

Oracle® Enterprise Manager Ops Center

Feature Reference Guide

12c Release 1 (12.1.4.0.0)

E27511-08

June 2013

E27511-08

Copyright © 2007, 2013, Oracle and/or its affiliates. All rights reserved.

Primary Authors: Barbara Higgins, Laura Hartman, Shanthi Srinivasan, Owen Allen, Uma Shankar, Richa Agarwala, Luis Ramos Veja, Karen Orozco Sanchez

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xxvii
Audience.....	xxvii
Documentation Accessibility	xxvii
Related Documents	xxvii
Conventions	xxvii

Part I Getting Started

1 Introduction

Oracle Enterprise Manager Ops Center in the Data Center.....	1-1
Roles and Connection Modes	1-1
User Preferences and Roles	1-2
Start In and Start Page Preferences.....	1-2
Time Preferences	1-2
Membership Graph Preferences	1-2
User Preferences Summary.....	1-3
Preferences By Role.....	1-3
How To Use This Collection	1-3

2 Asset Management

Introduction to Asset Management	2-1
Roles for Asset Management	2-1
Actions for Asset Management.....	2-2
Location of Asset Management Information in the User Interface	2-2
Discovering and Managing Assets	2-3
Declaring Servers for OS Provisioning	2-4
Declaring Servers for Service Processor Configuration	2-5
Adding Assets Using a Discovery Profile	2-6
Finding Assets	2-6
Creating a Discovery Profile.....	2-7
Editing a Discovery Profile	2-8
Copying a Discovery Profile.....	2-8
Deleting a Discovery Profile.....	2-9
Installing Agent Controllers From the Command Line	2-9
Using Management Credentials	2-16

Upgrading Management Credentials From a Prior Version	2-16
Updating Management Credentials	2-16
Creating Management Credentials	2-16
Using Custom SSH Keys for OS Discovery	2-16
Editing Management Credentials.....	2-18
Copying Management Credentials	2-18
Deleting Management Credentials.....	2-18
Editing Asset Attributes	2-18
Using Access Points	2-18
Deleting Assets	2-19
Special Discovery and Management Procedures	2-19
Windows Systems.....	2-20
Servers	2-21
Sun ZFS Storage Appliance	2-21
Oracle Solaris Cluster	2-28
Using Tags	2-29
Adding Tags.....	2-29
Viewing Tags	2-29
Grouping Assets Using Tags	2-29
Deleting Tags	2-29
Using Groups	2-29
Types of Groups	2-30
System Groups	2-30
User-Defined Groups	2-30
Viewing Group Data.....	2-31
Creating a Group.....	2-31
Editing a Group	2-32
Adding Assets to a Group	2-32
Removing Assets from a Group.....	2-33
Moving Assets to a Group	2-33
Moving a Group	2-33
Deleting a Group	2-33
Related Resources for Asset Management	2-33

3 Jobs

Introduction to Jobs	3-1
Roles for Job Management	3-3
Job Management.....	3-4
Viewing Jobs	3-4
Viewing All Jobs and Jobs With Specific Status	3-4
Viewing Job Details.....	3-5
Monitoring Jobs for an Asset.....	3-6
Search for Jobs	3-7
Taking Actions on a Job	3-8
Answering Questions for a Job	3-8
Stopping a Job	3-8
Re-running a Completed Job	3-8

Re-running a Job on Failed Targets	3-8
Copying a Job.....	3-9
Deleting a Job.....	3-9
Debugging a Job Using the OCDoctor	3-9
Changing the Job Properties	3-9
Events for Job Management	3-10
Related Resources for Job Management	3-10

