to check whether any hardware asset can report its power utlization.

Checking Power Capability

1. Expand the All Assets section of the Navigation pane.

2. Select the server.

3. Click the Capabilities tab.

4. In the list of enabled capabilities, locate Report Power Usage.

Viewing Power Utilization

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. The current values are reported for the following attributes:

- Wattage

- System Load for an OS

- Utilization Per Cent for an Oracle VM Server for SPARC

- Incoming air temperature and outgoing air temperatur

- Power Policy

- Cost Per KiloWatthour

- 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 Editing the Energy Cost for more information.

- Total Power Cost for one day

The data over time is represented in the following graphs:

- Power Consumption and Utilization: By default, the graph shows the power consumed in the last day in watts. If the server is shut down, the graph shows any existing historical data.

- Temperature and Fan Speed: By default, the graph shows the incoming air temperature and the outgoing air temperature in Fahrenheit, and the average fan speed in RPM. Click on any point on the graph to see that data for that point in time.

By default, the graphs are in Live mode, which reports new information every five seconds. Click on the Live button in the graph's toolbar to make the information static. This enables you to change the period of the graphs to one of the following:

- One hour (1H) – One point for every five minutes

- One day (1D) – One point for every five minutes

- Five days (5D – One point for every five 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 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 6W 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 1D 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.

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 heterogenous 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 heterogenous 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.

To View Power Utilization Charts

1. Select a hardware asset from the Assets section in the Navigation pane. You can also select a group or virtual pool from the Assets section if the group contains an hardware asset.

2. Select the Energy tab to view the current data. The Power Utilization and Consumption graph shows the power use in the last hour. The Temperature and Fan Speed graph shows the incoming air temperature, the outgoing air temperature, and the fan speed.

3. To see historical data, click on the Live button to stop updating the data.

4. Click on the Display field and select one of the time periods. Both graphs change immediately.

5. To see the cost of the power use, click on the Chart Options drop-down list and select the Chart Cost option.

6. To change the type of graph, click on the Charts tab and select either Bar or Area for the type of graph.

To Export Power Utilization Charts

1. Expand Assets in the Navigation pane.

2. Expand the hardware type and select the hardware asset.

3. Select the Energy tab to display the graphs. In the graph's toolbar, click the Export to CSV/XML icon. or Select the Charts tab and then click the Export Chart Data button. The Export Data window is displayed.

4. Select the format in which you want to store the data, either CSV and XML format.

5. If you have already set the time period of the chart, select the option Current View for the Time Period. If you want the data for six months, select the 6 Months option for the Time Period.

6. Click Export to store the data. The data is exported and saved in the directory where you are running the user interface.

Configuring Power Utilization

The CPUs of a server have the ability to manage power consumption.

To Set Policy for an Asset's Power Utilization

1. Select a server, a group of servers, or some members of a group of servers from the Assets list in the Navigation pane.

2. Click Set Power Policy in the Actions pane. If a selected server does not support power configuration, the action is not available. The Modify/Change Power Control Settings popup is displayed.

3. Click one of the options for configuring a power policy:

   - Set Elastic Mode to enable CPU power management. This option conserves power but decreases performance.

   - Set Performance Mode to disable CPU power management. This option increases performance but increases power consumption.

4. Click Close.

Using a Hardware Monitoring Profile

A hardware monitoring profile is a set of rules applied to a hardware asset. If a status changes or a threshold is crossed, an alert is created. Enterprise Manager Ops Center provides default profiles for each asset type. You can create new profiles or modify existing profiles.

To Display the Current Hardware Monitoring Profile:

1. Expand Assets section of the Navigation pane.

2. Expand All Assets.

3. Select a hardware asset.

4. Click on the Monitoring tab. For each rule in the profile, the Monitor tab shows the name of the rule, the limits of the rule, and whether the rule is in effect.

To Apply a Hardware Monitoring Profile:

1. Expand Assets section of the Navigation pane.

2. Expand All Assets.

3. Select a hardware asset.

4. Click on the Monitoring tab.

5. Click Apply a Monitoring Profile in the Action pane. The Apply Monitoring Profile wizard starts, with the selected hardware asset specified as the target.

6. Click on the Profile list and select a profile from the list. To see details of the profile, click the icon.

7. Click Apply. The asset you select is now monitored, according to the rules in the profile.

Managing Locator Lights

You can activate or deactivate LED locator lights on managed servers and blades to locate a specific asset among many of the same type. This can simplify physical maintenance tasks.

To Activate Locator Lights

1. Expand Assets in the Navigation pane.

2. Expand All Assets and select one hardware asset or a homogeneous group.

3. Click Locator Lights On in the Actions pane. The LED locator lights on the asset or assets are activated.

To Deactivate Locator Lights

1. Expand Assets in the Navigation pane.

2. Expand All Assets and select one hardware asset or a homogeneous group.

3. Click Locator Lights Off in the Actions pane. The LED locator lights on the asset or assets are deactivated.

Resetting a Server

You can reset a server or a set of servers.

To Reset a Server

1. Select a server or server group from either the Navigation pane or the Membership Graph.

2. Click Reset Server(s) to reset the system. For a group, select the list of servers from the group and click the Reset Server(s) icon. A Reset pop-up window appears with the following options:

   ■ Reset

   ■ Force Reset

   ■ Reset with Network Boot

3. Click the appropriate option as required to reset the selected hardware.

Powering an Asset On and Off

You can use Enterprise Manager Ops Center to power on and power off a server or chassis. Stopping a server initiates a graceful shutdown of the operating system and subsequent power-off of the server. If no operating system is installed, you must force a shutdown of the server.

You can power on a managed server or a server group. If boot PROMs are configured, the servers will boot.

To Power On a Server or Chassis

1. Expand Assets in the Navigation pane.

2. Select a hardware asset.

3. Choose Power On in the Actions pane. For server groups, select the servers from the list and click the Power On icon.

4. To start the server, choose Default Power On.

5. To start the server and use a manual network boot, choose Power On with Network Boot. A manual network boot is required for OS-based or manually discovered servers.

6. Click OK. A job is submitted. Jobs initiated on groups of servers will run longer than jobs on individual servers.

To Power Off a Server or Chassis

1. Expand Assets in the Navigation pane.

2. Select a hardware asset.

3. Choose Power Off in the Actions pane. For server groups, select the servers from the list and click the Power Off icon. A Power OFF pop-up window opens. The following options are available:

   ■ Power OFF

   ■ Force Power OFF

   ■ Emergency Power OFF (This is available only for Chassis.)

4. Click the appropriate option. A job is initiated. The selected asset is be powered off.

Accessing the Serial Console

From the Enterprise Manager Ops Center UI, you can get access to a managed hardware asset's operating system through the asset's serial console. The Enterprise Manager Ops Center UI opens the asset's serial console, starts an ssh session, and logs into the operating system using the stored credentials for the asset's service processor. You can then issue operating system commands but you cannot issue service processor commands.

Before You Begin

- Enable ssh. Use Custom Discovery to discover the OS and the SPs on the hardware so that you can enable ssh access. If you have already discovered and managed an asset and now want to use the serial console, re-discover the asset using Custom Discovery.

- Verify network access. Verify that the Enterprise Controller can use ssh on Port 22 so that the Enterprise Controller can reach the asset's Proxy Controller or agent.

- Verify that your role gives you permission to access and change the asset.

To Access the Serial Console

1. Expand All Assets in the Assets section of the Navigation pane.

2. (Optional) Filter the assets by selecting an asset type in the View window in the center pane.

3. Select the hardware asset. The asset details are displayed in the center pane.

4. Click the Console tab.

5. Click Enable Console to activate the console. A job is submitted to activate the console and log you into the operating system.

6. Issue operating system commands.

7. (Optional) Click Undock to detach the console from the UI and move it to another location on your monitor.

8. To close the console, press the ESC+(keys.

7

OS Monitoring

Use Enterprise Manager Ops Center to manage and monitor your Windows, Oracle Solaris, and Linux operating systems because an OS is a managed asset. You can monitor the status, activity, and usage of the operating systems. An agent embedded with the OS asset monitors the following variables :

- Connectivity
- File system status
- CPU usage
- Memory usage
- Network usage
- Monitoring variables

To monitor an OS, Enterprise Manager Ops Center compares the value of an OS attribute to a threshold value at regular intervals and reports a status.

- Good – The asset is working properly.
- Unknown – An error occurred while attempting to retrieve information about the health of the asset. The asset is connected but does not return any information.
- Unreachable – The asset's agent IP address cannot be contacted. This indicates network problems or, if the asset is hardware, it might be in Standby mode.
- Warning – The value of an attribute has exceeded the warning threshold value. Take action to prevent a problem.
- Critical – The value of an antribute has exceeded the critical threshold. Take action to correct the problem.

Enterprise Manager Ops Center provides default thresholds in its OS monitoring profiles. You can adjust the thresholds to meet your data center guidelines and you can set different threshold for each operating system.

Historical data includes CPU, memory, I/O and power data. You can create reports and can graph the historical data for trend analysis and forecasting. See Reports for more information.

 To see the default profile for OS Monitoring, see Monitoring Profiles. To change the threshold values for OS monitoring, see Editing A Monitoring Rule.

Monitoring an OS

Enterprise Manager Ops Center can monitor Oracle Solaris, Linux, and Microsoft Windows operating systems. For Windows systems, the attributes that are monitored

are determined by what has been configured in the Windows Management Instrumentation (WMI) remote monitoring capabilities. You must configure the WMI on the Windows system to manage or monitor its OS.

To Monitor an OS

1. Select an OS asset in the Navigation pane to see its current status. The Dashboard page in the center pane includes the Current Alert Status field.

2. For a status of Warning or Critical, scroll down the Dashboard page to view the Status table. This table displays the five (5) most recent problems on the asset and the category and name of the threshold that has been exceeded.

3. Click the Summary tab to see a table of all the OS monitoring variables and the status of each one.

4. Click on the Monitoring tab to see a list of all the monitoring variables, both hardware and OS, the latest value of each one, and the values of their thresholds. These threshold values are set by applying a monitoring profile.

5. Click on the Utilization tab to see the OS monitoring variables by category and in the context of time. The Summary view displays the version of the operating system and the following list of categories:

   - CPU Utilization
   - Memory Utilization
   - Swap Utilization
   - Network Bandwidth Utilization
   - File System Utilization
   - I/O Utilization
   - System Load

6. Click one of the categories to see that category's monitoring variable's status and a chart of its value for the last 24 hours. You can change the duration of the chart to another time period.

7. To export chart data, click the Charts tab and then select the category. The same chart as the one on the Utilization tab is displayed. Click the Export Chart Data button to choose options for creating a file.

8. To compare the history of different monitoring variables, click the Charts tab.

   - Click the option to include all lines on a single chart.
   - Select two or more categories of monitoring variables.
   - If you prefer, change the time period.

Rebooting an OS

Use the following procedure to reboot the OS.

To Reboot an OS

1. Select the OS from the Asset section of the Navigation pane.

2. Click Reboot in the Actions pane.

3. Confirm the reboot action.

Managing Boot Environments

The boot disk is the disk from which the Oracle Solaris OS kernal loads. You can partition the boot disk and copy an active OS environment to a fully bootable alternate OS disk or boot environment. You can update or perform routine maintenance on the alternate boot environment (ABE) while the active boot environment (BE) is running, or fall back to an alternate boot environment if there is a problem with the active boot environment. System administrators often create an alternate boot environment to test OS updates before deploying them to the active environment.

Use the Oracle Solaris Live Upgrade's Synchronize Boot Environments feature to create an alternate boot environment identical to the active boot environment, or use the `lucreate` command to create a custom alternate boot environment. Using Enterprise Manager Ops Center, you can use its Update Profiles or a custom profile to create an alternate boot environment.

All boot environments associated with an Oracle Solaris OS, including alternate boot environments, are displayed in the user interface. All boot environments are available for management, but you must make an ABE active to manage it. You can monitor only the active boot environment so to monitor an ABE, you must make it the active environment.

See About Oracle Solaris Live Upgrade for an overview of how Live Upgrade and alternate boot environments work in Enterprise Manager Ops Center . See for Additional Resources information about how to use the `lucreate` command.

Displaying Boot Environment Details

An alternate boot environment (ABE) appears as an attribute of the OS asset in the Asset pane. However, you cannot view the details of the alternate boot environment until you boot from it, making it the active or live boot environment. The following types of information is available for all boot environments:

- Name and description
- Status
    - Active or Live – Current boot environment
    - Inactive – An ABE that is available for activation
    - Invalid – An ABE that is not available for activation
- File System location and description
- Device size, mount location and options
- Date and time that the ABE was last synchronized

You can activate an ABE or schedule when to activate it. You can synchronize (sync) the boot environments.

To Display Boot Environment Details

1. Expand the Assets section of the navigation pane.

2. Click an Oracle Solaris OS. The OS Details are displayed in the center pane. If the OS has an ABE, the center pane includes a Boot Environments tab.

3. Click the Boot Environments tab to see all associated boot environments. The active boot environment is displayed at the top of the pane and shows the Active status. The alternate boot environments are displayed after the active boot environment. Shared files and directories appear in the Synch List.

4. Click the icon next to the Status to display details, including a list of file systems, the type of file system, device size, and the mount details.

Synchronizing Boot Environments

Before you activate an ABE, you synchronize it with the live boot environment to capture any changes that have been made. For example, after you install applications on your active boot environment, the current ABE does not contain these applications. When you synchronize boot environments, you create a new alternate boot environment that is identical to the active boot environment.

To Synchronize Boot Environments

1. Expand the Assets section of the navigation pane.

2. Click an Oracle Solaris OS. The OS Details are displayed in the center pane, which includes a Boot Environments tab.

3. Click the Boot Environments tab to see all associated boot environments. The active boot environment is displayed at the top of the pane and shows the Active status. The alternate boot environments are displayed after the active boot environment. Shared files and directories appear in the Synch List.

4. Click Synchronize Boot Environment in the Actions pane. A new alternate boot environment is created.

Activating a Boot Environment

When you activate an alternate boot environment, the current boot environment becomes an alternate boot environment. Before you activate an ABE, synchronize it with the current boot environment.

To Activate a Boot Environment

1. Expand the Assets section of the navigation pane.

2. Click an Oracle Solaris OS. The OS Details are displayed in the center pane, which includes a Boot Environments tab.

3. Click the Boot Environments tab to see all associated boot environments. The active boot environment is displayed at the top of the pane and shows the Active status. The alternate boot environments are displayed after the active boot environment. You can activate a boot environment that has an Inactive status.

4. Click the alternate boot environment..

5. Click Activate Boot Environment and Reboot in the Actions pane.

8

Reports

Use the Reports feature to consolidate changes to hardware, software, and job conditions. You can gather detailed information about job history, firmware, OS updates, and problems and then export that information to a file in CSV or PDF output.

> **Note:** Reports can consume significant disk space. To release disk space, move or delete report results. Reports in PDF and CSV files are stored in the `/var/opt/sun/xvm/reports` directory. You can relocate this directory to a dedicated disk.

About Charts

Enterprise Manager Ops Center collects information every five minutes on every managed asset and displays the last hour of data on the asset's Summary tab and OS Details page. To see utilization data over longer periods of time, up to six months, create one of the following charts, if it is available for the asset type:

- Power Utilization: Servers, chassis

- CPU Utilization: Operating system, operating system for a virtual machine, virtual host, virtual pool

- Memory Utilization: Operating system, virtual host, virtual pool

- Network Utilization: Operating system, operating system for a virtual machine, virtual host, virtual 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, data is collected 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

- After six months: Data is deleted.

Figure 8–1 is a sample chart is for a managed OS asset.

Figure 8–1   Sample Chart for an Managed OS

Group Charts

In addition to displaying charts for individual assets, you can see charts for a group of assets: the five most-utilized assets and the five least-utilized assets in the group. In the Summary page of a group, hourly usage is displayed as a bar chart. You can create the following charts for the managed groups:

- Server Group

  - Power

- Chassis Group

- OS Group

- CPU Utilization

- Network Utilization

- System Load

- Memory Utilization

- Host Group

  - CPU Utilization

  - Network Utilization

  - Allocated Memory

  - Power

Creating Charts

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.

By default, the graphs show one day of data. Click on the Live button in the graph's toolbar 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 data for these time periods is stored separately. For example, if a server has been managed for two hours and you select the 6W 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 1D view, the graph can display 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 toa 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 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.

To Display Charts

1. Expand Assets in the Navigation pane.

2. Select an asset. The Summary page is displayed in the center pane, which includes hourly charts of CPU, Network, and Memory utilization.

3. Click the Charts tab in the center pane.

4. Select the plot type. You can select resources such as CPU, network, memory file system, and system load.

5. Select the Chart option to display the charts as thumbnail images or single charts.

   - Select the thumbnail image option to see all resource utilization in one page.

   - Select the single chart option to compare multiple variables in one chart.

6. Select the type of chart. The available type depends on the type of asset.

7. Select the time interval represented in the X-axis.

8. Select how you want to measure utilization: as a percentage or in an appropriate unit of measure.

To Export Chart Data

1. Display the charts.

2. Click Export Chart Data.

3. Select the format and time period, and then click Export. The data is exported.

Creating System Information Reports

Create a system information report to obtain the information on assets such as OS, server, chassis, logical domains, global zone, non-global zone, and M-Series server.

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 and Description

   - Resource Type – Select an asset type

   - Schedule – Select Create Schedule to schedule the report.

   - Output Format – Select CSV or PDF format.

   - Select Targets – Add one or more targets by selecting them in the list of Available Items and clicking Add to Target List.

   When you have specified the report's contents, click Next. The attributes for the asset type are listed.

5. Select one or more attributes from the list and click Add or click Add All to choose all the attributes. Click Next.

6. To set a filter for an attribute, select the attribute and specify its condition. Click the Add icon to set filters for other attributes. When you are finished setting the filters, click Next.

7. If you selected Schedule in Step 4, the Report Schedule page is displayed. Select a schedule option, then click Next:

    ■ Now – Runs the report immediately.

    ■ Start Date – 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 on a recurring basis. Select the Start, End Time and the 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 AM, End Time at 12.00 AM and the Number of hours between runs as 2, the report is run at 6.00 AM, 8.00 AM, 10.00 AM and 12.00 AM.

8. Review the Summary.

    ■ To save the report as a template, click Save and Close. Templates are listed in the Report Templates section of the center pane.

    ■ To generate the report without saving it as a template, click Run and Close. Reports are listed in the Report Results section of 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 And Exporting Report Results for more information about generating a compliance job from the result.

About Update Reports

Use Update reports to check for new software updates and security advisories. You can get a general report, or test a system or installed package for available fixes. For auditing purposes, you can 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.

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 OS updates apply.

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 Linux and Oracle Solaris operating systems:

■ Change History – Provides a history of OS 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 should have these incidents installed. CVE IDs are unique, common identifiers for publicly known security vulnerabilities.

- Distribution Update – Provides basic information of all known distribution and local incidents.

- Package Compliance – Provides the details of the selected packages on your managed system that are compliant or not compliant with the latest recommended version available.

- RSC – Recommended Software Configuration – Provides information about the system compliance for installing a specific application, such as the Oracle 11g Database, on a Oracle Solaris or Linux OS.

- Service Pack Compliance (Linux only) – Provides information on incidents created by the publication and release of a service pack by a vendor. This helps to determine whether your system has the latest service packs released by the vendor.