Part II Configure

4 Monitoring Rules and Policies

Introduction to Monitoring	4-1
Roles for Monitoring	4-1
Actions for Monitoring Rules and Policies	4-2
Location of Monitoring Rules and Policies in the User Interface	4-2
Monitoring Rules	4-3
User-Defined Rule Parameters.....	4-4
Enabled and Active Rules	4-4
Editing Monitoring Rules	4-5
Using Historical Data to Determine Threshold Limits for a Specific Asset	4-6
Adding Monitoring Rules.....	4-7
Monitoring Policies.....	4-8
Default Monitoring Policies.....	4-9
Available Monitoring Policies	4-9
Monitoring Policy Details	4-10
Applying a Monitoring Policy to a Group	4-12
Editing the Monitoring Configuration of an Asset Bound to a Monitoring Policy	4-12
Creating a Group of Assets According to Monitoring Policy.....	4-12
Extracting a Monitoring Policy	4-13
Adding or Copying a Monitoring Policy.....	4-14
Deleting a Monitoring Policy	4-14
Related Resources for Monitoring Rules and Policies	4-15

5 Software Libraries

Introduction to Software Libraries and Repositories.....	5-1
Roles for Software Libraries	5-1
Actions for Software Libraries.....	5-2
Location of Software Library Information in the User Interface.....	5-2
Knowledge Base and Parent Repository.....	5-3
EC Library for Oracle Solaris 10, 9, and 8 and Oracle Linux OS.....	5-3
Publishers and Parent of the Oracle Solaris 11 Repository	5-3
Using Software Libraries	5-4
Viewing the Contents of a Software Library	5-4
Creating a Software Library	5-5
Libraries for Oracle Solaris 11	5-5
Oracle Solaris 11 Software Update Library	5-6

Summary of Oracle Solaris 11 Software Update Library	5-6
Content of the Oracle Solaris 11 Software Update Library	5-7
Configuring Parent Repositories	5-7
Using an Alternate IPS Repository	5-8
Adding Content.....	5-8
Deleting Oracle Solaris 11 Software Update Library	5-8
Libraries for Oracle Solaris 10, 9, 8 and Linux	5-8
Images.....	5-9
Images for OS Provisioning.....	5-10
Images for Firmware Updates.....	5-10
Uploading or Importing Images	5-11
Importing an Image	5-12
Uploading an Image	5-13
Uploading Firmware Images.....	5-13
Working with Firmware for Power Distribution Units.....	5-13
Uploading Firmware With Metadata.....	5-14
Uploading Firmware Without Metadata.....	5-14
Creating a Firmware Profile for PDU Firmware Updates	5-14
Local Content	5-14
Local Categories	5-15
Uploading a Local Action	5-15
Uploading a Local Software Package.....	5-16
Uploading a Local Configuration File.....	5-17
Uploading Software in Bulk	5-18
Uploading Local Software in Bulk	5-18
Viewing Results of a Bulk Upload Operation.....	5-19
Using Local Content	5-19
Editing Local Content.....	5-20
Deleting Local Content.....	5-20
Backing Up Images and Local Content	5-21
Related Resources for Software Libraries	5-21

6 Storage Libraries

Introduction to Storage Management	6-1
Roles for Storage Management.....	6-1
Actions for Storage Management.....	6-2
Location of Storage Information in the User Interface	6-2
Storage Libraries.....	6-3
Network Attached Storage (NAS) Storage Libraries	6-4
Storage Libraries for Server Pools.....	6-5
Storage Libraries for a Virtual Datacenter.....	6-5
Storage Libraries for Oracle Solaris Zones	6-6
Storage Libraries for Oracle VM Server for SPARC.....	6-6
Storage Libraries and Repositories for Oracle VM Server for x86	6-6
Types of Storage for Libraries	6-8
File Systems Libraries	6-8
Viewing Local Libraries	6-8

Editing the Attributes of a Local Library.....	6-8
Creating a Local Library	6-8
Deleting a Local Library.....	6-9
Block Storage.....	6-9
Dynamic Block Storage	6-9
Static Block Storage.....	6-10
Selecting LUNs For the Block Storage Library	6-10
Adding Capacity to Dynamic Block Storage Libraries.....	6-11
Adding Storage to a Logical Domain.....	6-12
Create a LUN	6-13
Clone a LUN	6-14
Oracle VM Storage Connect Plug-ins	6-15
Storage Hardware	6-15
RAID Controller	6-15
NAS Storage Appliance.....	6-16
Setting Up a Share on an NFS Server	6-16
Setting Up an NFS Client.....	6-16
Creating a NAS Storage Library	6-17
File Server	6-18
Storage Array	6-18
Storage Server: Sun ZFS Storage Appliance.....	6-18
Managing a Sun ZFS Storage Appliance	6-18
Displaying Volume Capacity	6-19
Auto Service Requests for the Sun ZFS Storage Appliance.....	6-19
Provisioning and Updating a Sun ZFS Storage Appliance.....	6-19
Exadata Storage Server.....	6-20
Storage Profiles	6-20
High Availability for Storage Resources	6-20
Related Resources for Storage Management	6-20

7 Networks

Introduction to Network Management	7-1
Roles Required for Network Management	7-2
Actions for a Network and Network Domain	7-2
Location of Network Information in the User Interface	7-3
Fabrics	7-3
Network Domains	7-4
Default Network Domain	7-5
User-Defined Network Domains	7-5
Editing Attributes of a User-Defined Network Domain	7-6
Networks	7-6
Requirements for a Network	7-7
Public Networks and Private Networks	7-7
Networks for Server Pools.....	7-8
Networks for Oracle Solaris Zones.....	7-8
Networks for Oracle VM Servers.....	7-9
Networks for Virtual Datacenters.....	7-9

Bandwidth Management.....	7-9
Managing the Bandwidth Flows for a Data Link.....	7-9
Properties of Bandwidth Flow	7-10
IP Multipathing Groups	7-10
Creating IPMP Groups	7-11
Link Aggregation or Bond	7-13
Creating Link Aggregation.....	7-14
Properties of a Network	7-14
IPv4 and IPv6 Protocols	7-15
Routing Mode	7-15
Static Route for the Network	7-15
Address Allocation Method	7-16
Maximum Transmission Unit (MTU)	7-16
VLAN and VLAN Tags.....	7-17
Network Utilization	7-17
Network Connectivity	7-17
Network Hardware	7-18
Network Interface Card (NIC)	7-18
Network Switches	7-18
Network Profiles	7-19
Oracle Enterprise Manager Ops Center's Networks	7-19
Network Switch Configuration.....	7-19
Separate Management, Provisioning, and Data Networks.....	7-19
Combined Management and Provisioning Network and a Separate Data Network	7-21
Combined Provisioning and Data Network and a Separate Management Network	7-23
Combined Provisioning, Data, and Management Network	7-24
Firewall Rules	7-26
Network Port Requirements and Protocols	7-26
Related Resources for Networks	7-29

8 Plans and Profiles

Introduction to Plans and Profiles	8-1
Roles for Plans and Profiles	8-2
Actions for Deployment Plans.....	8-2
Editing a Profile	8-2
Editing a Deployment Plan.....	8-3
Copying a Deployment Plan	8-3
Deploying a Plan	8-4
Deleting a Deployment Plan.....	8-4
Location of Profiles and Plans in the User Interface	8-5
Overview of Version Control.....	8-5
Operational Profiles and Plans.....	8-5
Actions for Operational Profiles and Plans	8-6
Creating an Operational Profile and Plan	8-6
Editing an Operational Profile or Plan.....	8-7
Copying an Operational Profile	8-7
Copying an Operational Plan.....	8-8

Deleting an Operational Profile or Plan.....	8-8
Profiles and Policies	8-8
About Policies.....	8-9
About Profiles.....	8-9
Deployment Plans	8-10
About Simple Deployment Plans.....	8-11
About Multi-Step Deployment Plans.....	8-12
About Complex Deployment Plans.....	8-12
Related Resources for Plan Management	8-13