- Solaris Update Compliance (Oracle Solaris OS only) – Provides information on whether an Oracle Solaris system is compliant with a specific update.

- Baseline Analysis (Oracle Solaris OS only) – Helps to check the compliance of systems against newly released Oracle Solaris baselines.

See Firmware And Os Update Reports for more details about each type of report and how to create and export a report.

Creating Problem Reports

The following types of problem reports are available:

- The Problem Summary Report is a historical report that summarizes information about all alerts and problems for a specified category, such as alarm state, alarm owner, asset type, date range, severity levels, and affected asset groups.

  Figure 8–2 is an example of a Problem Summary Report for a critical problem.

Figure 8–2 Problem Summary Report

ORACLE Enterprise Manager 11g Ops Center

Critical

Severity	State	ID	Description	Source	Owner	Creation Date
Critical	Unassigned	1057	CPU Usage Percentage has reached 11.51 on hs-x4100-2	hs-x4100-2	-	10/21/2010 10:24:42 MDT

- The Problem Detail Report contains detailed information about one or more problems. 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.

  Figure 8–3 is an example of a Detail Report for System Load Problems. The first page of the report is a summary of the report.

Figure 8–3   Summary of Problem Report

ORACLE Enterprise Manager 11g Ops Center

Report Name: System Load Detail Report
Description:
Run Date: Tue Oct 26 13:15:20 MDT 2010
Report Type: Problem Detail Report

Severity	State	ID	Description	Source	Owner	Creation Date
Warning	Unassigned	1080	System Load 5-Minute Average Per CPU has reached 1.06 on xvm-t2000-9	xvm-t2000-9	-	10/22/2010 11:23:47 MDT
Warning	Unassigned	1078	System Load 15-Minute Average Per CPU has reached 1.27 on xvm-t2000-9	xvm-t2000-9	-	10/22/2010 11:18:46 MDT
Warning	Unassigned	631	System Load 5-Minute Average Per CPU has reached 1.295 on xvm-t2000-3	xvm-t2000-3	-	10/21/2010 01:51:56 MDT
Warning	Unassigned	629	System Load 15-Minute Average Per CPU has reached 1.045 on xvm-t2000-3	xvm-t2000-3	-	10/21/2010 01:21:57 MDT
Warning	Unassigned	608	System Load 15-Minute Average Per CPU core has reached 1.13 on xvm-280-4	xvm-280-4	-	10/20/2010 23:06:13 MDT
Warning	Unassigned	611	System Load 15-Minute Average Per CPU core has reached 1.02 on xvm-280-2	xvm-280-2	-	10/20/2010 23:16:14 MDT

After the summary page, each problem has four pages:

- Details of the problem, as shown in Figure 8–4
- Suggested actions, if any
- Alerts History, as shown in Figure 8–5
- Any annotations that are associated with the Problem

Figure 8–4   Details of a Problem

ORACLE Enterprise Manager 11g Ops Center

Problem Details

Resource:	xvm-t2000-9 : OperatingSystem
Model:	-
Serial Number:	-
Contract Status:	-
Problem ID:	1080
Problem Detection Date:	10/22/2010 11:23:47 MDT
Severity:	Warning
State:	Unassigned
Owner:	-
Assign Date:	-
Problem Description:	System Load 5-Minute Average Per CPU has reached 1.06 on xvm-t2000-9
URL:	-
Ops Center Server:	xvm-x4150-18
Problem Link:	-

This figure shows detailed information about the first problem in the summary.

****

Figure 8–5   Alerts History of a Problem
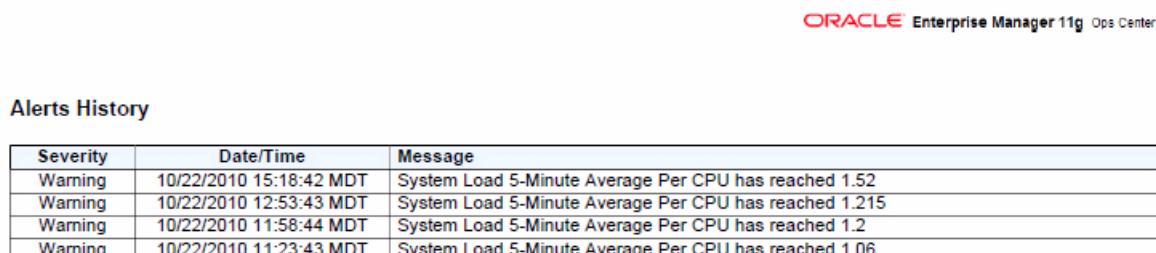
When you create a report, you can save the report as a template, or you can generate the report. After a report is created, you can view the report, re-run the report to get updated information, or save it as a template.

To Create a Problem Summary Report

1.   Expand Reports in the Navigation pane.

2.   Click on Problem Reports, and then click Problem Summary Report.

3.   Click Create Problem Summary Report in the Actions pane.

4.   Define the report parameters by providing the following information, then click Next:

   ■   Name and description for the report.

   ■   Choose the targets, either Select All Assets or Select Specific Assets.

   ■   If you do not want to the report to run on a schedule, click the Create Schedule check box to clear it.

   ■   Choose the output format, either CSV or PDF.

5.   Define the problem parameters by providing the following information, then click Next:

   ■   To create a historical report of all problems and the date that each problem was detected, select All Creation Dates.

   ■   To create a summary report for problems 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 be included 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-separated list for multiple criteria. For example, FileSystemUsage or FileSystemUsage,SwapUsage.

6.   If you selected Schedule in Step 4, the Report Schedule page is displayed. Select a schedule option, then click Next:

   ■   Now – Runs the report immediately.

   ■   Start Date – 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 on a recurring basis. Select the Start, End Time and the Number of hours between runs.

7. Review the Summary.

   ■ To save the report as a template, click Save and Close. Templates are listed in the Report Templates section of the center pane.

   ■ To generate the report without saving it as a template, click Run and Close. Reports are listed in the Report Results section of the center pane.

8. 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.

To Create a Problem Detail Report

1. Expand Reports in the Navigation pane.

2. Click on Problem Reports, and then click Problem Detail Report.

3. Click Create Problem Detail Report in the Actions pane.

4. Define the report parameters by providing the following information, then click Next:

   ■ Name and description for the report.

   ■ Choose the output format, either CSV or PDF.

   ■ Choose the targets by highlighting one or more problems in the list. Use Ctrl+Enter to select multiple problems. The search field enables you to search for a specific problem. Click a table heading to change the sort order, or remove a column from the table.

5. Review the Summary.

   ■ To save the report as a template, click Save and Close. Templates are listed in the Report Templates section of the center pane.

   ■ To run the report without saving it as a template, click Run and Close. Reports are listed in the Report Results section of the center pane.

6. To view the report, select the report and then click one of the icons to choose the format: View interactively, View CSV, or View PDF.

9

Problem Management

Problem management in Enterprise Manager Ops Center consists of several components that are designed to work together to simplify managing problems for a large number of assets. The components include monitoring rules, suggested actions, and methods for automating problem 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 monitoring is enabled, it generates alerts, which connect to both the problem management and notification features.

How Alerts and Problems Work

Each time an asset's attribute that is being monitored does not meet its monitoring rule, a new alert is generated and is reported as a problem in the Message Center. If a value for an attribute exceeds its monitoring rule and then later meets the rule, the alert is cleared automatically. If the attribute does not meet the rule in the next monitoring period, a new alert is generated.

A problem reported in the Message Center is not cleared automatically in response to a change in the attribute's value. To remove a problem from the Message Center, you must close the problem or take no action on it for seven (7) days. See Closing a Problem for more information about manually closing a problem.

If an attribute's value is moving in and out of its monitoring rule's parameters, alerts are generated and cleared continuously. However, a new problem is only generated if the original problem is closed. If a problem is not yet closed, the new alerts are aggregated into the existing problem.

Problems with no activity for seven (7) days are closed automatically and do not appear in the Message Center or in the Asset view. If any action istaken on a problem, such as adding an annotation, the counter is reset.

When a problem is closed, its status changes to Closed, the problem is deleted from the list of active problems, and the problem is no longer displayed in the UI. You can retrieve information about a closed problem for 60 days by using the public API. After 60 days, closed problems 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:** If the monitoring condition is still true after the problem is closed, a new alert is raised and a new problem is created.

Problem Severity Badges

When an asset has a problem, the severity badge appears next to the asset icon in the Asset hierarchy. If it is the highest severity problem in the membership of a group, the severity badge also appears next to the parent assets. In Figure 9–1, the OS for xvmbrm-t5140-2 has a critical problem. The critical problem 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 9–1   Critcal Problem Badge
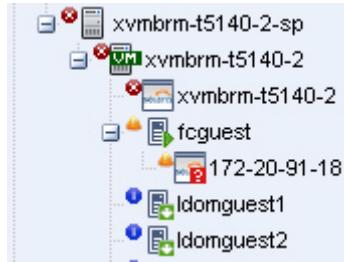
If the problem is the only critical problem, the badge is removed when the problem is acknowledged, marked repaired, or closed. If open problems are still present, the next highest severity badge is displayed. For example, if both a Critical and a Warning problem is detected and the Critical problem is acknowledged, the Critical badge is replaced with the Warning badge because that is now the highest level unacknowledged problem.

Figure 9–2   Warning Problem Badge
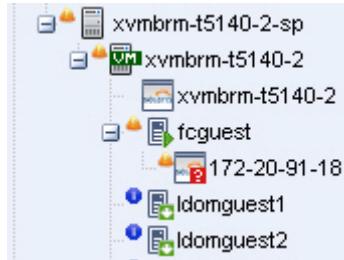
See the *Enterprise Manager Ops Center Advanced User's Guide* for more information about creating and maintaining monitoring rules, profiles, plans, and the problems knowledge base.

Methods of Problem Management

Enterprise Manager Ops Center uses an help desk approach to managing problems. The following are the key tools available for taking action on a problem:

- Message Center – View the status of problems and assign problems.

- 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 a problem.

- Operational Plans – Deploy a shell script against a specific asset, or asset sub-type to automate problem resolution.

- Problems Knowledge Base – Collect comments and suggested actions for known issues for future use.

If you want to receive e-mail or pager notification each time a problem is reported in the Message Center, you can create notification rules to send a message advising you of a new critical or warning problem.

The Message Center contains a detailed list of unassigned problems, problems assigned to you, and problems assigned to other users.

You can manage problems from either the Message Center or from the Asset view. You can view and add comments and annotations, take action on a problem, and closeproblems.

- The Message Center provides a list of all problems. Select a problem to see its details and activity.

- From the Asset tree in the Navigation pane, select the asset and then click the Problems tab to see a list of problems for that asset.

Any user that has the Manage or Admin role for the asset can take action on the problem. The user to which the problem is assigned must also have the Manage or Admin role. If the icon is not active, the user does not have the appropriate role.

See User and Role Management for more information about user roles and authorizations

Viewing Unresolved Problems

You can display a history of unresolved problems for a specific asset, for a user, or unassigned problems. Problems fall into one of the following severity levels:

- Critical

- Warning

- Informational, or Info

To View Unresolved Problems

1. Select the view:

    To view unresolved problems for a specific asset, click the asset in the Navigation pane, then click the Problems tab in the center pane.

    To view unresolved problems from the Message Center, click one of the following:

    - Unassigned Problems

    - My Problems

    - Problems Assigned to Others

        The number of unresolved problems for an asset appears in a bar chart and in a summary by severity. All Unresolved problems appear in a table.

2. (Optional) To view details of a problem, hover or click the problem in the Unresolved Problems table.

3. (Optional) To view the alerts that make up the problem, click the problem in the Unresolved Problems table, then click the Alerts icon in the center pane.

Viewing Problem Details

A problem consists of one or more alerts. You can view problem details, including the individual alerts that are part of the problem. The following problem details are available:

- How long the problem has been open, or the duration

- When the problem was assigned

- The number of suggested actions for the problem

- Who is assigned to the problem

- Which resource is affected, or the cause

- A description of the problem

To View Problem Details

1. Select Assets in the Navigation pane.

2. Select an asset that has a problem badge next to the icon. The Dashboard page is displayed with the status of the asset.

3. Click the Problems tab.

4. Hover over the problem to display the problem details.

5. To display the alerts that are associated with the problem, click the Alerts sub-tab or click the Alerts icon in the center pane. The alerts that make up the problem are displayed, including the current and highest alert status, and the alert history.

Assigning a Problem

You can assign a problem to a user who has Manage or Admin role for the asset.

Assigning a problem might affect the asset's Problem severity badge. If a problem was previously acknowledged or marked as being repaired, its severity was not propagated up to antecedent assets in the navigation pane. After assigning a problem (to a user or to no one), the severity is propagated up again to antecedent assets in the navigation pane.

To Assign a Problem

1. To display a problem from the Message Center, click Message Center, then click Unassigned Problems in the navigation pane.

   To display a problem from the asset view, click the asset in the Navigation pane, then click the Problems tab.

2. Select one or more problems in the center pane, then click the Assign Problem(s) 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 move an assigned problem back to the Unassigned Problems queue, select No One from the list.

4. (Optional) Add a note in the text field.

5. Click Assign Problems.

Acknowledging Problems

Acknowledging a problem indicates that you are investigating the issue. You can acknowledge a problem if you have the Admin or Manage role for the asset on which the problem is identified.

When an asset has a problem, a severity badge appears next to the asset in the Asset hierarchy. If it is the highest severity problem in the membership, it also appears next to the parent assets. In Figure 9–3, the OS for xvmbrm-t5140-2 has a critical problem. The critical problem 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 9–3   Critcal Problem Badge on Asset and Parent Assets
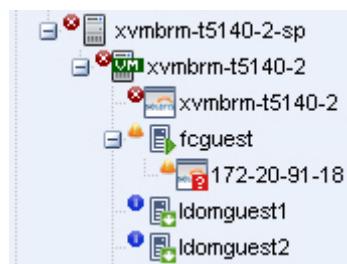

If the problem was previously in an Unassigned state or was assigned to someone else, the severity was taken into account in the computation of the highest severity to propagate up to antecedent assets in the navigation pane. When you acknowledge a problem, it is moved into your queue in the Message Center and the severity is no longer propagated up to antecedent assets in the navigation pane.

When you acknowledge the Critical problem, the badge is replaced with the Warning badge because that is now the highest level unacknowledged problem.

Figure 9–4   Effect of Acknowledging a Critical Problem
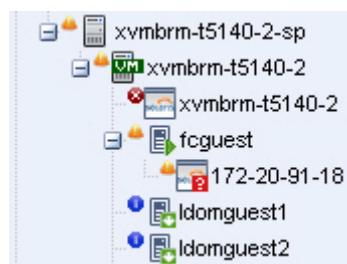
 To Acknowledge Problems

1. From the Message Center, click Message Center, then click one of the following:

   ■   Unassigned Problems

   ■   My Problems

   ■   Problems Assigned to Others

   OR

   From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Select one or more problems, then click the Acknowledge Problems(s) icon in the center pane.

Adding an Annotation

Annotations are defined by the asset type. The annotation can be a comment, a suggested action, or you can refer to an operational profile. Any user can add an annotation to a Problem. To add an entry to the Problems Knowledge Base requires Enterprise Manager Ops Center Admin permissions

See Adding An Annotation To The Problems KB for information about annotating the Problems Knowledge Base.

To Add an Annotation

1. From the Message Center, click Message Center, then click one of the following:

   - Unassigned Problems

   - My Problems

   - Problems Assigned to Others

   OR

   From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Select the problem, 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 Problem is open.

5. The Synopsis field is completed based on the annotation type. Edit the synopsis, as needed. There is no character limit in the UI, but the API allows for 80 characters.

   > **Note:** If 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 Problems Knowledge Base and include the annotation for every problem of this type and severity, click the check box.

   > **Note:** You must have the Enterprise Manager Ops Center Admin role to complete this operation.

8. Click Save and Execute or click Save.

Displaying Annotations

Annotations are automated operations with associated scripts, suggested fixes or actions, or text-only comments. You can associate annotations with a problem instance or an asset type. The Problems Knowledge Base contains your annotations, by asset type, and stores the information on the Enterprise Controller. This annotation can be viewed by browsing the Problems Knowledge Base.

Example 9–1   Example of Using Annotations

```
The CPU usage on a Sun Fire x4150 host is exceeded and a problem is generated. The
problem is assigned to Lee. Lee is concerned because these systems are often used
to host Oracle Solaris Zones. Lee adds the following comment to the problem: "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 Problems Knowledge
Base and is displayed the next time CPU usage is exceeded on a global zone asset
type.
```

To Display Annotations for a Problem

1. From the Message Center, click Message Center, then click one of the following:

   ■   Unassigned Problems

   ■   My Problems

   ■   Problems Assigned to Others

   OR

   From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Click the problem in the center pane.

3. Click the View Annotations icon.

To Display Annotations for an Asset Type Using the Problems Knowledge Base

1. Click Plan Management.

2. Expand Problems Knowledge Base in the Navigation pane, then select the asset type. The Annotations associated with the asset type are displayed 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 a problem as fixed, or when you close a problem. You can also use annotations to build a Problems Knowledge Base that contains a mixture of comments, suggested actions, and automated actions. To add a comment, see Adding an Annotation.

To View Comments

1. From the Message Center, click Message Center, then click one of the following:

   ■   Unassigned Problems

   ■   My Problems

- Problems Assigned to Others

OR

From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Click the problem in the center pane, then click the View Comments icon.

Taking Action on a Problem

If you have the Manage or Administration role for an asset that has an Open Problem, you can correct some problems by using a script or command. In these situations, you might want to associate an automated action with some known issues. For other problems, you might want to review the issue before deciding on the appropriate action.

The problem management functionality provides varying levels of control. You can build the Problems 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 Problem or use an operational plan to correct the problem. If 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 Problem.

See Problems Knowledge Base to add a suggested action or automated action in the Problems Knowledge Base. See Operational Plans for information about creating an operational profile.

To Take Action on a Problem

1. From the Message Center, click one of the following:

   - Unassigned Problems

   - My Problems

   - Problems Assigned to Others

   OR

   From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Select the problem.

3. Click the Take Action(s) on a Problem icon in the center pane.

4. Select the action that you want to perform:

   - If the Problems Knowledge Base has provided a suggested action for the problem, select Execute the Selected Suggested Action option and then select the action from the table.

   - If 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, type 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 problem 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.

Using Maintenance Mode

Maintenance mode is designed to disable assets from generating problems temporarily. This mode is useful when you plan to power off a hardware asset or reconfigure a system manually, and you do not want problems to appear in the user interface.

> **Note:**   Alerts are still generated when the asset is maintenance mode. View alerts by selecting the Alerts tab, which is a subtab of the Problems tab.

When an asset is placed in maintenance mode, the severity badge of unassigned and assigned problems affecting the asset and its children is not propagated up the asset membership hierarchy in the navigation pane.

When the maintenance operations are completed, use the Remove From Maintenance action to begin monitoring the asset. When the asset is removed from maintenance mode, the severity badge appears in the asset membership hierarchy in the navigation pane.

To Use Maintenance Mode

1. To remove an asset from monitoring, highlight the asset in the Navigation pane, then click Place in Maintenance in the Actions pane.

2. To begin monitoring the asset, highlight the asset in the Navigation pane, then click Remove From Maintenance in the Actions pane.

Marking a Problem Repaired

The software cannot determine if a problem is repaired. However, you can open a known problem and manually add a note with the repair details and mark the problem as repaired. You must have the Manage or Admin role for the asset to perform this task.