Part III Operate and Maintain

9 Incidents

Introduction to Incident Lifecycle Management	9-1
Roles for Incident Management	9-2
Actions for Incident Management	9-3
Actions for Incidents.....	9-3
Actions for Service Requests.....	9-3
Actions for Incidents Knowledge Base.....	9-4
Location of Incident and Service Request Information in the User Interface	9-4
Incidents	9-4
About Incident Severity Badges.....	9-5
Message Center	9-6
Incidents Dashboards in the Message Center.....	9-8
Annotations	9-10
About Annotation Views.....	9-11
Incidents Knowledge Base	9-12
Annotations in the Incidents Knowledge Base	9-13
Incident Lifecycle Management	9-13
Methods of Incident Management.....	9-14
Viewing Unresolved Incidents.....	9-15
Viewing Incident Details.....	9-15
Assigning an Incident.....	9-15
Acknowledging Incidents.....	9-16
Adding an Annotation.....	9-17
Displaying Annotations.....	9-18
Viewing Comments.....	9-18
Taking Action on a Incident.....	9-18
Maintenance Mode.....	9-19
Marking an Incident Repaired.....	9-20
Closing an Incident.....	9-20
Oracle Services and Service Requests	9-21
Requirements for Oracle Services.....	9-22
Viewing Contract and Warranty Information.....	9-22
Viewing Service Requests.....	9-23
Filing a Service Request.....	9-24

Automated Service Requests	9-24
Related Resources for Incident Management	9-24

10 Reports

Introduction to Reports	10-1
Roles for Reports	10-2
Actions for Reports	10-2
Managing Report Templates	10-2
Generating a Report from a Report Template	10-3
Editing a Report Template.....	10-3
Deleting a Report Template.....	10-3
Viewing and Exporting Report Results	10-3
Viewing the Report Result.....	10-3
Saving the Report Result.....	10-4
Location of Report Information in the User Interface	10-4
Creating a Firmware Report	10-5
Creating an Operating System Report	10-6
Update Compliance Reports	10-7
Linux and Oracle Solaris OS Update Reports	10-7
Creating a Change History Report	10-8
Creating a Baseline Analysis Report	10-9
Creating a Profile Analysis Report	10-12
Creating a Recommended Software Configuration Report.....	10-15
Creating an Oracle Solaris Update Compliance Report	10-16
Creating an Incident Compliance Report	10-17
Incident Compliance Report for Oracle Solaris or Linux.....	10-17
Creating an Incident Compliance Report for Microsoft Windows.....	10-18
Creating a Host Compliance Report	10-19
Host Compliance Report for Oracle Solaris or Linux.....	10-20
Creating Host Compliance Report for Microsoft Windows	10-20
Creating a Common Vulnerability and Exposure (CVE) Report.....	10-21
Creating a System Catalog Report.....	10-23
Creating Additional Operating System Reports	10-23
Creating a Distribution Update Report	10-24
Creating a Service Pack Compliance Report.....	10-24
Creating a Package Compliance Report	10-25
Creating Incident Reports	10-25
Creating System Information Reports	10-27
Creating Exalogic Reports	10-28
Related Resources for Reports	10-29

11 Hardware

Introduction to Managing Hardware Assets	11-1
Configuring Hardware Assets	11-1
Monitoring Hardware Assets.....	11-2
Maintaining Hardware Assets	11-2
Roles for Hardware Management	11-2

Actions for Hardware Management	11-3
Location of Hardware Information in the User Interface.....	11-3
Profiles for Hardware Management	11-4
Hardware Resource Profiles.....	11-5
Firmware Provisioning Profiles	11-5
Configuring the Service Processor.....	11-5
Creating a Service Processor Configuration Profile and Plan	11-5
Creating a BIOS Configuration Profile and Plan.....	11-6
Creating a Snapshot of a Service Processor Configuration.....	11-8
Applying a Snapshot of a Server Processor Configuration	11-8
Configuring a RAID Controller	11-9
Configuring a Dynamic System Domain	11-9
Configuring a Rack and Placing Components	11-10
Hardware Monitoring.....	11-11
Hardware Status.....	11-11
Groups of Hardware Assets	11-11
Connectivity Status	11-11
Service Processor Details.....	11-12
RAID Controller Details.....	11-12
Sun ZFS Storage Appliance Details	11-12
Server Details.....	11-13
Chassis Details	11-14
M-Series Details	11-14
Switch Details	11-16
Rack Details	11-17
PDU Details.....	11-17
Oracle Solaris Cluster Details.....	11-18
Oracle Exalogic Details.....	11-18
Monitoring Power Utilization	11-18
Energy Tab	11-18
Charts Tab	11-20
Maintaining Hardware Assets	11-20
Setting and Changing the Power Policy	11-20
Replacing a Failed Power Distribution Unit in a Rack	11-21
Installing and Upgrading Oracle Solaris Cluster	11-22
Firmware Provisioning.....	11-22
Firmware Profiles.....	11-22
Firmware Compliance Reports	11-23
Updating Firmware	11-23
Related Resources for Hardware Management.....	11-25