If the problem 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 problem as repaired, its severity badge will not be displayed in assets list in the navigation pane.

To Mark a Problem Repaired

1. From the Message Center, click Message Center, then click one of the following:

    ■   Unassigned Problems

- My Problems

- Problems Assigned to Others

OR

From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Select one or more problems, then click the Mark Problem(s) as Repaired icon in the center pane.

3. (Optional) Select the problem, then add a Note.

4. Click Tag Problems as Being Repaired.

Closing a Problem

Closing a problem changes the state to Closed, removes the problem from the list of active problems, and no longer displays it in the UI. You can retrieve information about a closed problem for 60 days by using the public API. After 60 days closed problems are permanently deleted.

To edit the time limit, you must edit the value in the public API. You can also disable the feature by setting the value for the number of days to 0.

> **Note:** Problems with no activity for seven (7) days are closed automatically by Ops Center, and do not appear in the UI. You can edit this value in the public API.

After the problem is closed, its severity badge is not displayed in the asset hierarchy.

To Close a Problem

1. From the Message Center, click Message Center, then click one of the following:

- Unassigned Problems

- My Problems

- Problems Assigned to Others

OR

From the asset view, click the asset in the Assets section of the Navigation pane, then click the Problems tab.

2. Select one or more problems, then click the Close Problem(s) icon in the center pane.

3. (Optional) Select the problem, then add a Note.

4. (Optional) To temporarily disable the monitoring rule that identified the problem, click the Action check box, then define when the monitors will be enabled.

5. Click Close Problems.

10

Virtual Pools

A virtual pool is a group of one or more virtual hosts with the same processor architecture that have access to the same virtual and physical networks and storage resources. Virtual pools provide load balancing, high availability capabilities, and sharing of some resources for all members of the pool.

A virtual host refers to Oracle VM Servers that are managed through Enterprise Manager Ops Center. A guest in a virtual pool refers to a logical domain running on an Oracle VM Server.

Virtual pools are resource pools of virtual hosts that share compatible chip architecture, which facilitates actions such as moving guests between hosts. You can also apply resource configurations and policies to them. The policies that you establish for a virtual pool manage many of the CPU utilization and resource balancing functions. Operations to the virtual pool are delegated to the individual virtual hosts in the virtual pool.

To manage the guests within a virtual pool, you can perform warm and live guest migration and you can balance all of the guests' load among the members of the virtual 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. If a virtual host system shuts down for any reason, such as a hardware failure, you can start the guests on another virtual host in the same pool. All guests in the virtual pool can access the images contained in the virtual pool's storage library.

Each virtual host contains a hypervisor and its local resources and network connections. Virtual hosts in a virtual pool share network and storage libraries and several virtual pools can share the same networks and storage resources.

 You can move a virtual host to another virtual pool or add a stand-alone or individually managed virtual host to a virtual pool.

> **Note:** virtual host can belong to only one virtual pool at a time.

Requirements for a Virtual Pool

To be added to a virtual pool, a virtual host must meet the following requirements:

- Must have the same chip architecture as the other virtual hosts in the virtual pool to support load balancing and guest migration. Enterprise Manager Ops Center and Oracle VM Server for SPARC support SPARC T-Series chip architectures.

- Must be running at least the Oracle VM Server for SPARC version 1.2 software.

  Must be running on at least the Solaris 10 10/09 OS and must meet specific patch and firmware requirements.

- Must be on the same physical network as other virtual hosts in the virtual pool. The virtual hosts can be on several physical networks through multiple NICs, but must have at least one network in common with the other virtual hosts in the virtual pool. A good practice is to use a dedicated physical migration network.

- The virtual host's guests must have been created using Enterprise Manager Ops Center. Any guests created through the native CLI cannot be managed by Enterprise Manager Ops Center.

- The virtual host's guests must be in a shutdown state. Shutdown guests are no longer associated with the virtual host. After you add the virtual host to the virtual pool, you can associate the guest with any virtual host in the virtual pool.

Policies for a Virtual Pool

When you create a virtual pool, you define guest placement, and auto balancing polices. When you create guests, you define the guest resource consumption, including physical and virtual CPUs. You can edit the policies at any time.

Placement Policy

The Placement Policy determines the preferred virtual host for new guests within the virtual pool and how the virtual pool is balanced. The placement policy is defined when a virtual pool is created.The following are the placement policy options:

- Place the guest on the virtual host with the lowest relative load, based on the lowest memory and CPU utilization. The calculation is based on a combination of the average load during one hour, one day and three weeks.

- Place the guest on the virtual 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.

- Place the guest on the virtual host that is consuming the least power.

To change the placement policy for a virtual pool, see Editing the Configuration of a Virtual Pool.

Example 10–1   Example of Least Allocated Virtual Host

A virtual pool has two virtual hosts with different CPU and memory allocations. Virtual Host A has eight available CPUs and 16 GB of memory. Virtual Host B has four available CPUs and 8 GB of memory.

Virtual 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.

Virtual Host A's total static allocation is four vCPUs and four GB of memory.

Virtual host B has one guest:

- Guest W has three vCPU and 4096 MB of memory.

Virtual Host B's total static allocation is three vCPUs and four GB of memory.

The allocation percentage for Virtual 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 Virtual Host B is: CPU allocation is 3 vCPUs/4 physical CPUs or 75%. Memory allocation is 4 GB/8 GB or 50%

Therefore, Virtual Host A is the less or least allocated of the virtual hosts.

How Power Minimization Policy Works

The power minimization policy can be used only when the virtual host has the capacity to report its power usage. To report power usage, a virtual host must have ILOM version 2.x or 3.x.

This power minimization policy consolidates guests into the optimal set of virtual hosts. The idle virtual hosts are powered off or set to low-power mode on explicit approval.

Automatic Load Balancing Policy

Use the Automatic Load Balancing Policy to schedule load balancing within a virtual 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 hosts in the virtual pool every Saturday at midnight according to the defined placement policy.

> **Note:** The day and time are in the Enterprise Controller's time zone.

If you do not want to balance the virtual pool's load automatically, you can schedule a reminder to rebalance the virtual pool's load manually.

Libraries for a Virtual Pool

Enterprise Manager Ops Center uses storage libraries to store ISO images and guest metadata. Software libraries are also used to provide storage disks for the guest operating systems and for guests' data.

The virtual hosts in a virtual pool can get access to and share any storage and networks associated with the virtual pool. If you add a virtual host to a virtual pool, the libraries associated with that virtual host become available to all the other virtual hosts in the virtual pool.

Networks for a Virtual Pool

Networks enable virtual hosts and guests to communicate with each other or through the Internet. Each network has its own values for parameters such as an IP address, subnet, time server, NIS, and DHCP information. Guests in the network communicate with each other or with the Internet through these hosts.

A managed network depends on the physical network interface card (PNIC) that is available to the virtual host. You can create one managed network for each physical network interface card. If one host has two PNICs, it is a good practice to create two managed networks: a management network and a data network. This configuration lets you place all the guests on the data network, and reserving the management network for access to internal resources of the data center.

 When you assign a network to a virtual pool, the network becomes accessible to all virtual hosts in the pool. All virtual hosts in a virtual pool belong to the same set of networks. At least one network must be assigned to a virtual pool to ensure that when

you migrate a guest from one virtual host to another virtual host within the pool, the guest can still access the network.

Guests communicate with the networks through their virtual host. When you add a virtual host to a virtual pool, the virtual host is also configured for all the associated networks for the virtual pool. In this way, the virtual host can access all of the networks defined for the pool and be an active member of the pool.

You can associate a network multiple times on the same Oracle VM Server. The multiple connections are made through different switches on the Oracle VM Server. For each network connection, a virtual switch is created for the network that connects to only one NIC at a time. The naming pattern for the switches relies on the Oracle VM Server. If an Oracle VM Server is connected to a network, for example network 1.1.1.0/24, the name of the virtual switches is 1.1.1.0_24, 1.1.1.0_24_1, 1.1.1.0_24_2 and 1.1.1.0_24_3. When a network connection is made to the Oracle VM Server, the virtual switch created is incremented.

You can assign a network multiple times to a virtual pool. The network is deployed the same number of times on each virtual host of the virtual pool. Also, the switches that are created for each network connection are same on the virtual hosts. This ensures that the logical domain can be migrated across the virtual hosts in the virtual pool. To migrate a logical domain, the switch name must be the same on the source and target virtual hosts..

See Managing Networks of a Virtual Pool for detailed information about networks in a virtual pool.

Viewing the Configuration of a Virtual Pool

The Virtual Pool Dashboard displays the characteristics and uses for a virtual pool.

To Display Virtual Pool Details

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the All Assets drop-down list.

3. Select a virtual pool. The Dashboard page for the virtual pool is displayed in the center pane with the following tabs:

   - Summary – Details about the virtual pool, including the policies, attached virtual hosts, guests, and resource load.

   - Associated Libraries – List of libraries for the virtual pool, their sizes, the guests that use each library, and the contents of each library. Library contents include image type, modification date, and size.

   - Available Networks – List of virtual networks for the virtual pool, including network name, IP address, netmask, and description.

   - Charts – Enables you to create CPU, memory, and network usage charts.

   - Logs – Provides a record of recent activity.

Creating a Virtual Pool

You can create a new virtual pool with one or more virtual hosts that share the physical and virtual networks and storage resources.

To add a virtual host to an existing virtual pool, see Adding an Oracle VM Server to a Virtual Pool.

Before You Begin

- Review Requirements for a Virtual Pool.

- Decide on the placement policy for this virtual host's guests.

- Decide on the balancing policy for this virtual host, if any.

- Identify the virtual host to be added to the new virtual pool:

    - Be a managed virtual host or a pooled host.

    - Have no guests.

    - Be in a healthy state.

    - Cannot be the only virtual host in another virtual pool.

    - Have a compatible chip architecture with the other virtual hosts that you wan to add to the virtual pool.

    - Be compatible with the CPU frequency of other Oracle VM Server

To Create a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the drop-down list.

3. Click Create Virtual Pool in the Actions pane.

4. Provide a name and description for the new virtual pool, and assign tags. Tags are keywords that you can use to help identify the virtual pool. Click Next.

5. Select how new guests are allocated to a virtual host in the virtual pool, according to one of the following policies:

    - Lowest relative load

    - Least allocated CPU and memory resources

    - Least power consumption

6. Select how loads are balanced within the virtual pool according to one of the following policies:

    - Manual balancing, the default

    - Automatic balancing, which reviews the load on the host for the interval and day that you request and migrates the guests, as needed, to balance the load. You can require administrator approval before the guests are moved.

7. Assign one or more networks to the virtual pool from the list of available networks.

8. Specify the number of connections for the selected network. The number of connections is also the number of virtual switches.

9. To allow guests to move between virtual hosts, click the Migration check box to designate the network as a Migration Network. Migration networks enable you to isolate guest migration traffic to a subset of the Networks in the pool. You can designate more than one network as a migration network.

10. Select an Oracle VM Server from the list to define the architecture of the virtual pool. Click Next.

11. Select one or more Oracle VM Servers from the list of virtual hosts that match the chip architecture to add to virtual pool. Click Next.

> **Note:** If the Oracle VM Server you select is already in another virtual pool, the server moves from its current virtual pool to the new virtual pool. An Oracle VM Server can belong to only one virtual pool.

12. For each network connection requested in Step8, specify the NIC and IP address to create the virtual switch.

13. For each IP address, select either Use Static IP and provide the IP address or select Assign by DHCP to obtain the IP address from the DHCP server as needed. Click Next.

14. From the list of available storage libraries, select ast least one NAS storage library to associate with the virtual pool. Click Next.

15. Review the summary and details of the new virtual pool. Click Finish to submit a job that creates the new virtual pool.

Editing the Configuration of a Virtual Pool

You can edit the following virtual pool attributes:

- Name
- Description
- Tags
- Placement Policy
- Auto Balancing Policy
- Load Balancing Schedule

Before you change a policy for a virtual pool, review the relationship between the physical and virtual CPUs, that is, its placement policy. Each virtual pool has a placement policy that defines how guests are placed:

- Place the guest on the virtual host with the lowest relative load.
- Place the guest on the virtual host with the lowest allocated CPU and memory.
- Place the guest on the virtual host which is consuming the least power.

Auto-balancing changes the allocation of resources defined by the placement policy. The load is balanced by migrating a guest to a different virtual host in the same virtual pool

You can choose to balance the load on a schedule that you define, or you can balance the load manually. The Auto Balance policy is set for each virtual pool.

To Edit Virtual Pool Attributes

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the All Assets drop-down list.

3. Select a virtual pool. The Dashboard page for the virtual pool is displayed in the center pane.

4. Click Edit Virtual Pool Attributes in the Actions pane.

5. To change the name or description of the virtual pool or to change its tags, enter new text in the text box.

6. To change the Placement Policy or Auto Balancing Policy, select the new policy. Select the date interval, time interval, and time of day to change the auto balance or to send a reminder to balance the load manually.

7. Click Save

Adding an Oracle VM Server to a Virtual Pool

An Oracle VM Server can belong to only one virtual pool. You can add an Oracle VM Server to a virtual pool at any time if it meets the following requirements:

- A stand-alone Oracle VM Server, that is, an individually managed virtual host

- Does not have running or suspended logical domains.

- In a healthy state

- Cannot be the only Oracle VM Server in another virtual pool

- Has a compatible chip architecture with the other virtual hosts in the virtual pool

- Compatible with the CPU frequency of other Oracle VM Servers in the virtual pool.

> **Note:** If the Oracle VM Server you add to the virtual pool has network connections that are not available to the virtual pool, those connections are broken and no longer available to the Oracle VM Server.

To Add an Oracle VM Server to a Virtual Pool

1. Click Virtual Pools in the Assets view of the Navigation pane.

2. Select the Virtual Pool to which you want to add the Oracle VM Server.

3. Click Add Oracle VM Server in the Actions pane.

4. Select one Oracle VM Server from the list of eligible Oracle VM Servers in the table. Click Next.

5. Review the summary of the virtual pool and the new virtual host. Click Finish to submit a job that adds the Oracle VM Server to the virtual pool.

Removing an Oracle VM Server From a Virtual Pool

You cannot remove an Oracle VM Server from a virtual pool if it is the only virtual host in the virtual pool or if it has any running or suspended guests.

When you remove an Oracle VM Server from a virtual pool, the server becomes an individually managed server. The virtual pool's libraries and networks are no longer available to the server. Instead, the Oracle VM Server uses the default local library and management network.

Before You Begin

- Shut down all the guests associated with a virtul pool.

- If the Oracle VM Server's state is unknown, you can force the removal of the server from the virtual pool. A forced removal shuts down and disassociates all guests from the server.

To Remove an Oracle VM Server From a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool in the drop-down list.

3. Select the virtual pool.

4. Select the Summary tab in the center pane. The list of Oracle VM Servers in the virtual pool are listed.

5. Select an Oracle VM Server from the list.

6. Click the Remove from Virtual Pool icon.

7. In the window, click Remove from Virtual Pool to start a job that removes the Oracle VM Server from the virtual pool . If the host state is unknown, click Force Remove Virtual Host in the dialog box.

Managing Storage Libraries of a Virtual Pool

When you create a new NAS storage library, you can also associate it with a virtual pool. Use the following procedure to associate an existing storage library with a virtual pool.

To Associate a Storage Library With a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the drop-down list. All the virtual pools are listed.

3. Select a virtual pool.

4. Click Associate Libraries from the Actions pane. The Associate Library window is displayed.

5. Select the libraries from the list.

6. Click Associate Library. The selected libraries are associated with the virtual pool.

To Disassociate a Library from Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the drop-down list.

3. Select the virtual pool from the list.

4. Select the Libraries tab in the center pane. The libraries that are associated with the virtual pool are listed.

5. Select the library that you want to disassociate from the virtual pool. Check the usage and contents of the library before disassociating it.

6. Click the Disassociate Library icon. The Disassociate Library window is displayed.

7. Click Disassociate Library to confirm unmounting the library.

Managing Networks of a Virtual Pool

You can assign networks to virtual pool and define the number of connections for each network. For each network connection, a virtual switch is created. You need to define the NIC and the IP address for each network connection.

To Assign a Network to a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the drop-down list.

3. Select the virtual pool.

4. Select Attach Network from the Actions pane. The Assign Network wizard is displayed.

5. Select a network from the list and enter the number of connections for the network. You can add more than one network. Click Next.

6. Select a NIC from the list and specify how it connects to the network. The list of NICs includes NICs that are not used by other networks and NICs that are bound to other networks, identified by the network's VLAN tag. You can select a NIC bound to a network if the network has a different VLAN tag. Select NICs so that all the virtual hosts have the same virtual switch name. The attach network job fails if the switch names.

7. For each network connection, choose the Static IP Adress option and provide an IP address, or select the Assign to DHCP option to allocate the IP address as needed. Enter the IP address in the DHCP Client ID field. Click Next.

8. Review the selections you have maded. Click Finish to assign the networks. The networks are assigned to the virtual pool and are available to all the guests of all the virtual hosts in the pool.

To Unbind a Network From a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool from the drop-down list..

3. Select the virtual pool.

4. Select the Networks tab in the center pane. The networks for the virtual pool are listed.

5. Select the network that you want unbind.

6. Click the Unbind Network From Virtual Pool icon.

7. Click Unbind to confirm disassociating the network.

Deleting a Virtual Pool

Deleting a virtual pool removesthe association among the virtual hosts and removes all pool-specific library and network connections. The virtual hosts in the virtual pool become individual virtul servers relying on default network connections..

To Delete a Virtual Pool

1. Select Assets from the Navigation pane.

2. Select Virtual Pool in thedrop-down list.

3. Select the virtual pool.

4. Click Delete Virtual Pool in the Actions pane.

5. Click Delete Virtual Pool to confirm action.

11

Oracle VM Server for SPARC

Oracle Solaris VM Server for SPARC technology, formerly known as Logical Domains or LDoms, 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.

With Enterprise Manager Ops Center, you can monitor and manage all of the virtual machines from a single browser user interface. You can use virtual pool technology to maximize capacity without overloading a server. To balance the load, you can move a virtual machine from one physical system to another system either automatically or rmanually, from the user interface.

In this release you can perform the following tasks:

- Provision the system to create Control Domain

- Reboot, and shut down the Oracle VM Server

- Monitor and interpret the performance of Oracle VM Server

- Create and provision logical domains

- Manage logical domains, including editing, migrating, starting, rebooting, and shutting down domains

- Monitor and interpret the performance of logical domain

About Oracle VM Server for SPARC

You use hardware virtualization to create multiple virtual machines on a single piece of physical hardware. Unlike Oracle Solaris Zones, where the operating system is the same for each non-global zone, virtual machines can run on full instances of different operating systems, or different versions of the same operating system. These instances are called logical domains.

Oracle VM Server for SPARC technology, formerly known as Logical Domains or LDoms, 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. This technology allows you to allocate a system's various resources, 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 a boot environment, 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.

Oracle VM Server for SPARC functionality is especially useful in the following scenarios:

- To combine several small UNIX and Linux servers, you can create logical domains to host the applications running on those servers.

- To use different OS kernels in the same system, you can create logical domains to host different OS kernels.

- To provide maximum isolation and security, logical domains provide separate OS and hardware resources.

- To replace multiple legacy servers, set up a one-to-one mapping from legacy servers to the logical domains.

See Additional Resources for links to the Oracle VM Server for SPARC product documentation.

Enterprise Manager Ops Center and Oracle VM Server

Enterprise Manager Ops Center can manage logical domains that are created through the its user interface. When you use Enterprise Manager Ops Center to provision the Oracle VM Server for SPARC software, Enterprise Manager Ops Center installs the Oracle Solaris OS, the control domain, and an agent on the target system. This provisioning action removes existing virtualization software, including any logical domains.

> **Note:** You can discover logical domains that were created through the native CLI and display them in the Assets tree. However, you cannot use Enterprise Manager Ops Center to perform any actions on those domains.

Requirements for Oracle VM Server for SPARC

Enterprise Manager Ops Center supports Oracle VM Server for SPARC 1.2, 1.3, 2.0, and 2.1 version. The Oracle VM Server must meet specific hardware, OS, and firmware requirements.

> **Note:** The OS in the control domain must have the default locale set to C. The control domain provisioning is supported on Oracle SPARC or x86 Proxy Controllers, but not on Linux Proxy Controllers.

> **Note:** 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.

Architecture of Oracle VM Server for SPARC

Figure 11–1 shows a sample architecture that contains a control domain and three logical domains. The system administrator created a logical domain for different areas of company. The first domain deploys a financial application; the second domain deploys a database application; the third domain deploys a web application. CPU threads and memory resources are assigned to each domains, depending on the requirements.

Figure 11–1   Sample Architecture for Oracle VM Server for SPARC
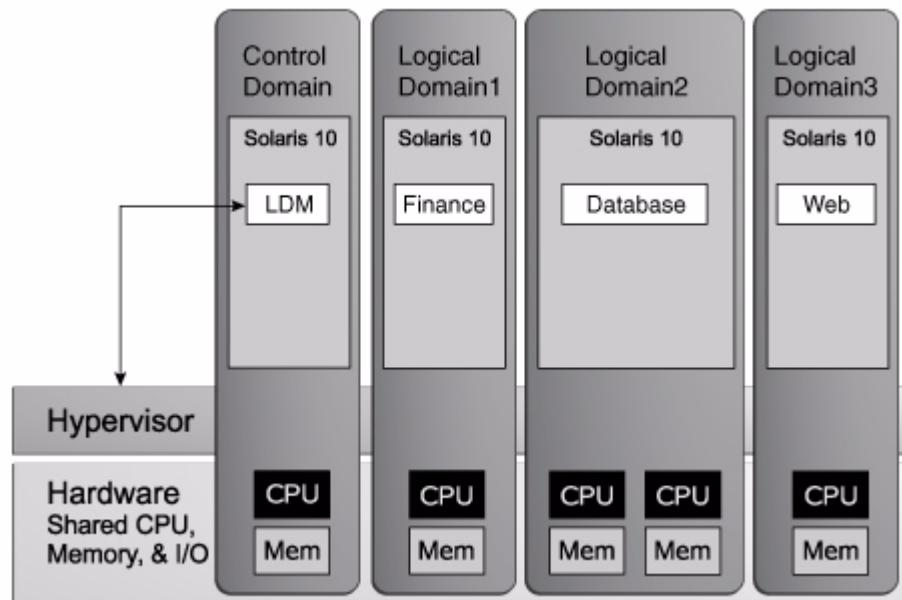
The control domain handles the Oracle VM Server environment and creates the logical domains. The hypervisor is a firmware layer on the FLash PROM of the motherboard, operating between the OS and the hardware. The hypervisor provides a set of support functions to the OS, so that the OS does not need to call hardware functions.

Roles for Domains

Oracle VM Server for SPARC software's control domain, or virtual host, createw and manages the logical domains and allocates virtual resources to other domains. In addition to control domains and logical domains, the software can create domains for specific roles, such as a service domain and an IO domain.

An individual domain can have one or more of these roles, such as combining the functions of an I/O and service domain.   In Enterprise Manager Ops Center, the control, service, and I/O domains are in the same domain.

- **Control Domain**

  The Oracle VM Server virtual host runs in the control domain. There can be only one control domain for each server. The first domain created after installing the Oracle VM Server software is the control domain and is named the primary domain. The control domain contains the SUNWldm packages, including the Oracle VM Server host application and the host daemon (ldmd) process required for managing logical domains. If you have an OS and other software running on a

server and you install the Oracle VM Server software, that server becomes your control domain.

> **Note:** Zones are not supported in the control domain. You must not run user applications in the control domain and you cannot migrate, delete, or shut down a control domain.

- Logical Domain

  The logical domain is a complete virtualized environment that has no ownership of physical I/O or virtual devices. This domain is managed by the control domain and uses services from the I/O and service domains. The logical domain must run an OS that supports both the sun4v platform and the virtual devices presented by the hypervisor.

- Service Domain

  The service domain provides virtual device services such as a virtual network switch, a virtual console concentrator, or a virtual disk service, to other domains.

- I/O Domain

  The I/O domain has direct access to physical I/O devices, such as a network card in a PCI Express controller. This domain shares the devices with others in the form of virtual devices when the I/O domain is also the control domain. The number of I/O domains that you can have depends on your platform architecture. For example, if you are using an UltraSPARC T1 processor, you can have a maximum of two I/O domains, with one of them being the control domain.

Configuration for an Oracle VM Server for SPARC

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 Creating an OS Profile for Oracle VM Server.

> **Note:** Enterprise Manager Ops Center does not support use of the LDoms configuration manager in LDoms 1.2. Do not attempt to reconfigure the control domain and logical domains using the native CLI of Logical Domain software.

The recommended minimum configurations for the control domain are described in the following sections.

CPU Threads

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.

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 contains at least one thread from the parent core. Crypto units cannot be split as CPU threads are. For example, you have assigned the crypto unit for the first CPU core to the control domain. If 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 if 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.

Example 11–1   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::

- 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 th econtrol domain's RAM to 8 GB.

Additional Configurations

The following configuration for the control domain is performed during installation of Oracle VM Server for SPARC and does not need to be included in the profile:

- A virtual network switch for each subnet named *subnet-address_prefix-length*, for example, 10.17.7.0_24

- A virtual disk server named *primary-vds0*

- A virtual console concentrator is named *primary-vcc0*

Creating an OS Profile for Oracle VM Server

Create a profile to install the Oracle VM Server for SPARC software, the OS, and to set the parameters for the control domain. Use this profile in a deployment plan that provisions a SPARC server's service processor.

> **Note:** The version of Oracle VM Server for SPARC to be installed depends on the target systems. After the provisioning job starts, an information problem displays the version that is installed on the target server.

See Hardware And Provisioning Profiles and Deployment Plans for more information about creating profiles and plans, and applying the plans.

To Create an OS Profile for Oracle VM Server

1. Select Plan Management from the Navigation pane.

2. Select OS Provisioning in the Profiles and Policies tree. The OS Provisioning page is displayed in the center pane.

3. Select Create Profile in the Actions pane. The Create Profile-OS Provisioning wizard starts.

4. To identify the new profile, enter a name and a description for the profile.

5. To create a deployment plan that containsthe new profile, select the Create a deployment plan for this profile option .

6. For subtype, select Oracle VM Server. Click Next.

7. From the OS Image list, select the OS image that you want to install.

8. Select a Distribution Type. The minimum requirement is End User distribution. Distributions that are lower than End User require additional package dependencies and the provisioning job can fail if the packages are not present.

9. (Optional) To add a script to the profile, select Include Custom Scripts and then click the Add icon. Enter the location of the script, which must be accessible from the Enterprise Controller. Specify if you want to execute the script before or after the provisioning operation.

10. Specify the following parameters and then click Next:

    - Language – Select a Language from the list.

    - TimeZone – 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 console serial port device.

    - Console Baud Rate– To monitor the installation using a serial connection, select the baud rate.

    - NFS4 Domain – Enter the NFS4 domain name that the target system will use. A dynamic NFSv4 domain name enables the name to be derived at run time, based on the naming service configuration. If you prefer, enter a static domain name.

    - Password –Enter the root password for the root user on systems provisioned using this profile. Re-enter the password for confirmation.

11. 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 Enterprise Manager Ops Center cannot control the network boot process remotely on these systems.

12. Select Automatically Manage with Oracle Enterprise Manager Ops Center to install an agent on the system. Click Next.

13. JET modules can be used to provision an OS. The base_config, custom, and flash JET modules are always installed. To specify additional JET modules that you have installed on the Proxy Controller to perform OS provisioning operations, enter a comma-separated list of the JET modules.

   a. Click the Add icon to add JET name-value pairs. The JET parameters customize how this profile provisions the target systems.

   b. Enter the name of the JET parameter that you want to add in the Name field.

   c. Enter the value that you want to assign to the JET parameter in the Value field.

   When you are finished, click Next.

14. Specify the resources that you want to assign to the control domain. The remaining resources are available for the logical domains.

Figure 11–2 Parameters for a Control Domain
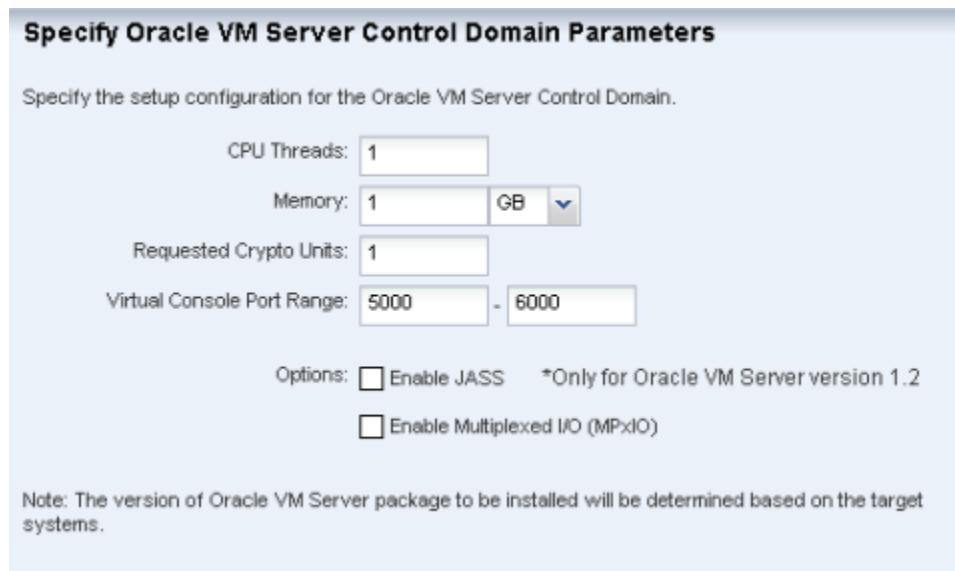
- CPU Threads – Specify the number of CPU threads that you want to assign to the control domain.

- Memory – Specify the amount of memory that you want to assign to the control domain.

- 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.

- Enable JASS – Select this check box to harden the system by installing the SUNWjass package.

   Note: JASS is not supported for Oracle VM Server for SPARC 1.3 or higher versions. This option is disregarded.

- Enable Multiplexed I/O (MPxIO) – Select this check box to enable Fibre Channel connectivity for the control domain. This action enables the Fibre Channel ports on the system that is configured for storage.

Click Next.

15. 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. 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, either ufs, unnamed, 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 (MB) – 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,

16. Specify the name service, domain name and the corresponding name server. You can select the following name service:

- 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. If 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.

- None – Select None when there is no naming service configured.

Click Next.

17. Select one of the network interface options for the target system:

- Use Link Aggregation – Click Next to go to Step 19 and Step 20

- Use an IPMP group – Click Next to go to Step 21

- None – Click Next to go to Step 18

18. Select a DHCP-enabled network interface for the boot interface from the list of all managed networks. Click the Add icon to add more networks.

   a. For each network, select a NIC from the list of available interfaces.

   b. Select the Address Allocation Method for the network except the boot interface. If you select Use Static IP, you must provide the IP address when you apply the profile so that the IP address can be assigned to the target system after provisioning.

19. To specify and configure the Link Aggregation, enter the following:

    a. Link Aggregation Name – The name of the Link Aggregation is set to "aggr". Add an integer to differentiate this aggegation from other aggregations.

    b. Network – Select a network from the list.

    c. NICs – List the physical interfaces of the selected network that must be configured as a single unit.

    Click Next to configure the link aggregation.

20. Configure the IEEE 802.3ad Link Aggregation with the following parameters:

    a. Load Balancing Policy – Define the policy for outgoing traffic.

    b. Aggregation Mode and Switches –- If the aggregation topology connects through a switch, determine whether the switch supports the Link Aggregation Control Protocol (LACP). If the switch supports LACP, you must configure LACP for the switch and the aggregation. Define one of the modes in which LACP must operate.

    c. MAC Address Policy – Define whether the MAC address of the NICs are fixed.

21. In the Specify IPMP Group step, define the following information:

    ■ IPMP Group Name – Provide a name for the IPMP group.

    ■ Network – Select a network from the list.

    ■ Failure Detection – The Link-based detection is always embedded. If you want to include Probe-based detection, select the Probe option.to specify the IPMP interfaces. Test address is not required for probe-based failure detection.

    Click Next.

22. To specify an IPMP interfaces, define the following information. The data and test addresses are assigned when the profile is applied.

    ■ Specify the interfaces that will be part of the IPMP group.

    ■ Define the interfaces as Failover or Standby.

    ■ Configure additional IP addresses for the interfaces.

    Click Next.

23. Review the Summary of your selections. Click Finish to crete the profile.

See Deployment Plans and Plans for creating a plan with this profile and apply the plan to provision the Oracle VM Server.

Provisioning Oracle VM Server for SPARC

You can install the software on bare-metal or on systems that is already configured with logical domains. When you provision Oracle VM Server through Enterprise Manager Ops Center, any previous configuration is removed from the service processor by resetting the service processor 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 virtual 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 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.

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.

- Configures DHCP and installs the OS.

- 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.

- Depending on the version, installs the SUNWJass package to harden the system.

- Enables the Fibre Channel ports on the system that is configured for storage.

- Installs and configures the agent. Do not install the agent manually on the Oracle VM Server host.

Before You Begin

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 Requirements for Oracle VM Server for SPARC .

2. Discover the target system on which you want to install Oracle VM Server for SPARC software.

3. Configure DHCP services on Proxy Controllers. See CONFIGURING DHCP.

4. Create a profile for provisioning Oracle VM Server for SPARC. See Creating an OS Profile for Oracle VM Server.

5. Create a deployment plan with the profile created for provisioning Oracle VM Server for SPARC.

   See Deployment Plans for information about creating a plan with this profile .

6. (Optional) Adjust the amount of time allowed for the job. The default time is three hours.

   a. In the Enterprise Controller, edit the /var/opt/sun/xvm/satellite.properties file and locate the osp.default_ timeout property.

   ```
   OS Provisioning
   #################################
   # Default Timeout of OSP Job
   osp.default_timeout=180
   ```

   b. Change the value of the property to the desired number of minutes.

   c. Save and close the file. You must restart the Enterprise Controller after changing the value.

7. The Oracle VM Server for SPARC version to be provisioned is selected depending on the type of target server and firmware version. You can define the version to be provisioned explicitly. The version value is read during the provisioning job.

    a. In the Enterprise Controller, edit the
        `/var/opt/sun/xvm/satellite.properties` file.

    b. Add the following entry with one of the valid versions: 1.2, 1.3, 2.0, and 2.1.

        `osp.ldom_version=<version>`

    c. Save and close the file. You do not need to restart the Enterprise Controller.

        To allow Enterprise Controller to determine the proper version of Oracle VM Server, edit the properties file and remove this entry.

To Apply the Profile on the Target

You apply a profile to a target system by applying a deployment plan that contains the profile. The plan can contain other profiles too. You can apply the Oracle VM Server provisioning profile in one of the following ways:

Method 1: Select the Target

1. Select the server that will become the Oracle VM Server host from the Assets sections.

2. Click Install Server option from Deploy in the Actions pane. The Install Server window is displayed. The list of available plans for provisioning OS is also displayed.

3. Select a plan from the list or, to create a new plan, click Create New Plan and use the following procedure.

4. Click Apply Plan. The job to provision Oracle VM Server for SPARC starts.

To create a new deployment plan:

1. Click Create New Plan. The Select Plan Type window is displayed.

2. Select one plan type, either Provision OS or Install Server.

3. Click Create New Plan. The Create Deployment Plan window is displayed.

See Creating A Deployment Plan.

Method 2: Select the Plan

1. Select the deployment plan from the Plan Management section in the Navigation pane.

2. Select Apply Deployment Plan from the Actions pane. The Select Assets window is displayed.

3. Select a server from the list of available targets and click Add to Target List. You can add more than one target to apply the plan.

4. Select whether you want to apply the plan with minimal user interaction or override any profile values. In either case, the Provision OS wizard starts. If you chose to override profile values, select options and edit the parameters for each step in the wizard. If you chose minimal user interaction option, you specify the resources.

5. Specify the network resources for each target.

6. Enter the IP address for each target.

7. Review the summary of the parameters for provisioning the Oracle VM Server.

8. Click Apply the deployment plan on the selected targets.

Initiating Manual Net Boot

A target system requires a manual net boot operation if Enterprise Manager Ops Center cannot perform a remote network boot process. Profiles that provision this type of system contain the Manual Net Boot option. You must initiate the net boot operation when the Oracle VM Server provisioning job is completed.

1. Log in to the service processor of the target system.

2. If the target system is running, enter `halt`.

3. Enter the following command

   ```
   boot net:dhcp - install
   ```

Provisioning the Oracle VM Server Version

The version of Oracle VM Server for SPARC software installed on the target server is determined during the execution of the job, unless you have defined the version explicitly. The version installed depends on the following conditions:

- If the target is a T1 server, the Oracle VM Server for SPARC 1.2 version is installed. The Oracle VM Server SPARC versions 1.3 and 2.0 are not supported for T1 servers.

- The system-firmware versions also determines the version of Oracle VM Server for SPARC that is installed.

  - For Firmware 7.1.x , Oracle VM Server for SPARC 1.2 is installed.

  - For Firmware 7.2.0 to 7.2.6 , Oracle VM Server for SPARC 1.3 is installed.

  - For Firmware 8.0Firmware 7.4.0 on UltraSPARC T2 and UltraSPARC T2 Plus Servers – Oracle VM Server for SPARC 2.1.0 or higher , Oracle VM Server for SPARC 2.0 is installed.

  - Firmware 8.1.0 on SPARC T3 Servers – Oracle VM Server for SPARC 2.1.

- The JASS that is used for hardening the system does not apply to Oracle VM Server for SPARC 1.3 and 2.0 versions. If you have selected the Enable JASS option, the option is skipped when installing Oracle VM Server for SPARC 1.3, 2.0, and 2.1 versions. A notification is sent if this even occurs.

Managing Oracle VM Server

To manage the resources for Oracle VM Server, you edit the resource configuration such as CPU Threads, Crypto Units, and memory. You can associate storage libraries with the Oracle VM Server and monitor the performance of logical domains..

Viewing Information About an Oracle VM Server

When you select an Oracle VM Server from the Assets tree, Enterprise Manager Ops Center displays information about it in the center pane. From the center pane, you can select tabs and actions to manage the Oracle VM Server.

To View Information About an Oracle VM Server

1. Select an Oracle VM Server from the Assets section in the Navigation pane.

2. Click on one of the tabs in the center pane to view details of the Oracle VM Server's configuration.

- Dashboard

    - Summary – The summary of the group to which the Oracle VM Server belongs.