12 Operating Systems

Introduction to Operating Systems	12-1
Roles for Operating Systems.....	12-3
Actions for Operating Systems.....	12-4
Location of Operating System Information in the User Interface.....	12-5
Viewing Operating Systems	12-7

Operating System Profiles	12-8
Using Agent Management for Operating Systems	12-9
Monitoring Operating Systems	12-11
Using Analytics	12-12
Displaying Analytics Information	12-13
Displaying the Analytics Summary.....	12-14
Displaying the Processes View	12-15
Displaying the Services View	12-16
Displaying Thresholds	12-16
Displaying Historical Data	12-18
Displaying Metrics	12-19
Displaying and Creating Charts	12-21
Displaying Virtualization Analytics.....	12-23
Customizing the Analytics View	12-23
Provisioning Operating Systems	12-24
About Dynamic Host Configuration Protocol (DHCP).....	12-25
About WAN Boot.....	12-25
Requirements for a WAN Boot Connection.....	12-26
Checking Your Client OBP for WAN Boot Support	12-27
Setting Up a WAN Boot Connection.....	12-27
Disabling and Enabling WAN Boot	12-27
Adding or Re-Aliasing a Device for WAN Boot	12-28
Importing Images.....	12-28
Overview of OS Provisioning Profiles	12-29
Specifying Common Linux Parameters.....	12-29
About Oracle Solaris or Linux OS Provisioning Profiles	12-30
NVRAMRC File.....	12-31
About Oracle VM Server for SPARC OS Provisioning Profile.....	12-31
About JumpStart Enterprise Toolkit (JET) for Oracle Solaris.....	12-31
About Provisioning Plans	12-32
About Deployment Plans That Provision an Operating System.....	12-33
Provisioning Linux and Oracle Solaris 8, 9, and 10.....	12-33
Provisioning Oracle Solaris 11.....	12-34
Oracle Solaris 11 Provisioning Profiles.....	12-35
Steps for Oracle Solaris 11 Provisioning Plan	12-35
Defining Network Interface Names for Oracle Solaris 11.....	12-35
Selecting an Alternate Network Interface for Provisioning Logical Domains.....	12-36
Overview of Oracle Solaris Boot Environments	12-36
Understanding the Differences Between Oracle Solaris 11 and Oracle Solaris 10 Boot Environments	12-37
Monitoring Boot Environments	12-37
Viewing Boot Environments	12-38
Managing Boot Environments	12-39
Activating and Reboot a Boot Environment	12-40
Deleting a Boot Environment.....	12-40
Overview of Oracle Solaris 11 Boot Environments	12-41
Displaying Oracle Solaris 11 Boot Environment Details.....	12-41
Displaying Total ZPools Utilization for Oracle Solaris 11 Boot Environments	12-42

Displaying Oracle Solaris 11 Boot Environments	12-42
Displaying Snapshots for Oracle Solaris 11 Boot Environments	12-43
Displaying File Systems for Oracle Solaris 11 Boot Environments	12-43
Displaying Associated Zone Boot Environments	12-43
About Boot Environment Profiles and Plans for Oracle Solaris 11.....	12-44
Creating an Oracle Solaris 11 Boot Environment	12-45
Overview of Oracle Solaris 10 Boot Environments	12-45
Requirements for Oracle Solaris 10 Live Upgrade and Oracle Solaris Zones	12-47
Displaying Boot Environment Details for Oracle Solaris 10.....	12-47
Displaying Total ZPools Utilization for Oracle Solaris 10 and Earlier Boot Environments	12-48
Displaying Boot Environments for Oracle Solaris 10 and Earlier.....	12-48
Displaying File Systems for Oracle Solaris 10 and Earlier Boot Environments.....	12-48
About Boot Environment Profiles and Plans for Oracle Solaris 10 and Earlier	12-49
Determining Your Boot Environment Policy.....	12-49
Determining Your Live Upgrade Feature Requirements and Options.....	12-49
Overview of Boot Environments Profiles and Plans	12-50
Creating an Oracle Solaris 10 Boot Environment	12-50
Defining Deployment Options for Oracle Solaris 10	12-51
Creating an Oracle Solaris 10 Boot Environment From a Deployment Plan	12-52
Creating an Oracle Solaris 10 Boot Environment from an Operational Plan.....	12-52
Creating an Update Profile for Oracle Solaris 10 Boot Environment.....	12-53
Creating an Oracle Solaris 10 Boot Environment From an Update Profile	12-53
Creating an Oracle Solaris 10 Boot Environment With an OS Update Job.....	12-54
Synchronizing Oracle Solaris 10 Boot Environments	12-55
Activating a Boot Environment.....	12-56
Related Resources	12-56

13 Operating System Updates

Introduction to Operating System Updates	13-1
Requirements for Updating Operating Systems	13-3
Methods of Running an Update Job	13-4
Options Available When Running an Update Job	13-4
Roles for Operating System Updates	13-4
Actions for Operating System Updates	13-5
Location of Operating System Updates in the User Interface	13-5
Using System Catalogs.....	13-6
Viewing and Modifying a Catalog	13-6
Comparing System Catalogs	13-6
About Operating System Update Reports.....	13-7
Creating Update Policies.....	13-7
Creating Update Profiles.....	13-9
Updating Oracle Solaris 11 Operating Systems	13-11
Updating Oracle Solaris 11	13-11
About Oracle Solaris 11 Update Deployment Plan.....	13-12
Updating Oracle Solaris 8, 9, and 10 and Linux Operating Systems.....	13-13
Updating an Operating System From a Deployment Plan	13-13

About Operating System Update Jobs	13-14
Updating an Operating System by Modifying a System Catalog.....	13-14
Updating an Operating System From an Operating System Report Result	13-15
Using a System Catalog	13-15
Updating an Oracle Solaris Boot Environment	13-16
Updating a Boot Environment Using an Update Job	13-16
Updating Microsoft Windows Operating Systems	13-19
About Windows OS Update Jobs	13-19
Modify the Registry	13-19
Configure Oracle Enterprise Manager Ops Center for Updating the Windows Operating System	13-20
Creating an Update Job for the Windows Operating System	13-20
Related Resources	13-21

Part IV Virtualize and Cloud Management

14 Oracle Solaris Zones

Introduction to Oracle Solaris Zones	14-1
Global and Non-global Zones	14-2
Types of Non-Global Zones.....	14-2
Roles for Zone Management	14-2
Zone Management Functions	14-3
Preparing Your Global Zone	14-3
Associating Storage Libraries	14-4
Managing Global Zone Networks	14-4
Creating IPMP Groups	14-7
Creating Link Aggregation	14-9
Modifying and Detaching a Network from the Global Zone	14-10
Discovering and Managing Existing Zones	14-11
Deleting or Unmanaging a Global Zone.....	14-12
Unmanaged Zone Storage Information	14-12
Outline of Zone Creation	14-13
Determining Zone Requirements	14-13
Requirements for Zones on Oracle Solaris 10 OS.....	14-14
Requirements for Zones on Oracle Solaris 11 OS.....	14-14
Zone Configuration Parameters	14-15
Creating a Zone Profile	14-16
Creating and Deploying Zone Plans	14-19
Zpool and File System of Zones.....	14-21
Modify Zone Configuration	14-23
Managing Zones	14-24
Replicating Zones.....	14-25
Adding Storage to Zones	14-27
Moving Zone Storage	14-27
Adding File Systems to Zones.....	14-29
Connect and Disconnect Networks	14-30
Automatic Recovery	14-30