    - Membership Graphs – A graphical representation of the group members and its relationships.

    - Status – A graphical representation of unassigned problems and a table of recent problems.

- Summary

    - Name

    - Description

    - Status

    - CPU Info

    - Available CPU Threads

    - Available Crypto Units

    - Available memory

    - Version of Logical Domains Manager

    - Control Domain specification

    - Resource utilization of Oracle VM Server

    - Logical domains, and their status and specification

    - Virtual pool

    - Reachability

- Libraries

    The Libraries tab displays thestorage libraries that are associated with the Oracle VM Server. When you select a library, the logical domains that are associated with this library are listed under Usage section. You can also view the contents of the library from the Contents section.

    See Creating A Local Library.

    The icons perform the following actions::

    - Associate Libraries – Add an existing storge library to the Oracle VM Server. See Associating Libraries With Oracle VM Server

    - New Local Library – Create a storage library on the Enterprise Controller.

    - Edit Local Library – Change the attributes of an existing local library.

    - Delete Local Library – Remove the local library.

    - Disassociate Library – Remove the connection between the library and the Oracle VM Server.

- Networks

    The Networks tab lists the networks that are connected to the Oracle VM Server. You can attach the Oracle VM Server to more than one network. When you select a

network from the list, the connectivity information and the logical domains using the network are displayed.

The icons perform the following actions::

- Attach Network – Connect the Oracle VM Server to the selected network.

- Unbind Network from Virtual Pool – Disconnect the selected network from Oracle VM Server.

- Modify Physical Connectivity – Edit the connectivity attributes of the selected network.

- Refresh DHCP Connectivity – Re-connect to the DHCP server.

See Managing Networks of Oracle VM Server for detailed procedures.

Note: You cannot modify or delete the management network of the Oracle VM Server.

- Problems

   The Problems tab lists the problems that are reported for the Oracle VM Server. It provides the list of unresolved problems and a graphical representation of the problem composition.

   To see the corresponding alerts for each problem, see Problem Management for more detailed information.

- Monitoring

   The Monitoring tab 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 Monitoring Profiles.

- The Charts tab displays the CPU, memory and network utilization of the Oracle VM Server. The data is collected in five-minute intervals and is displayed graphically. The utilization data is provided for different time intervals. See About Charts for more information about reading charts.

- Jobs

   The Jobs tab lists the current and completed jobs of the Oracle VM Server. See Job Management for more information about managing jobs.)

- Configuration

   The Configuration tab displays information such as remote logging, routing, NFS domain name, and name service information for the Oracle VM Server.

Editing the Attributes of Oracle VM Server

You can modify the attributes of Oracle VM Server such as name, description, tags, CPU Threads or Crypto Units (MAUs), and memory size.

The Oracle VM Server must have Oracle Solaris 10 10/09 or later provisioned to edit the attributes.

To Edit the Attributes of an Oracle VM Server

1. Click Assets and then select the Oracle VM Server in the Navigation pane.

2. Select the Summary tab in the center pane.

3. Select Edit Attributes from the Actions pane.

4. Change the values for any of the followingattributes.

   ■ Name and Descriptiong

   ■ Tags

   ■ CPU Threads

   ■ Memory size

   ■ Crypto Units

5. Click the Save icon to save the changes.

When you modify the memory size, Oracle VM Server is rebooted unless it is running Oracle VM Server for SPARC 2.0 version.

Rebooting an Oracle VM Server Host

An Oracle VM Server can be rebooted regardless of the state of the logical domains. During the reboot, the logical domains might suspend temporarily but keep their current states.

To Reboot an Oracle VM Server Host

1. Select an Oracle VM Server from the Assets section in the Navigation pane.

2. Click Reboot Oracle VM Server Host in the Actions pane.

3. To confirm the action, click Reboot Oracle VM Server in the window. A job is submitted to reboot the Oracle VM Server. . If you were logged into the server console, you must log in again.

Associating Libraries With Oracle VM Server

Oracle VM Server uses storage libraries and software librarie to create logical domains and to maintain their data.

To Associate Libraries With Oracle VM Server

1. Select an Oracle VM Server from the Assets section in the Navigation pane.

2. Select Associates Libraries in the Actions pane. As an alternative, you can select the Libraries tab in the center pane and then click the Associate Libraries icon. The Associate Libraries window is displayed.

3. In the Available Libraries table, select the libraries that you want to associate with the Oracle VM Server.

4. Click Associate Libraries. A job is submitted to associate the selected libraries with the Oracle VM Server.

Managing Networks of Oracle VM Server

You can assign networks to Oracle VM Server and also unbind the networks. You can now assign a network multiple times to an Oracle VM Server to connect the logical domains to a network multiple times. Oracle VM Server creates a virtual switch for each network connections. Each switch is numbered and then incremented for each connection created.

See About IPMP Groups and Aggregated Links for detailed information.

Enterprise Manager Ops Center provides the option to attach networks to stand-alone Oracle VM Servers. When you select a stand-alone Oracle VM Server in the Assets tree, the Attach Network option is available in the Actions pane. This option is not available for Oracle VM Servers that belong to a virtual pool. See Managing Networks of a Virtual Pool for assigning networks to virtual pool.

To Attach Networks to Stand-Alone Oracle VM Servers

1. Select the Oracle VM Server from the Assets tree in the Navigation pane.

2. Select Attach Networks in the Actions pane.

3. Select a network to assign to the Oracle VM Server.

4. Enter the number of connections to the network. To assign more networks, click the Add icon. Click Next.

5. For each connection, define the interface:

    a. Select the NICs from the list for each connection.

6. Specify the Address Allocation for each network connection:

    - For Static IP, specify the IP address.

    - For Assign By DHCP, specifiy the IP address of the DHCP server used by the Oracle VM Server.

      Click Next.

7. Review the attributes and then click Finish to assign the network to Oracle VM Server.

Creating Logical Domains

A logical domain is a virtual machine that has its own operating system and identity within a single 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.

Enterprise Manager Ops Center can create logical domains and install Oracle Solaris OS on them. Using profiles and deployment plans, you can to create more than one domain simultaneously and then save the configuration for future logical domains.

To create logical domains:

- Create Logical Domain Profile and Plan – You create a profile which defines the configuration of the logical domain. Using the profile, you create a deployment plan. You apply the plan on an Oracle VM Server to create logical domains. The logical domains do not have the Oracle Solaris OS. You then select each logical domain and apply an OS provisioning plan to install the OS. See Provisioning 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..

Creating Logical Domain Profiles

A logical domain profile contains the requirements and configuration for an entire logical domain, including the CPU Threads, memory, storage, and network details . When a deployment plan applies the profile, you create a logical domain. You must provision Oracle Solaris OS on each domain separately. As an alternative, you can include the logical domain profile along with an OS provisioning profile in a complex plan.

See Logical Domain Profiles for detailed procedures for creating logical domain profiles and plans.

Provisioning Logical Domains

After you create a logical domain, you must provision an OS.

Before You Begin

- Ensure that you have imported the appropriate Solaris 10 (at least Solaris 10 8/07) ISO image into the library.

- Create or identify an OS provisioning profile for SPARC. See Creating an OS Profile for Oracle VM Server.

- Create or identify a deployment plan with the OS provisioning profile.

To Provision a Logical Domain

1. Select a logical domain from the Assets in the Navigation pane.

2. Select Install Server in the Actions pane. The Install Server window is displayed with the list of deployment plans for logical domains. If no plans are listed, click Create New Plan to create an OS provisioning plan. See Creating an OS Profile for Oracle VM Server to create a plan.

3. Select an available plan from the list.

4. Select whether you want to apply the plan with minimal user interaction or override any profile values.

5. Click Apply Plan. The Provision OS wizard launches. You can apply the plan in the following methods:

    - Minimal user interaction – All the OS provisioning configuration values are applied but you are prompted to enter the network resource assignment such as the IP address of the system. Enter the IP address and click Next.

    - Override profile values – You apply each configuration value in the profile. Accept or change the values and then click Next.

6. Review the information in the Summary step and click Apply to install the OS on the logical domain.

Managing Logical Domains

To manage the logical domains that are created through Enterprise Manager Ops Center, you can do the following operations:

- View and edit logical domain information.

- Start a logical domain

- Shut down a logical domain

- Reboot a logical domain

- Modify the attributes of a logical domain

- Migrate logical domains

- Add storage to logical domains

- Connect networks to logical domains

Viewing Logical Domain Information

You can manage and monitor the logical domain using the tabbed windows in the center pane.

To View Logical Domain Information

1. Select a logical domain from the Assets pane. The Dashboard is displayed in the center pane with the following information:

   - Domain Name

     Description

     Status

     Running Time

     Operating System

     Tag

     Metadata Library

     Boot Device

     Membership Graph

     Problem Status

2. Select one of the tabs.

   - Summary tab

     - Name, details and status

     - CPU Threads and utilization

     - Crypto units

     - Memory size

     - Boot device

     - Virtual disk images

     - Hourly chart of CPU and Memory Utilization

   - Console tab

     Use the Console tab to connect to the virtual console of the logical domain. Click Enable Console to activate the console. See Using the Logical Domain Console.

   - Network tab

     The Network tab displays the networks that are assigned to this logical domain with their details such as NIC Name, IP Address, MAC Address, and Network. See Managing Logical Domain Networks.

- Storage tab

  The total storage size of all the virtual disks that are associated with the this logical domain and a table showing each virtual disk is displayed. These virtual disks can be from NFS or local libraries, or LUNs from the Fibre Channel library. When you select a LUN , details such as name, vendor, GUID, product, status, and revision are displayed.You can add virtual disks and delete any of the virtual disks.

- Problems tab

  The Problems tab shows reported problems and lists unresolved critical problems. You can take the suggested actions and review the history of the alert. See Problem Management.

- Monitoring tab

  You can set monitoring rules to monitor the performance of the logical domain. You can view the set monitoring rules and the status of the rules set.

  To change the monitoring values, see Monitoring Profiles.

- Chart tab

  The Charts tab displays the CPU and network utilization of the logical domain. The data is collected in five-minute intervals and is displayed graphically. The utilization data is provided for different time intervals. See About Charts for more information about reading charts.

- Jobs tab

  The Jobs tab lists the current and completed jobs of the logical domain. See Job Management for more information about managing jobs.

Editing the Attributes of a Logical Domain

You can change the attributes and resource assignments of a logical domain.

To Edit the Attributes of a Logical Domain

1. Select the logical domain in the Assets section of the Navigation pane. The Dashboard of the logical domain is displayed.

2. Select the Summary tab in the center pane.

3. In the Actions pane, click Edit Attributes.

4. Modify the values of Description, Tag, CPU threads, Crypto units, and Memory as required.

5. Click the Save icon to save the changes.

Starting a Logical Domain

To start a logical domain, you must select a stand-alone Oracle VM Server or virtual pool that has sufficient resources to host the logical domain. When you start a logical domain, the logical domain is powered on and the OS is booted.You can change the following during the start process:

- Virtual pool wher e the Oracle VM Server resides

- Oracle VM Server host of the logical domain

- Network used by the logical domain

To Start a Logical Domain

1. Select the logical domain that you want to start from the Assets section of the Navigation pane.

2. Select Start Guest from the Actions pane. The Start Guest wizard is displayed.

3. Select the virtual pool in which you want the logical domain to run. Click Next.

> **Note:** If any network is assigned to the logical domain but not associated with the virtual pool, the network is not available after the logical domain starts.

4. Select the Oracle VM Server that is the host for the logical domain or choose a different Oracle VM Server. You can view the load of each Oracle VM Servers and select an Oracle VM Server from the table. Click Next.

5. Specify the network interfaces for the logical domain. Specify the virtual switch that connects to the logical domain's MAC address. You cannot connect the same virtual switch to two different MAC addresses

   (Optional) You can select Do Not Connect if you do not want a connect a logical domain MAC address to any virtual switch.

   Click Next .

6. Choose whether to start the logical domain immediately or to specify the date and time to start. Click Next.

7. Review the selected parameters and click Start to start the logical domain. A job is submitted to start the logical domain.

Managing Logical Domain Networks

You can connect to a logical domain to a network multiple times on the same Oracle VM Server. For each connection, a virtual switch is created in the Oracle VM Server but you must specify the NIC each virtual switch uses. This association is done when attaching networks to Oracle VM Server.

You can now create IPMP groups in logical domains manually. It takes approximately 15 minutes for a new IPMP group to display in the UI.

You can also aggregate the NICs for the logical domain and create an aggegated link.

See About IPMP Groups and Aggregated Links for information about IPMP and link aggregation support in Enterprise Manager Ops Center. For more information about creating IPMP and link aggregations, see Additional Resources.

To Connect a Logical Domain to a Network

1. Select a logical domain that is in shutdown state in the Assets section of the Navigation pane.

2. Click Connect Guest to Network in the Action pane. The Connect Guests to Network window is displayed. The networks that are not currently connected to the selected logical domain are displayed in the table.

3. Select a network that you want to connect to the logical domain. Use Shift and Ctrl keys to select more than one network.

4. Click Connect to Network. The selected network is connected to the logical domain

To Disconnect a Logical Domain from a Network

1. Select a logical domain that is in shutdown state in the Assets section of the Navigation pane.

2. Select the Network tab in the center pane. The list of connected networks are listed.

3. Select the network that you want to unbind.

4. Click the Disconnect Selected Guest From Network icon. The Unbind a Network window is displayed.

5. Select the network from the list.

6. Click Disconnect From Network to unbind the network from the logical domain. The selected network is disconnected from the logical domain.

Migrating a Logical Domain

You can migrate a logical domain from an Oracle VM Server to another virtual host within the same virtual pool. This type of migration is called warm migration. The logical domain is shut down or suspended and then its static configuration is moved to the target host. The logical domain is then restarted on the target host using the same virtual image. Warm migration does not maintain the live run state of the operating system, that is, going through a save-to-disk cycle during the migration. Shared storage is required for this implementation.

Before You Begin

- Verify that the source and target hosts are in the same virtual pool.

- Verify that the source and target hosts are in the same network.

- Verify that the source and target host have access to the storage library, that is, the disk image or guest metadata, for the logical domain .

- For Oracle VM Server for SPARC 1.2 version, you can migrate with only 1 MAU.

  For Oracle VM Server for SPARC 1.3 and 2.0 versions, you can migrate with any number of MAUs.

- Verify that the target host has sufficient resources (CPU threads, memory) ato support the logical domain.

  > **Note:** For Oracle VM Server for SPARC 2.1, you cannot migrate a logical domain to a previous version of the server such as 2.0, 1.3 or 1.2. You can migrate only within the same version of Oracle VM Server for SPARC.

To Migrate a Logical Domain

1. Select the logical domain from the Assets section of the Navigation pane.

2. Click Migrate Guest in the Actions pane.. The Migrate Logical Domain wizard is displayed.

3. Select the target Oracle VM Server, the new host of the logical domain. The eligible target Oracle VM Servers is displayed in order of decreasing preference. Click Next.

4. Review the migration Summary and click Finish to migrate the logical domain.

Rebooting a Logical Domain

You can reboot a logical domain from the Enterprise Manager Ops Center UI.

To Reboot a Logical Domain

1. Select the logical domainfromt he Assets section of the Navigation pane..

2. Click Reboot Guest in the Actions panet. The Reboot Logical Domain window is displayed

3. Click Reboot Logical Domain to confirm. A job is submitted to reboot the logical domain.

Managing Storage of a Logical Domain

You can use the options Add Storage to Logical Domain and Remove Storage From selected Logical Domain to manage storage of logical domains.

Before You Begin

Ensure that the logical domain is in the shut down state to enable the storage options on the UI.

To Add Storage to Logical Domain

1. Select the logical domain from the Assets section of the Navigation pane.

2. Select Add Storage to Logical Domain from the Actions pane. The Add Storage window is displayed.

3. Click the Add icon to add a storage to the logical domain.

4. Select an associated library from the list. The library can be local, NAS, SAN or local device library.

5. If the selected library is local or NAS, specify the size of the disk storage. You cannot specify the disk size for SAN and local device library.

6. (Optional) You can add storage from more than one associated library. The total additional storage disk will be displayed for Additional Storage Specified. The total value for logical domain storage is updated.

7. Click Finish. The selected storage will be added to the logical domain.

To Remove Storage from Logical Domain

1. Select the logical domain from the Assets section of the Navigation pane.

2. Select the Storage tab in the center pane. The list of storage is displayed.

3. Select the storage that you want to remove. The Remove Storage From Selected Logical Domain icon is enabled.

4. Click the Remove Storage From Logical Domain icon.

5. Click Delete to confirm the action. The selected storage will be removed from the logical domain.

Shutting Down a Logical Domain

When you shut down a logical domain, it is disconnected from its networks, disassociated from the Oracle VM Server control domain, and all the associated resources are released. You can restart the logical domain with the same configuration

that is maintained in its storage library or you can restart the logical domain on a different Oracle VM Server.

To Shut Down a Logical Domain

1. Select a logical domain from the Assets section of the Navigation pane.

2. Click Shutdown Guest in the Actions pane.

3. Click Shutdown Logical Domain. A job is submitted to shut down the logical domain.

Deleting 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 the associated resources are released and the domain configuration is removed from the library. All the references to the logical domain, including its disk image and snapshots are removed from the system. However, the logical domain profile and the plan that used the profile remains.

To Delete a Logical Domain

1. Select the logical domain from the Assets section of the Navigation pane..

2. Click Delete Guest in the Action pane.. The Delete Logical Domain window is displayed.

3. Click Delete Logical Domain. A job is submitted to delete the logical domain.

Using the Logical Domain Console

You can attach to the logical domain console from witin the Enterprise Manager Ops Center UI. You enable the console connection and then you connect to the console.

To Open a Logical Domain Console

1. Select the logical domain from the Assets section in the Navigation pane.

2. Select the Console tab in the center pane.

3. Click the Enable the Console Connection icon. A job is initiated to enable the console connection.

4. Click the Connect to the Console icon. The connection is established and the console appears on the UI.

5. Enter the user name and password to log into the logical domain.

6. (Optional) Click the Undock the Console icon to move the console window separately from the browser window.

7. (Optional) The default connection time is set to 120 minutes. To change the expiration time, click the Set Console Connection Timeout icon and enter a new time limit in minutes. The console connection is enabled for the new time limit. After the time limit, you must enable the console connection again.

> **Note:**   If 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.

To Close a Logical Domain Console

1. Select the logical domain from the Assets tree in the Navigation pane.

2. Select the Console tab in the center pane. The logical domain console that is already enabled is displayed.

3. Click Disable the Console Connection.

> **Note:** If the logical domain is not in view in the Assets tree, the logical domain is automatically logged out but the connection exists until the connection time expires.

12

Oracle Solaris Zones

Oracle Solaris Zones, previously known as Solaris Containers, are an integral part of the Oracle Solaris 10 operating system (OS). Zones isolate software applications and services using flexible software-defined boundaries. You can use zones to create private execution environments within a single instance of the Oracle Solaris 10 OS. Each environment has its own identity that is separate from the underlying hardware. Each environment behaves independently as if running on its own system, making consolidation simple, safe, and secure.

You build Oracle Solaris Zones by using the following technologies:

- Solaris Resource Manager
- Solaris Zones partitioning technology

About Oracle Solaris Zones

A zone is a virtualized operating system environment created within a single instance of the Solaris OS. Within a zone, the operating system is represented to the applications as virtual operating system environments that are isolated and secure. The applications run in different zones with complete isolation, while the underlying operating system resources are centrally managed and administered. This section describes the varieties of zones.