Migrating Zones	14-30
Migrating a Physical Oracle Solaris System into a Zone.....	14-30
Migrating Zones to a Different Machine	14-31
Script to Migrate a Zone With Dependencies	14-36
Script Requirements.....	14-37
Recovering Zones	14-38
Zones Server Pool.....	14-40
Updating Zones	14-40
Install Packages and Patches on Zones.....	14-40
Configure patchadd and pkgadd Commands	14-41
Editing the .uce.rc File	14-41
Updating a Global Zone.....	14-41
Updating a Non-Global Zones.....	14-42
Zone Parallel Patching.....	14-43
Related Zone Operations	14-43
Related Resources	14-44

15 Oracle VM Server for SPARC

Introduction to Oracle VM Server for SPARC	15-1
Roles for Managing Oracle VM Server for SPARC.....	15-2
Available Management Features.....	15-2
Prerequisites for Provisioning Oracle VM Server for SPARC	15-3
Operating System.....	15-3
Hardware and Firmware	15-3
Recommended Minimum Configuration	15-4
CPU Threads	15-4
Crypto Units.....	15-4
RAM	15-5
Creating an OS Provisioning Profile.....	15-5
Oracle VM Server for SPARC Parameters.....	15-5
To Create an Oracle VM Server for SPARC Provisioning Profile.....	15-6
Create Link Aggregation in Oracle VM Server for SPARC	15-9
Create and Apply Deployment Plan for Oracle VM Server for SPARC	15-10
Installing Oracle VM Server for SPARC	15-11
Selection of Oracle VM Server Version.....	15-12
Important Notes for Installation	15-13
Manual Net Boot Initiation	15-13
Additional Configurations.....	15-13
Oracle VM Server for SPARC Management.....	15-13
Edit Configuration and Tags	15-14
Associate Storage Libraries.....	15-14
Attaching Networks.....	15-14
Reboot Oracle VM Server.....	15-17
Performance Management.....	15-17
Server Pools	15-18
About Logical Domains	15-18
Creating a Logical Domain Profile	15-19

Creating and Deploying a Logical Domain Plan	15-21
Applying Logical Domain Deployment Plan	15-22
OS Provisioning on Logical Domains	15-22
Complex Plan.....	15-23
Management of Logical Domains	15-23
Edit Logical Domain Configuration	15-24
Shut Down a Logical Domain	15-24
Starting a Logical Domain	15-24
Add Storage to Logical Domain.....	15-25
Managing Logical Domain Networks.....	15-25
Migrate Logical Domains.....	15-26
Setting User Accounts for Migration	15-26
Migrating a Logical Domain	15-27
Migrating Multiple Logical Domains	15-27
Automatic Recovery of Logical Domains	15-27
Recovering Logical Domains.....	15-29
Delete a Logical Domain	15-29
Connect to Logical Domain Console.....	15-29
Manage Existing Oracle VM Server for SPARC Environments.....	15-29
Related Resources	15-31

16 Oracle VM Server for x86

Introduction to Oracle VM Server for x86	16-1
Roles for Oracle VM Server for x86 Management	16-2
Integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center	16-2
Installing Oracle VM Manager.....	16-3
Discovering Oracle VM Manager	16-3
Discover Oracle VM Servers.....	16-6
Discovery Profile for Oracle VM Server	16-6
Discovering Oracle VM Servers Using SSH.....	16-6
Installing Oracle VM Servers	16-7
Creating OS Provisioning Profile for Oracle VM Server	16-7
Applying Deployment Plan for Oracle VM Server.....	16-9
Administration of Oracle VM Manager.....	16-10
Manage Ownership of Oracle VM Servers.....	16-11
Discover Storage Resources.....	16-11
Manage Storage Resources	16-12
Manage Networks.....	16-13
Attaching Networks.....	16-14
Manage Oracle VM Servers.....	16-14
Edit Oracle VM Server Information	16-15
Editing IPMI Configuration.....	16-15
Placing Oracle VM Server in Maintenance Mode	16-15
Updating Oracle VM Server	16-16
Create Server Pool	16-17
Server Pool Policies.....	16-17
Creating Oracle VM Server for x86 Server Pools.....	16-18