Global and Non-Global Zones

The Oracle Solaris 10 operating system supports two types of zones:

- The global zone is the default operating system and has control over all the processes. 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 any other zones.

Each Oracle Solaris 10 system contains a global zone. The global zone is both the default zone for the system and the zone used for system-wide administrative control. All processes run in the global zone if no non-global zones are created.

- Booting the global zone is equivalent to booting the system hardware.
- Each zone, including the global zone, is assigned a zone name. The global zone always has the name "global".
- Each zone is assigned a unique numeric identifier. The global zone always has the identifier ID 0.

- Each zone has a path to its root directory that is relative to the global zone's root directory.

- The global zone is the only zone from which a non-global zone can be configured and installed.

You can 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 can are managed independently of each other. For example, different versions of the same application can be run in separate zones.

Enterprise Manager Ops Center provides complete solution for managing your virtualized environment. Use Enterprise Manager Ops Center to do the following:

- Create zones

- Discover and manage existing zones

- Manage zones such as booting, rebooting, shutting down, cloning, migrating and deleting zones.

- Manage zone file systems, storage, networks and per-zone resource utilization.

Types of Non-Global Zones

You can create various types of non-global zones for different purposes.

By File Systems

Use the global zone to create the following types of non-global zones:

- **Sparse Root Zone** – Contains a read/write copy of a portion of the file system that exists on the global zone. Other file systems are mounted read-only from the global zone as loop-back virtual file systems. When a sparse root zone is created, the global administrator selects which file systems to share with the sparse root zone in addition to 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.

- **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** – Supports different versions of Solaris OS for running applications. For example, you can install Solaris 8 or 9 in a branded zone.

Greenfield and Brownfield Zones

Zones that are created using Enterprise Manager Ops Center are referred to as greenfield zones and represented with the icon in the user interface. Zones that are were created using command line interface (CLI) are referred to as brownfield zones and represented with the icon in the user interface. For greenfield zones, the Enterprise Manager Ops Center has full access to the zone configuration data and provides complete management of the zone. For brownfield zones, Enterprise Manager Ops Center cannot detect the zone configuration and provides less support than the greenfield zones. The actions that are not available to brownfield zones are:

- Edit zone configuration

- Add file system

- Add storage

- Connect to networks

- Migrate a zone

Depending on the requirement to manage a brownfield zone's file system, storage, networks or to migrate the zone, you can convert it to greenfield zone. You can use one of the following options to convert a brownfield zone:

- Adopt zone

- Clone zone

Limitations of Enterprise Manager Ops Center's Zone Management

- Zones created by Enterprise Manager Ops Center must be administered only by Enterprise Manager Ops Center. You must not make changes or perform zone tasks using other utilities such as CLI. If you mix the management tools, Enterprise Manager Ops Center suspends the management of the zones.

- If you use Live Upgrade to update the OS in a zone, you cannot use Enterprise Manager Ops Center to manage the zone. ABE is not supported for zones that were created using Enterprise Manager Ops Center. Live Upgrade operations are available only for global zone as a whole and not for individual zones. When you migrate a zone, ABE is not supported. If you have ABEs for zones, it is recommended to be removed before migration.

- Zone management tasks such as Add File Systems, Add Storage, Connect Network and Edit Zone Configuration are enabled only if the zone is in the shut down state. However, the Migrate Zone and Clone Zone actions are enabled even when the zone is running.

Managing and Monitoring Global Zones

When you select a global zone in the Assets tree, Enterprise Manager Ops Center displays a set of tabs in the center pane:

- Dashboard

  - Name

  - Description

  - Status

  - Server name

  - Running Time

  - Operating System

  - OS Health

  - Total CPU Threads

  - Tags

  - Membership Graph

  - Problem Status

- Compliance reports

- Summary

  - Name

  - Description

  - Status

  - Server name

  - Running Time

  - Operating System

  - Name service

  - Total CPU threads

  - Total CPU and memory utilization

  - Non-global zones and its status

- Libraries

  The Libraries tab lists all the storage and software libraries associated with the global zone. It also displays the local devices that are attached to the global zone. You can select the associated library and see the contents of the library and the non-global zones that are associated with it. You have the option to associate a library or to disassociate alibrary from the global zone.

- Utilization

  The Utilization tab displays charts for CPU, memory, swap, IO, and file system utilization. Kernal statistics such as memory allocation, paging, IPC, and CPU details are also displayed. The the top ten processes for CPU and memory usage are also included.

- Networks

  The Network tab lists all the networks that are assigned to the global zone. You have the option to assign more networks, unassign a network, and modify the physical connectivity of a network. The IPMP groups and link aggregation are also defined in this tab. For each assigned network, you can view the additional connectivity information and the zones that are using this network.

- Problems

  The Problems tab shows the unresolved problems that are reported for the global zone and a graphical representation of the problem composition. You can also view the corresponding alerts for each problem reported. See Problem Management.

- Monitoring

  The Monitoring tab shows the rules and values that are set for zone activity.

  See Monitoring Profiles and Rulesfor detailed information about setting monitoring actions.

- Charts

  The Charts tab displays the CPU, memory and network utilization of the global zone. The data is collected in five-minute intervals and is displayed graphically. The utilization data is provided for different time intervals. See About Charts for

more information about reading charts. The utilization data can be viewed in the following ways:

- Total – Aggregated utilization data for all the zone including the global zone.

- Kernal – Utilization data for the kernal process only.

- Global – Utilization data for the global zone only.

The following utilization data are displayed:

- CPU utilization – You can display the shared CPU utilization for each zone in the global zone.

- Network utilization – The chart displays the total utilization for the all the zones including the global zone. Network usage is not available for each zone.

- Memory utilization – The memory utilization is the physical, virtual, and locked memory in use by each zone.

- File System Utilization – . The file system lists the root file system and other file systems corresponding to the zone paths in the global zone. You can select a file system and display the utilization rate for a historical period.

- Terminal

    The Terminal tab give you access to the console for a zone. Click the Enable the Console Connection icon to invoke the zlogin to the zone. Log into the zone using the zone user name and password. The zone console exists until you log out of Enterprise Manager Ops Center, click the Disable the Console Connection icon, or the global zone is shut down. You can click the Undock the Console icon to undock the console session from the Enterprise Manager Ops Center, making it available separately for your use.

- Job

    The Jobs tab lists the current and completed jobs for the global zone. See Job Management for more information about managing jobs.

- Configuration

    The Configuration tab lists the access points for the global zone. Access points are the resources associated with the operating system and are created as a result of discovery or when the agent starts up. This tab lists the Proxy Controller that connects the global zone to the Enterprise Controller. It also shows the status of the connection and the protocol used for the connection. You have option to unmanage or delete the discovered resources.

Managing and Monitoring Non-Global Zones

When you select a non-global zone in the Assets tree, Enterprise Manager Ops Center provides the following tabs in the center pane :

- Dashboard

    - Name

    - Description

    - Status

    - Running Time

    - Operating System

- OS Health

- Host name

- Total CPU Threads

- Tags

- Membership Graph

- Problem Status

- Compliance reports

- Summary

  - Zone name and its state

  - Type of zone and the zone path

  - Hostname

  - Zone library and storage size

  - CPU model and utilization

  - File systems allocated to the zone

  - Storage size and usage

  - Hourly utilization chart of CPU and memory resources

- Console

  The Console tab gives you access to the console of a zons. Click the Enable the Console Connection icon to invoke the zlogin to the zone. Log into the zone using the zone username and password. The zone console exists until you log out of Enterprise Manager Ops Center, click the Disable the Console Connection icon, or the global zone is shut down. You can click the Undock the Console icon to undock the console session from the Enterprise Manager Ops Center, making it available separately for your use.

- Utilization

  The Utilization tab displays charts for CPU, memory, swap, IO, and file system utilization. Kernal statistics such as memory allocation, paging, IPC, and CPU details are also displayed. The the top ten processes for CPU and memory usage are also included.

- Networks

  For a greenfield zone, the Network tab lists all the networks that are assigned to the zone. You have the option to connect a network to a zone in the shutdown state and disconnect networks from a zone. The associated IPMP groups are also listed for the zone. This information is not available for brownfield zones.

- Storage

  For a greenfield zone, the Storage tab lists the root file system, the default file system of the zone. It also lists other file systems that were added to the zone. You cannot edit the read and write access to the root file system.

  The Storage tab also displays the virtual disks that make up the zpool, which is created when you create a zone. The zpool consists of the virtual disks and LUNs used for storage. File systems are implemented as ZFS file systems and mapped to the storage.

From the Storage tab, you can add a file system to the zone and change an existing file system, delete a file system from the zone, and add more storage to a zone.

This information is not available for a brownfield zone.

- Problems

  The Problems tab lists out the unresolved problems that are reported for the zones and a graphical representation of the problem composition. You can also view the corresponding alerts for each problem reported. See Problem Management for more detailed information.

- Monitoring

  The Monitoring tab shows the rules and values that are set for zone activity.

  See Monitoring Profiles and Rules for detailed information about setting monitoring actions.

- Charts

  The Charts tab displays the CPU, memory and network utilization of the global zone. The data is collected in five-minute intervals and is displayed graphically. The utilization data is provided for different time intervals. See About Charts for more information about reading charts. For a zone, the following utilization data are displayed:

  - CPU utilization – You can display the CPU utilization for the zone. The per-zone utilization data is collected by the extended accounting in the global zone.

  - Network utilization –   The chart displays the total utilization for the all the zones including the global zone. Network usage is not available for each zone.

  - Memory utilization – The memory utilization is the physical, virtual, and locked memory in use by each zone.

  - File System Utilization – Whenever you create a zone, a zpool is created with the zone guest image name. The file system lists the root file system and other file systems that are added to the zone. You can select a file system and display the utilization rate for a historical period. The OS Details page lists all the file systems and the current utilization rate.

- Job

  The Jobs tab lists the current and completed jobs for the zone. See Job Management about managing jobs.

- Configuration

  The Configuration tab displays the information about the zone's configuration, depending on whether the zone is a greenfield zone or a brownfield zone.

  For greenfield zones, the Configuration tab shows the CPU, memory caps, shared memory, scheduler and semaphore IDs.

  For brownfield zones, the Configuration tab shows the the access points for the zone. Access points are the resources associated with the operating system and are created as a result of discovery or when the agent starts up. This tab lists the Proxy Controller that connects the non-global zone to the Enterprise Controller. It also shows the status of the connection and the protocol used for the connection. You have option to unmanage or delete the discovered resources.

Adopting a Zone

To convert a brownfield zone to a greenfield zone, adopt the brownfield zone. When you adopt a zone, you get access to the zone configuration data and greater ability to manage and monitor the zone. When you adopt a zone, you enable the following actions for the zone: Add Storage, Add File Systems, Migrate Zone, Replicate Zone, and Connect Networks to the zone.

During the adoption process, the zone is shutdown and then rebooted after adoption. A new zpool is created for the zone and the existing file system is copied to the zpool. You cannot add or delete a file system during the adoption process.

> **Note:** You cannot change the zone name.

Before You Begin

- Verify that the zone is a brownfield zone.

- Verify that the zone is in the managed state.

- Verify that the zone is running.

To Adopt a Zone

1. Select the brownfield zone in the Navigation pane.

2. Click Adopt Zone in the Actions pane. A warning message is displayed.

3. Click OK to continue and to start the Adopt Zone wizard.

4. Select a local library or a NAS library to save the zone's metadata. The metadata is the zone's configuration, the virtual image. If you want to allow the zone to be migrated, choose a NAS library. Click Next.

5. (Optional) You can modify the zone descriptions and tags. Click Next.

6. Modify the Reserved and Quota size of the existing file system. Click Next.

7. Click the Add icon to associate a storage library with the zone.

8. Select a library from the list of libraries that are associated with the global zone. This library will store the data from the zone's operation. You can select the following types of library from the list. If you have the option to change the size, provide sufficient storage space to accommodate any new file systems added to the zone. Allocate 20% more than the required storage space for a good performance.

   - Local Library – It is the global zone's local storage disks. Select a local library and provide the virtual disk size. However, if you choose a local library for zone storage, the zone cannot be migrated.

   - Local Device – These are disks that are connected to the system that is running the global zone. Select a disk from the list. The size of the disks cannot be changed.

   - NAS Storage – Select a NAS Storage library from the list. Define the size of the virtual disk. You can modify the disk name.

   - Fibre Channel – Select a Fibre Channel library from the list. The LUNs for that library are listed under the Virtual Disk Name. Select a LUN from the list. You cannot modify the LUN size or create a new virtual disk image. The total size of the selected libraries are displayed as Currently Accumulated Storage in GB.

Click Next.

9. Review the summary of the attributes you chose and click Finish to adopt the zone.

When the adopt zone job completes, the following changes occur:

■ The brown icon of the zone changes to green.

■ The agent is installed on the zone and the zone will be in a managed state.

■ The actions for the zone such as Add Storage, Add File System, Replicate Zone and Connect Network are enabled in the actions pane. The Migrate Zone option is enabled only if the zone metadata and added storage are on the shared storage.

Editing Zone Configuration

> **Note:**   You can edit the configuration of the non-global zone only if it has been created using Enterprise Manager Ops Center.

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.

 To Edit Zone Configuration

1. Select the zone from the Assets section in the Navigation pane.

2. Click Shutdown Zone from the Actions pane. Shutdown Zone window is displayed.

3. Click Shutdown Zone to shut down the zone immediately.

4. Select the zone in the Assets section of the Navigation pane. The Dashboard of the zone is displayed in the center pane.

5. Select the Configuration tab in the center pane. The current configuration of the zone is displayed.

6. Click Edit Configuration in the Actions pane. The Edit Zone Configuration window is displayed. Edit the following zone attributs, as required:

   - CPU Model
   - Scheduler
   - Memory Caps
   - LightWeight Processes
   - Message IDs
   - Semaphore IDs
   - Shared Memory IDs
   - Shared Memory

7. Click the Save icon in the center pane to save the changes.

8. Select the zone in the Assets tree, then click Boot Zone in the Actions pane or click the Boot Zone icon in the Zones pane.

9. Click Boot Now or click Boot Later.

You can modify the zone description and tags.

To Edit Zone Attributes

1. Select the zone from the Assets section in the Navigation pane.

2. Click Shutdown Zone from the Actions pane. Shutdown Zone window is displayed.

3. Click Shutdown Zone to shut down the zone immediately.

4. Select the zone in the Assets section of the Navigation pane. The Dashboard of the zone is displayed in the center pane.

5. Select Edit Attributes in the Actions pane. The following zone attributes can be modified:

   - Description
   - Tags

6. Click the Save icon in the center pane to save the changes.

7. Select the zone in the Assets tree, then click Boot Zone in the Actions pane or click the Boot Zone icon in the Zones pane.

8. Click Boot Now or click Boot Later.

Adding Storage to Zones

You can add storage only to greenfield zones. The storage is added to the zpool that is created for the zone. You cannot delete storage from the zpool.

To Add Storage to Zones

1. Select the zone from the Assets section in the Navigation pane.

2. Click Shutdown Zone from the Actions pane. Shutdown Zone window is displayed.

3. Click Shutdown Zone to shut down the zone immediately.

4. Select the zone from the Assets section in the Navigation pane.

5. Click Add Storage in the Actions pane. As an alternative, select the Storage tab and click the Add Storage to Zone icon under Zpool and Storage. The Add Storage to the Zone window is displayed.

6. Click the Add icon.

7. Select a library from the list. The library can be a local storage, NAS, local device, or Fibre Channel library.

8. For a local and NAS library, enter a virtual disk name and specify the size of the disk. For a Fibre Channel library, select a LUN from the list. For local device library, select a disk from the Virtual Disk Name list. You cannot specify the size of the disk or the LUNs. The Additional Storage Specified shows the new total storage space..

9. Click Add Storage to add the specified storage to the zone.

10. Select the zone in the Assets tree, then click Boot Zone in the Actions pane or click the Boot Zone icon in the Zones pane.

11. Click Boot Now or click Boot Later.

Adding File Systems to Zones

You can add file systems only to greenfield zones.

To Add a File System to a Zone

1. Select the zone from the Assets section in the Navigation pane.

2. Click Shutdown Zone from the Actions pane. Shutdown Zone window is displayed.

3. Click Shutdown Zone to shut down the zone immediately.

4. Select the zone in the Assets tree.

5. Click Add File Systems from the Actions pane. The Add File Systems window is displayed.

6. Click the Add icon .

7. Enter the Reserved and Quota size for the new file system.

8. Click Add File Systems.

9. Select the zone in the Assets tree, then click Boot Zone in the Actions pane or click the Boot Zone icon in the Zones pane.

10. Click Boot Now or click Boot Later.

Adding Libraries to the Global Zone

You can associate one or more libraries with a global zone. The libraries are available to all non-global zones in the hierarchy of the global zone.

To Associate a Library With the Global Zone

1. Select the global zone from the Assets section of the Navigation pane. The Dashboard of the global zone is displayed.

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.

4. Click Associate Libraries. A job is submitted to associate the libraries with the global zone.

Disassociating Libraries From Global Zone

When you disassociate a library from global zone:

■   The libraries are disassociated 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.

To Disassociate Libraries From Global Zone

1. Select the global zone from the Assets section in the Navigation pane. The Dashboard of the selected global zone is displayed in the center pane.

2. Select Libraries tab from 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. The Disassociate Library window is displayed.

5. Click Disassociate Library to confirm unmounting the library.

Using Networks With Zones

This section describes how to connect, modify, and disconnect networks from zones.You can create networks associated with an IPMP group.

See Managed Networks for more information about IPMP and Link Aggregation.

Modifying the Network Interface of a Global Zone

You can change the permanent IP address or change the allocation method.

> **Note:**   If the network interface is an IPMP group, you cannot modify the network interface.

1. Select the global zone from the Assets section of the Navigation pane. The Dashboard of the global zone is displayed

2. Select the Networks tab in the center pane.   The Networks tab lists the networks that are assigned to the global zone.

3. Select the network that you want to modify.

4. Click the Modify Physical Connectivity icon. The Change VirtualHost Network/NIC Connection window is displayed.

5. If the network interface has a static IP address, you can change the permanent IP address or you can change the Address Allocation method to Use Ext. DHCP

6. Click Submit to save the changes.

Assigning Networks to a Global Zone

In Enterprise Manager Ops Center, networks are managed objects. You can identify and assign one or more networks to the global zone. These networks can then be used by the global zone's non-global zones.

When you assign a network to a global zone, you must define the mode of the network, either Shared IP or Exclusive IP:

- In Shared IP mode, the network interface is used by more than one zone. You define the network interface when you you assign the global zone to the network.

- In Exclusive IP mode, the network interface is dedicated to the zone. An exclusive network must be declared for the global zone when you assign network to the global zone. Then you configure the IP configuration for the non-global zone.

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.

To Assign a Network to a Global Zone

1. Select a global zone from the Assets section in the Navigation pane.

2. Select Assign Network in the Actions pane. The Assign Networks to Global Zone wizard is displayed.

3. Select the network that you want to associate with the global zone from the list of networks that are not assigned to the global zone. To assign more networks to the global zone, click the Add icon.

4. For each selected network , specify the IP mode, either Shared or Exclusive. Click Next.

   - If you selected Shared IP mode, click Next to go to Step 5.

   - If you selected Exclusive IP mode, click Next to go to Step 7.

5. In the Specify Networks Interfaces step, select a NIC. If the NIC is a member of an IPMP group, you can select the group.

6. Select the address allocation method. If you selected an IPMP group, address allocation is not required.

   - Use Static IP – Enter the IP address of the global zone.

   - Assigned by External DHCP – The global zone uses a DHCP server external to Enterprise Manager Ops Center to acquire an IP address.

   Click Next to go to Step 12.

7. Click Connect Network in the Actions pane or select the Networks tab in the center pane and click the Connect Network icon. The Connect Networks to Zone window is displayed.

8. Select the networks that you want to connect to the zone. All networks assigned to the global zone with the same IP mode and not yet connected to the zone are listed.

9. Click the Add icon. You can assign an exclusive IP zone multiple times to a network. You can then create an IPMP group inside the exclusive IP zone.

10. Specify the network interface of the selected networks.

11. Click Connect Networks.

12. Click Finish to submit the job that assigns the network to the global zone..

Removing a Network From a Global Zone

When you unbind a network from a global zone, the global zone's non-global zones are also disconnected from the network.

To Unbind a Network From a Global Zone

1. Select the global zone from the Assets tree in the Navigation pane. The Dashboard of the global zone is displayed.

2. Select the Networks tab in the center pane. The networks that are assigned to the global zone are listed.

3. Select the network that you want to disconnect. The Unbind Network icon is enabled. You can unbind only one network at a time.

4. Click the Unbind Network icon. The Unbind Network from Global Zone window displays the list of zones from which the network will be removed.

5. Click Disconnect Zones.. A job is submitted to remove the network from the global zone.

Using Networks With Non-Global Zones

Before You Begin

- To assign a network to a non-global zone, verify that the zone is a greenfield zone. You can disconnect the networks of brownfield zones and greenfield zones.

- Shut down the non-global zone using the Shutdown Zone action.

To Connect a Network to a Non-Global Zone

1. Select the zone in the Assets section of the Navigation pane.

2. Click Connect Network in the Actions pane. As an alternative, select the Networks tab in the center pane and click the Connect Network icon. The Connect Networks to Zone window is displayed.

3. Select the network that you want to connect to the non-global zone. The list of networks includes all networks assigned to the global zone that have the same IP mode (Shared IP or Exclusive IP) and that are not yet connected to the non-global zone. Click the Add icon to add more networks.

4. Specify the network interface for the selected networks. You can assign an exclusive IP network multiple times to the same zone and then create an IPMP group the network.

5. Click Connect Networks to submit the job to assign the networks to the zone.

To Disconnect a Network From a Non-Global Zone

1. Select the zone in the Assets section of the Navigation pane.

2. Select the Networks tab in the center pane. The networks that are assigned to the zone are listed.

3. Select the network that you want to disconnect from the zone. The Disconnect Zone Network option is enabled.

4. Click the Disconnect Zone Network icon. A job is initiated to disconnect the network from the zone.

Managing Zones

From within Enterprise Manager Ops Center, you can boot, reboot, shut down, halt, clone, attach a console to, and delete greenfield and brownfield zones.

Booting a Zone

Booting a zone places the zone in the running state, using the current configuration. The Boot Zone option is enabled only when the zone is in the ready or installed state, but not in the running state.

If you have the autoboot properties were set when the non-global zone was created, the zone boots whenever the global zone boots.

Before You Begin

Ensure that the zone is in the Ready state or Installed state.

To Boot a Zone

1. Select the zone from the Assets section of the Navigation pane. The Dashboard page of the zone is displayed.

2. Click Boot Zone in the Actions pane. The boot zone job is submitted.

An alternative procedure also boots the zone:

1. Select the global zone from the Assets section of the Navigation pane.

2. Select the Summary tab.

3. Select the non-global zone from the Zones section.

4. Click the Boot Zone icon. The boot zone job is submitted.

Rebooting a Zone

You can reboot a zone that is in the running state. This type of reboot shuts the zone down and then boots it. This is different from the `zoneadm reboot` command in which the zone is first halted and then booted.

To Reboot a Zone

1. Select a zone from the Asset section in the Navigation pane. The Dashboard page of the selected asset is displayed in the center pane.

2. Click Reboot Zone in the Actions pane.

3. Click the Reboot button to confirm the action. The reboot job is initiated.

An alternative procedure also reboots the zone:

1. Select the global zone from the Assets section of the Navigation pane.

2. Select the Summary tab.

3. Select the non-global zone from the Zones section.

4. Click the Reboot icon. The boot zone job is submitted.

Attaching to the Zone Console

You enable the console connection and then connect to the console. The default connection time is 120 minutes.

To Attach to a Zone Console

1. Select the zone from the Assets section in the Navigation pane. The Dashboard of the selected zone is displayed.

2. Select the Console tab in the center pane.

3. Click the Enable the Console Connection icon. A job is initiated to enable the console connection.

4. Click the Connect to the Console icon. The connection is established and the console appears on the UI.

5. Enter the zone user name and password to log into the zone.

6. (Optional) Click the Undock the Console icon to make the console available separately from the UI.

7. (Optional) Click the Set Console Connection Timeout icon to set the number of minutes for the connection. When the time expires, you must enable the console connection again.

> **Note:** If the zone is not in view in the Assets tree, the zone is automatically logged out. To continue, connect to the zone console again.

Disabling the Zone Console Connection

If the zone is not in view in the Assets tree, the zone is automatically logged out but the connection exists until the connection time expires or you disable the connection.

To Disable a Zone Console Connection

1. Select the zone from the Assets tree in the Navigation pane. The Dashboard of the zone is displayed.

2. Select the Console tab of the zone in the center pane. The zone console that is enabled is displayed.

3. Click Disable the Console Connection to disable the zone console.

Shutting Down a Zone

Shut a zone down in a graceful manner so that it is in a state that can be restarted.

To Shut Down a Zone

1. Select the Zone from the Assets section in the Navigation pane. The Dashboard page of the selected zone is displayed in the center pane.

2. Click Shutdown Zone from the Actions pane. Shutdown Zone window is displayed.

3. Click the Shutdown Zone button to confirm the shut down action.

An alternative procedure also shuts down the zone:

1. Select the global zone from the Assets section of the Navigation pane.

2. Select the Summary tab.

3. Select the non-global zone from the Zones section.

4. Click the Shutdown icon. The shutdown job is submitted.

The zone is in the Installed state.

Halting a Zone

When you halt a zone, it removes both the application environment and the virtual platform for a zone. Halting a zone changes the zone's state to the Installed state. All processes are killed, devices are unconfigured, network interfaces are unplumbed, file systems are unmounted, and the kernel data structures are destroyed.

To Halt a Zone

1. Select a zone from the Assets section in the Navigation pane.

2. Click Halt Zone in the Actions pane. The Halt Zone window is displayed.

3. Click the Halt Zone button to confirm the halt zone action.

An alternative procedure also halts the zone:

1. Select the global zone from the Assets section of the Navigation pane.

2. Select the Summary tab.

3. Select the non-global zone from the Zones section.

4. Click the Halt Zone icon. The halt zone job is submitted.

Deleting a 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 made available.

- Any exclusive IPs that were assigned to the zone are available for re-use

To Delete a Zone

1. Select a zone from the Assets section in the Navigation pane.

2. Select Delete Zone from the Actions pane. The Delete Zone window is displayed.

3. Click Delete Zone to confirm whether you want to delete the zone. The zone is deleted and removed from the Assets tree.

Creating Zones Using Enterprise Manager Ops Center

When you use Enterprise Manager Ops Center to create a zone, the zone is displayed in the Assets tree with a green badge and it is managed by Enterprise Manager Ops Center. You are able to change and edit the zone's configuration.

> **Note:** If you create or migrate a zone using Enterprise Manager Ops Center, you cannot use Solaris Live Upgrade and an alternate boot environment (ABE) to update the zone.

Before You Begin

Before you create a non-global zone, determine the requirements for the new zone.

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 the Solaris operating system. For example, you can create a branded zone to run Solaris 8 applications on your 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 12–1

2. Import the images into one of the software libraries associated with the global zone.

During the process of creating a branded zone, you will install the images in the global zone. For instructions, see the README files in the Solaris 8 or 9 Containers bundle.

Table 12–1   Packages for Branded Zones

Global Zone OS	Branded Non-Global Zone
Oracle Solaris 10 10/08 OS or later:	For Solaris 8, install SUNWs8brandk from Solaris 8 Containers 1.0.1.
	For Solaris 9, install SUNWs9brandk from Solaris 9 Containers 1.0.1.
A version of Solaris prior to Solaris 10 10/08 OS:	For Solaris 8, install SUNWs8brandr and SUNWs8brandu from Solaris 8 Containers 1.0.1.
	For Solaris 9, install SUNWs9brandr, SUNWs9brandu, and SUNWs9brandk from Solaris 9 Containers 1.0.1.

Other Attributes of a 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 boots immediately after you create it and whenever the global zone boots.

- Shared CPU or Dedicated 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.

  – 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.

- Importance – 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.

- 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 need 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 10 OS configurations.

To Create a Zone

1. Select the global zone in the Assets section of the Navigation pane. This global zone is the virtual host for the new non-global zone. The Summary page of the global zone is displayed in the center pane.

2. Click Create Zone in the Actions pane. The Create Zone wizard starts.

3. Select the type of zone. Click Next.

   To create the branded type of zone, install the OS image:

   a. Select the image from the list of images in the software libraries associated with the global zone.

   b. Select the Brand Type, either Solaris 8 or 9.

   c. Wait for at least 5 minutes for Enterprise Manager Ops Center to install the packages and refresh the list of available brands.

      d. (Optional) Specify the HostID if you want to run an application only on a specific host.

      e. (Optional) Specify the Machine Type, either Sun4u or Sun4v.

4. Enter a zone name and description for the zone.

5. Select a storage library from the list of libraries that are associated with the global zone to store the zone's image and metadata. Click Next.

6. Specify the CPU type and memory thresholds. The CPU type can be 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. Click Next.

7. Configure the language, time zone, terminal-type, atnd root password for the zone.

8. For the host name, enter the zone name you specified in Step 4. Provide a domain name for the NFSv4 Domain Name or accept the default value, `dynamic`, to allowthe naming service you will specify in Step 14 to determine the NFSv4 domain at run time. Click Next.

> **Note:** The host name of the zone must be resolved by the naming service that you select in Step 14. If it is not, the zone creation job cannot boot the new zone.

9. Set the boot properties for the new zone. You can set the properties so that the zone boots after it is created and whenever the global zone boots. Click Next.

10. Configure the file systems for the 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. Set 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.

11. Configure the storage for the zone's operations. Select a library and specify the virtual disks for zone storage. The library can be local storage, NAS storage or Fibre Channel storage. If the library is local or NAS storage, specify the virtual disk name and size of the disk. If the library is Fibre Channel, 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.

> **Note:** If you specify a local storage library, you cannot migrate the zone in the future.

> **Caution:** An Oracle Solaris Zone running with a zonepath on a NFS share is not a supported configuration. If you specify this type of storage, do not use the zone for production or non-experimental workload.

12. If the Currently Accumulated Storage value is less than the Recommended Storage Size value, you can add more storage resources to the zone. Click the Add icon to configure more storage resources When the Currently Accumulated Storage size as at least equal to the Recommended Storage size, click Next.

> **Note:** Provide sufficient storage space to accommodate the file systems added to the zone. Allocate 20% more than the required storage space for good performance.

> **Note:** If you plan to use Alternate Boot Environments with zones, provide sufficient storage space. Specify twice the size of the zone's file system for the root file system of the zone. For example, if the root file system is 8 GB, the virtual disk backing this file system must be 16 GB.

13. Assign at least one network to the zone. Select the zone's IP mode, either 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. If the network has a defined IP range, the Auto Allocate IP option is displayed with the zone's IP address from the range.If the network does not have a defined IP range, the Use Static IP option is displayed. Type an IP address for the zone in the Zone IP field.

- 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. If 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.

    If 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 type 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.

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. If you know the NIS server details, choose the Specify an NIS Server option to

provide the NIS server host name and its IP address. If 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.

15. Click Finish to start the job for creating a zone.

Creating zones with 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.

Cloning a Zone

Use this procedure 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 . You can clone either a greenfield or brownfield zone.

To Clone a Zone

1. Select the zone you want to copy from the Assets section in the Navigation pane.

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, timezone, 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. If 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 sytem 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 already defined in the source zone.

9. Accept or change the storage library. The library can be either a NAS storage or Fibre Channel library. If the library is NAS storage, specify the virtual disk name and size of the disk. If the library is Fibre Channel, 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. If the network has a defined IP range, the Auto Allocate IP option is displayed with the zone's IP address from the range.If the network does not have a defined IP range, the Use Static IP option is displayed. Type an IP address for the zone in the Zone IP field.

    ■ 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. If 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.

        If 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 type 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.

When the job successfully completes, the zone is displayed in the Assets section.

Migrating a Non-Global Zone

In 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. 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 Enterprise Manager Ops Center. See Migrating a Zone to a Different Machine

- Migrating a non-global zone that has dependencies that are not managed by Enterprise Manager Ops Center. This process uses a script to perform the migration. See Using a Script to Migrate a Zone With Dependencies

You can migrate a greenfield zone. If you want to migrate a brownfield zone that has been adopted to become a greenfield zone, see Using a Script to Migrate a Zone With Dependencies to migrate the additional inherited file systems.

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 12–1 diagrams the steps for a P2V conversion.

Figure 12–1   Workflow for Migrating an Existing Oracle Solaris System Into a Zone
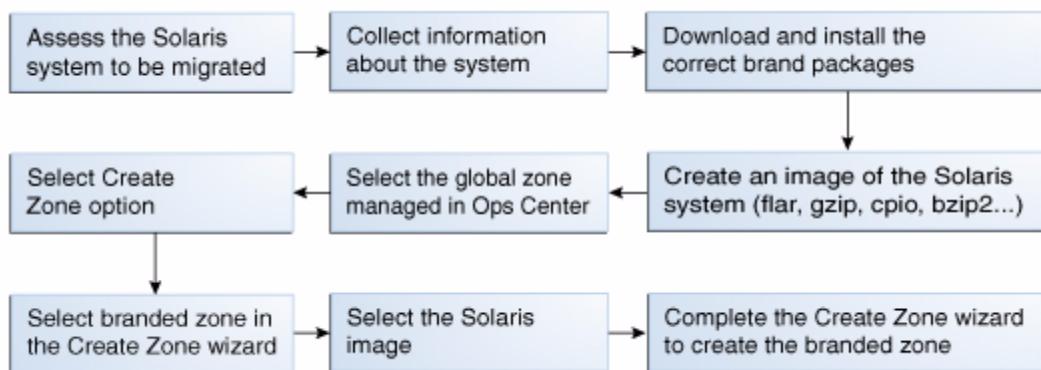
Before You Begin

- Assess the system to be migrated and collect information.

  See Addtional Resources 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 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

See Additional Resources for instructions to download and install the correct brand packages.

To Migrate a Physical Solaris System Into a Zone

1. Create an image of the Oracle Solaris system that you want to migrate.

2. Import the image into Enterprise Manager Ops Center.

   See Importing Images to import the image.

3. Select the destination global zone from the Assets section of the Navigation pane.

4. Click Create Zone in the Actions pane. The Create Zone wizard is displayed.

5. Select Branded Zone for the type of the zone.

6. Select the image that you have created in Step 1 from the Images list.

7. (Optional) Specify the HostID if you want to run applications that are associated with a specific host ID.

8. (Optional) Specify the machine type as Sun4U or Sun4V.

9. Complete the Create Zone wizard, according to To Create a Zone.

Migrating a Zone to a Different Machine

In Enterprise Manager Ops Center, you can migrate a non-global zone from its source global zone to a target global zone running on a different machine. 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 zone.

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 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. If 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. The patches that must be backed out of the zone before the update are also listed. If the target global zone does not support backout-on-attach capability, 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, if you change any other configuration, A warning message indicating that if you 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, Enterprise Manager Ops Center checks all global zones for compatibility with the source global zone. If there are compatible global zones, the Migrate Zone wizard is displayed. Otherwise, the following pop-up window is displayed.

*Figure 12–2   Message Displayed When No Global Zones are Compatible*
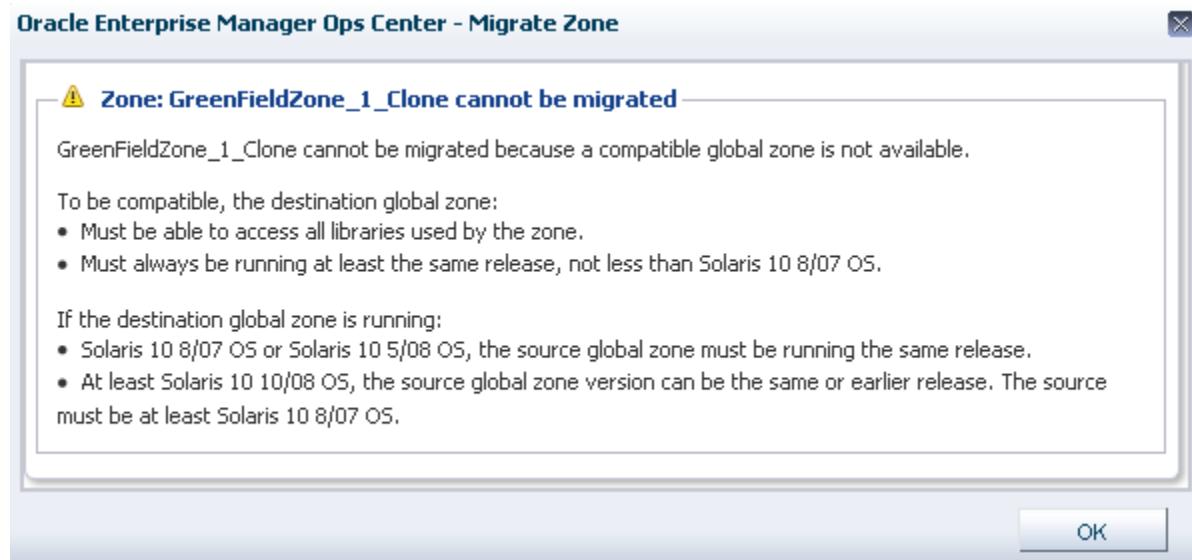
Figure 12–2 shows the compatible global zones, depending on the operating system release.

Table 12–2

Source Global Zone	Compatible Target Global Zone
Solaris 10 8/07	Solaris 10 8/07 Solaris 10 10/08 Solaris 10 5/09 Solaris 10 10/09
Solaris 10 5/08	Solaris 10 5/08 Solaris 10 10/08 Solaris 10 5/09 Solaris 10 10/09
Solaris 10 10/08	Solaris 10 10/08 Solaris 10 5/09 Solaris 10 10/09
Solaris 10 5/09	Solaris 10 5/09 Solaris 10 10/09

Before You Begin

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 network storage library.

- The non-global zone is in the running state.

- The source and target global zones and non-global zones were created through Enterprise Manager Ops Center.

To Migrate a Zone

1. Select a zone from the Assets section in the Navigation pane.

2. Click Migrate Zone from the Actions pane. A warning message indicating that if you change the zone configuration except for Name, Description, Tags, and NIC details , you must execute the `sys-unconfigure` command on the zone. Also, if 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 the target global zone from the list of compatible global zones. Each entry includes the number of non-global zones that are already running, total CPUs, and available dedicated CPUs. Click Next.

5. Review the migration test result. Continue the migration based on the following conditions on the target global zone:

   - If the target global zone does not support update-on-attach capability, continue to migrate the zone without updating the zone's patches and packages.

   - If the target global zone supports update-on-attach capability, update the zone's patches and packages to match the destination global zone.

6. (Optional) To save the migration test result, click Save Test Result As.

7. (Optional) To change the target global zone, click Select New Destination.

8. Change or accept the new name for the migrated zone. The new name is the current zone name with an integer appended. If the Zone Name is displayed in red, the zone name already exists in the target global zone and you must provide a unique name. Click Next.

9. 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 select Do Not Connect to that network.

   a. Select the network which is in yellow. The network interface details are displayed under Network Interface.

   b. Select a new network interface from the Network list or select Do Not Connect. The selected network can be either shared or dedicated. Enter the required network interface information for the selected network.

   Click Next.

10. Review the summary of the migration job. If there are no patches and packages to be backed out, continue with the migration. If there are patches or patches that must be removed, a warning is displayed .

- If the target global zone supports backout-on-attach capability, a warning is displayed that the a list of patches have been removed or downgraded.

- If 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. job.

If the target global zone does not support backout-on-attach capability, you must remove or downgrade the patches and packages manually before you can continue with the migration. Use the New Update OS Job action to back out the patches and packages, using this procedure:

To Back Out Patches and Packages

1. Create an update profile with the list of patches and packages that must be backed out.

  See Creating an Os Update Job.

2. Select the zone.

3. Create a New Update OS Job with the OS Update Profile.

The OS Update job removes or downgrades the patches and packages of the zone. Now you can repeat the migration job on the zone.

Using a Script to Migrate a Zone With Dependencies

You can migrate zones even if the zone has dependencies that are not managed by Enterprise Manager Ops Center, for example, if the zone uses storage that is not part of the zpool. You can use scripts to extend the zone migration feature to include the migration of the dependencies such as storage or other resources that is not managed by Enterprise Manager Ops Center.

If you have a adopted a brownfield zone with additional inherited file system, you can lose the additional file systems during migration. To avoid this, use this script to migrate the file systems.

You can develop your own scripts to migrate the dependencies and place it 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 if the script is found only on one of the global zone.

- 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 six times when the migration job is executed.

Table 12–3   Script Call to Environmental Variables

Call to the Script	Environmental Variables
On the source global zone, to verify the script existence and user dependencie	OEMOC_ZONENAME=source zonename OEMOC_PHASE=VERIFY OEMOC_OPERATION=MIGRATION
On the target global zone, to verify the script existence and user dependencies.	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=VERIFY OEMOC_OPERATION=MIGRATION
On the source global zone, before the zone is shutdown	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_RUNNING OEMOC_OPERATION=MIGRATION
On the source global zone, after the zone is shutdown.	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION
On the target global zone, before the zone is started.	OEMOC_ZONENAME=source zonename OEMOC_PHASE=POSTOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION
On the target global zone, after the zone is started	OEMOC_ZONENAME=source zonename OEMOC_PHASE=POSTOPERATION_RUNNING OEMOC_OPERATION=MIGRATION

During rollback phase, the same sequence is executed in reverse order. The environment variables values change accordingly :

- PREOPERATION becomes PREROLLBACK

- POSTOPERATION becomes POSTROLLBACK

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 Example 12–1, "Sample Script" for migrating the additional inherited file systems of an adopted zone. The sample script is based on the assumption that `/opt/` filesystem is the additional inherited file system. The script re-configures `/opt/` filesystem on the target global zone before rebooting the zone. The script re-configures the user inherited file systems during POSTOPERATION_NOTRUNNING phase

Example 12–1   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

If 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 Enterprise Manager Ops Center cannot get access to the information in the global zone and therefore cannot perform any compatibility checks. The procedure to move 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 Enterprise Manager Ops Center..

See the Enterprise Manager Ops Center Administration 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.

- Verify that the target global zone has access to he 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 is used. In the Enterprise Controller command prompt, enter the following command:

  ```
  export JAVA_HOME=/usr/jdk/latest
  ```

To Recover Greenfield Zones

1. Connect to the 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
    --------------------------------------------------------------------------------
    ----------------------------------------------------------------
    gzhost36             zone                 OK           False
    com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost
    gzhost44             zone                 OK           True
    com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost-4
    ```

5. List the zones that are running in the source global zone. For example, if 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_hosts -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, if 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> -D <target global
zone>

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 .

13
Cluster Monitoring

Enterprise Manager Ops Center can manage and monitor Oracle Solaris Clusters and their nodes. Enterprise Manager Ops Center manages Oracle Solaris Clusters in the same way as other assets. Cluster assets are of the following types:

- Cluster: Software that supports redundant services among the nodes in the cluster so that there is no single point of failure.

- Cluster Node: A server that is a fully-voting member of the cluster. A zone acts as a virtual node.

- Zone Cluster Group: A cluster of non-global zones that are hosted on a global cluster. At most, one node of a zone cluster can exist on a global cluster node. Each node in the zone cluster has the same name as the zone cluster.

- Zone Cluster Node: A non-global zone that is a member of a zone cluster group.

You can monitor, provision, and update the software on a cluster's nodes.

After the Oracle Solaris Clusters have been discovered, the Asset tree represents the cluster architecture. A cluster can contain the following components:

- Global Nodes

- Zone Clusters and their zone nodes

To see the default profile for cluster monitoring, see Displaying A List Of Monitoring Profiles. To see the profiles for updating a cluster, see About Cluster Profiles.

Terms Used in Oracle Solaris Clusters

Zone node A non-global zone hosted on a cluster node and is in the nodelist of a resource group. It is active in the resource group but it is not a full cluster node.

Nodelist Specifies the nodes of any type that the resource group can run on.

global cluster A cluster of Oracle Solaris 10 nodes and the host of multiple zone clusters.

resource group A collection of resources and policies that act as a data service

Public interconnect Network connections for data services, always uses IPMP

Private interconnect Network connections for cluster management and configuration

Dual Partition upgrade Upgrades one half of the cluster's nodes, then the other half, so that data service is not interrupted.

CMAS The cluster software's agent. Enterprise Manager Ops Center uses this agent to discover and manage cluster nodes.

Viewing Cluster Details

When you select an Oracle Solaris Cluster, the dashboard shows its components. To see the attributes of a cluster, hover the cursor to see ID, version, and package status. In addition, the status of its quorum service is displayed: device, version, possible votes, and current votes.

When you select the specific cluster, the dashboard shows its components, for example, global nodes and zone clusters. The icons for each type display a number, which indicates the number of nodes or the number of zones.

When you select the Global Node type, the dashboard shows each global node. To see the attributes of a global node, hover the cursor to see OS version, host name, IP address, current votes, and possible votes.

The Alerts tab displays the current alert status for the cluster. The Serial Console tab gives access to the cluster's systems.

14

Plans

Enterprise Manager Ops Center operations are executed in a consistent way when you use plans and the profiles that they contain. All users can apply a plan. Your ability to modify plans depends on your role and whether you are the author of a profile or a plan.

- Deployment plans contain the sequence of profiles that deploy operating systems and firmware, the specification for executing each step, and the resources for each step such as network addresses, host names and so on.

- Operational plans contain scripts and utilities. You can use these to fix common problems and build a custom problems knowledge base for your environment.

See DEPLOYMENT PLANS and OPERATIONAL PLANS for information about creating plans and operational profiles.

Viewing Plans

You can view the following details of a deployment or operational plan:

- Details of the plan such as name, description, version, steps in the plan, target type and last modification.

- Required resources for the execution of the plan.

- The result of executing the plan.

- The other plans that use the plan in their steps.

- Version history of the plan.

- The list of jobs that were executed using this plan.

To View a Plan

1. Select Plan Management section in the Navigation pane.

2. Select a plan from either Operation Plans or the Deployment Plans tree.

3. Expand the selected plan. All the plans for this type are listed.

4. Select a plan from the list.

5. Select one of the tabs in the center pane to view information about the plan:

   - Details

   - Results

   - Referrers

- Version History
- Jobs

Details Tab

The Details tab provide an overall view of the selected plan and its output. Each step of the plan is displayed with the associated profile or plan, result of the task, and any additional parameters . The Details tab displays the following general information about the plan:

- Name
- Version
- Description
- Target type
- Creation date
- Last modified date
- Template used
- Referrers

Results Tab

The Results tab specifes the number and types of assets that executed the plan. For each task of the plan, the expected result is defined. Some steps in the plan use the result from the previous step as the target for the next step.

Referrers Tab

The Referrers tab list the other plans that use the selected plan to complete their tasks..

Version History Tab

Each time a plan is modified, the new version is recored on the Version History tab. You can select a version of the plan and view the details of the plans, and apply the plan. To delete any version of the plan, you must have the appropriate authorization.

Jobs Tab

The Jobs tab shows the list of jobs that were executed using the plan. It displays the status of the job, version of the plan used, the owner of the job, and the time taken to complete the job. You can select a job and then select the following options:

- Re-Apply Deployment Plan – This option allows you to select the asset on which you want to run the plan and the corresponding plan steps to be executed.
- Show Job Details – This option opens the Job Details window, which displays the job information and composition.

Applying a Deployment Plan

A Deployment Plan is a plan that installs or updates an operating system or firmware for an asset type.

To Apply Deployment Plan:

1. ExpandPlan Management in the Navigation pane, then click Deployment Plans.

2. Select the deployment plan you want to use.

3. Click Apply Deployment Plan in the Actions pane. The Select Assets window is displayed.

4. Select a target from the list of available targets and click Add to Target List. You can add more than one target to apply the plan.

5. Select whether you want to apply the plan with minimal user interaction or override any profile values. Click Next.

6. Specify the network resources for each selected targets.

7. Enter the IP address for each selected target.

8. Review your selections and then click Next.

9. Click Apply. A job starts to deploy software on the selected targets according to the plan.

Applying an Operational Plan

An Operational Plan contains a script that defines an operation to be performed on a managed resource or group of resources.

To Apply an Operational Plan

1. Expand Plan Management in the Navigation pane, then click Operational Plans.

2. In the center pane, highlight the plan you want to apply.

3. Click the Apply Operational Plan icon in the center pane. The Select Assets - Operational Plans page appears.

4. Select the assets or groups that will be the targets of the Operational Plan, then click Add to Target List.

5. Specify how to apply the plan. You can apply the plan with minimal interaction, or you can choose the option to override any profile values. Click Next.

6. Specify any additional environment variables. You can apply the variables to all targets, or click the Targets tab to select specific targets.

7. Review your selections and then click Next.

8. Click Apply. A job starts to deploy software on the selected targets according to the plan.

15

Enterprise Manager Ops Center and Grid Control

The Oracle Enterprise Manager Ops Center and Enterprise Manager Grid Control applications monitor assets, or targets, in different sections of the data center stack:

- Oracle Enterprise Manager Grid Control manages applications, middleware, databases, and operating systems

- Oracle Enterprise Manager Ops Center manages operating systems, virtual machines, servers, and storage devices

Use plug-in sofware that configures Oracle Enterprise Manager Ops Center and Grid Control to work together to monitor the entire asset stack in your data center. When you configure the applications to work together, each application shares information with the other application:

- Each user interface includes information from the other application. Resources managed by Oracle Enterprise Manager Ops Center are visible on the Oracle Enterprise Manager Grid Control user interface and resources managed by Oracle Enterprise Manager Grid Control are visible on the Oracle Enterprise Manager Ops Center user interface.

- Alerts generated in one application are reported in the other application.

See the Oracle Enterprise Manager documentation library for more information about Oracle Enterprise Manager Grid Control.

> **Note:** If you installed and configured the connector for Ops Center 2.5 and Oracle Enterprise Manager, uninstall the connector before upgrading to Oracle Enterprise Manager Ops Center. Beginning with Oracle Enterprise Manager Ops Center 11g Release 1, the connector is not needed for integration.

Oracle Enterprise Manager Ops Center and Grid Control monitor assets, or targets, in different sections of the data center IT stack. Enterprise Manager Grid Control manages the applications, middleware, databases, and operating systems at the top of the stack and Enterprise Manager Ops Center manages operating systems, virtual machines, servers, and storage devices at the bottom of the stack.

Figure 15–1   Scope of Product Operations in IT Stack

You can link the monitoring functionality of both products. You can install Enterprise Manager Ops Center and Grid Control in your data center and then configure the applications to communicate with each other.Even though both products monitor operating system status, they do not share overlapping information. The OS status and alerts from one application do not display in the other application's user interface. Because there is no overlap, the OS monitoring parameters specific to Oracle Solaris Zones and Oracle VM Server for SPARC appear in Enterprise Manager Grid Control.

See the *Oracle Enterprise Manager System Monitoring Plug-In for Oracle Enterprise Manager Ops Center Guide* for the steps to install the plug-in software.

Configuring the Applications

Install Enterprise Manager Ops Center and Grid Control in your data center and then configure the applications to communicate with each other. You must have the following access:

- Enterprise Manager Grid Control user

or

- DB user on the Oracle Enterprise Manager Grid Control repository with select permissions for the management repository and Enterprise Configuration Management view

To display basic application, middleware, and database monitoring information for a target in Enterprise Manager Ops Center, configure the Enterprise Controller to communicate with the Enterprise Manager Grid Control repository. The monitoring information appears in the Targets tab in the Asset view. Click a target to view basic information. To view more details, double-click the problem and launch the target home page in Enterprise Manager Grid Control.

To display information collected by Enterprise Manager Ops Center in the Enterprise Manager Grid Control console, you must install and configure a plug-in application. The plug-in creates a new Infrastructure Stack target type in the console. Select

Infrastructure Stack targets to display information collected by Enterprise Manager Ops Center. To get more details about the problem, you can drill down into the Enterprise Manager Ops Center UI.

The following reports are available with the Infrastructure Stack plug-in:

- Topology report – Provides stack layers report

- Hardware sensors report – Provides hardware sensors details based on the service processor, including current value and status, and thresholds

- Configuration report – Provides basic configuration of each layer within the stack

To view reports, go to the target's home page, then click the Reports tab.

See the *Enterprise Manager Ops Center Administration Guide* for information about configuring the Enterprise Controller to communicate with the Enterprise Manager Grid Control repository.

Viewing Enterprise Manager Grid Control Alerts in Enterprise Manager Ops Center

When configured, you can view the status of targets that are being monitored by Oracle Enterprise Manager Grid Control in the Oracle Enterprise Manager Ops Center UI. You can drill down for greater detail by double-clicking on a Grid Control generated problem to launch the Oracle Enterprise Manager Grid Control console. You can also launch the console from Administration in the Navigation pane.

Both applications monitor the OS, but only Oracle Enterprise Manager Ops Center monitoring is displayed in the UI.

Viewing Oracle Enterprise Manager Grid Control Alerts

1. Expand Assets in the the Navigation pane.

2. Click an OS that has the Oracle Enterprise Manager Grid Control Agent installed. The Dashboard page is displayed. The OS and the members above the OS are monitored by Ops Center. Applications below the OS are monitored by Oracle Enterprise Manager Grid Control, as shown in the following graphic. If you have a Solaris OS installed with Oracle Containers or Oracle VM Server for SPARC, the zones or guests are monitored by Ops Center and will appear below the OS.

Figure 15–2 Dashboard for an OS Asset
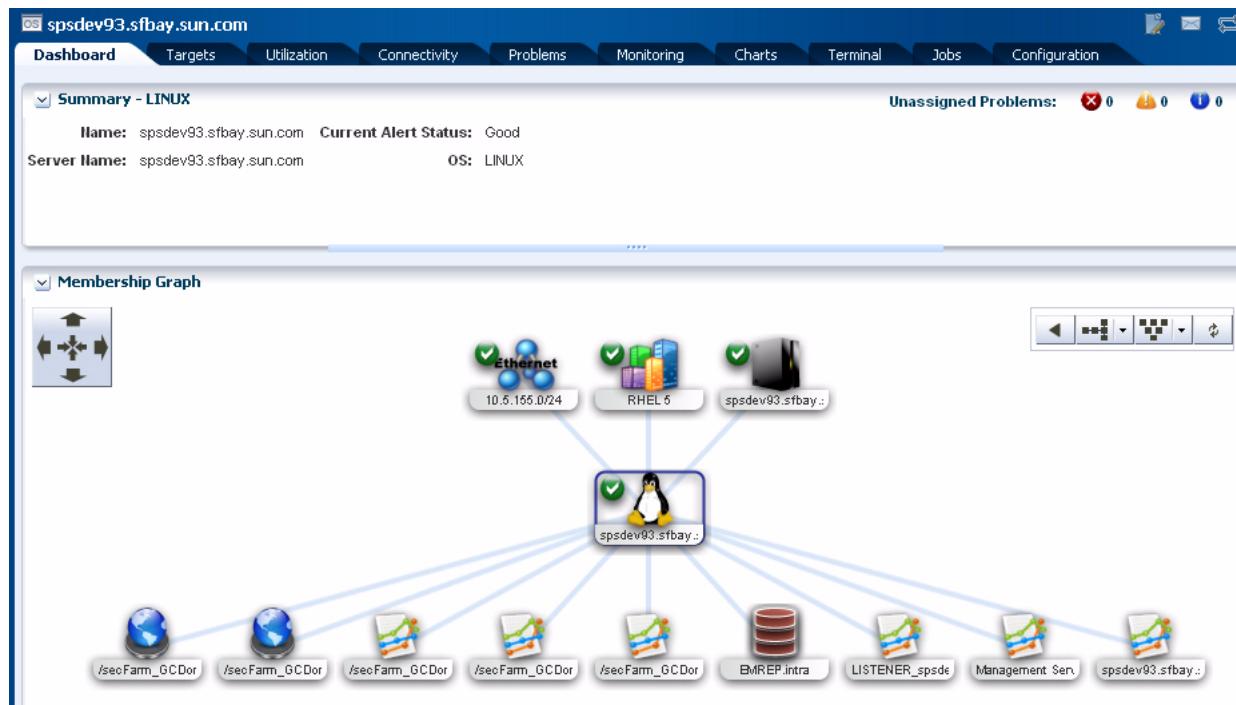
3. Click the Targets Tab to view Grid Control target information, including status alerts, and availability.

Figure 15–3   Targets Tab for an OS Asset

4. To view details about a specific target, either double-click the target or highlight the target, then click the Go To Grid Control icon.

Figure 15–4 is an example of a Database Instance target.

Figure 15–4   Details of a Grid Control Target

5.   For Alert history, you can choose from the following time frames in the drop-down list:

   ■   Current

   ■   Past 24 hours

   ■   Past 7 day history

   ■   Past 31 day history

For more information about Oracle Enterprise Manager Grid Control, see Additional Resources.

Viewing Enterprise Manager Ops Center Details in Enterprise Manager Grid Control

Before you can view details collected from the Enterprise Manager Ops Center in the Grid Control console, you must deploy and configure an infrastructure stack plug-in on your Oracle Enterprise Manager Grid Control Agents. Deploying the plug-in

creates a new Infrastructure Stack target type in the Oracle Enterprise Manager Grid Control console. To begin sharing monitoring information, you must define the targets.

Once configured, you can search for Infrastructure Stack target types to display information collected by the Oracle Enterprise Manager monitoring functionality and you can perform the following actions:

- View hardware alerts
- Run sensor reports
- Run OS update and baseline comparisons

See the Oracle Enterprise Manager documentation for information about displaying monitoring information by target type.