Create Virtual Machines	16-20
Virtualization Types	16-21
Creating Virtual Machine Profile.....	16-21
Deploying Virtual Machine Plan	16-24
Provisioning OS on Virtual Machines.....	16-25
Manage Virtual Machines	16-25
Edit Boot Order.....	16-26
Automatic Recovery	16-27
Related Resources	16-27

17 Server Pools

Introduction to Server Pools	17-1
Roles for Server Pool Management	17-2
Server Pool Policies	17-2
Placement Policy	17-2
Minimize Power Consumption Policy	17-3
Automatic Load Balancing Policy	17-3
Automatic Recovery	17-4
Server Pool Libraries	17-5
Server Pool Networks	17-5
Oracle VM Server for SPARC Server Pool	17-5
Server Pool Policies.....	17-6
Automatic Recovery	17-7
Creating an Oracle VM Server for SPARC Server Pool.....	17-8
Oracle Solaris Zones Server Pool	17-11
IP Stack in Server Pool.....	17-11
Server Pool Policies.....	17-12
Creating a Zones Server Pool	17-12
Oracle VM Server for x86 Server Pool	17-15
Server Pool Policies.....	17-16
Creating an Oracle VM Server for x86 Server Pools	17-17
Manage Server Pools	17-19
Editing Server Pool Parameters	17-19
Adding Virtualization Hosts	17-20
Associating Network Domains	17-21
Attaching Networks.....	17-23
Associating Libraries	17-24
Creating Guests	17-24
Migrating Multiple Guests.....	17-25
Migrating Zones	17-25
Migrating Logical Domains.....	17-26
Migrating Virtual Machines	17-26
Balancing Resources	17-26
Deleting Server Pool	17-27
Related Resources	17-27

18 Virtual Datacenters

Introduction to Cloud Management	18-2
Oracle Engineered Systems	18-2
Overview of Virtual Datacenter	18-3
Roles for Managing Virtual Datacenter	18-3
Required User Role	18-3
Actions Available in vDC Management	18-3
Virtual Datacenter Creation	18-4
Server Pool Setup	18-4
Storage Setup	18-5
Network Setup	18-6
CPU Oversubscription	18-7
Creating a Virtual Datacenter	18-8
vDC Management	18-10
Update vDC Configuration	18-10
Resource Management	18-10
Delete vDC	18-12
Create and Manage Cloud Users	18-12
Adding and Removing Cloud Users to an Account	18-12
Manage Cloud User Roles	18-13
Account Creation	18-14
Creating an Account	18-14
Account Management	18-15
Update Account Configuration	18-15
Account Resource Management	18-15
Maintaining OS Images	18-15
Delete an Account	18-16
Create and Manage vServer Types	18-16
Creating a vServer Type	18-17
Update vServer Types	18-18
Delete vServer Type	18-18
Overview of Cloud Users	18-18
Roles for Cloud User Tasks	18-19
Actions Available for a Cloud User	18-19
View Account Quotas and Virtual Resources	18-21
Creating and Managing vServers	18-22
Before You Begin	18-23
Creating vServers	18-24
Managing vServers	18-25
Creating vNets	18-26
Creating a Private vNet	18-27
Managing vNets	18-27
Managing Server Templates	18-28
Uploading Server Templates	18-28
Creating Server Templates	18-29
Server Template Management	18-30
Volume Creation	18-30

Creating Volumes.....	18-31
Volume Management	18-32
Creating Snapshots	18-32
Snapshot Management	18-33
Creating Distribution Groups	18-33
Distribution Group Management	18-34
Related Resources	18-34

Part V Command Line Interface

19 Command Line Interface

CLI Features.....	19-1
Asset Discovery, Management, and Grouping.....	19-1
Operating System Provisioning and Patching.....	19-2
Firmware Provisioning.....	19-3
Administration	19-3
Starting the CLI	19-3
Invoking the Command Line Interface.....	19-3
Connecting From the Enterprise Controller.....	19-4
Connecting From a Remote System	19-4
Using the CLI.....	19-4
Defining User Variables	19-5
Using the CLI One-Liner.....	19-5
General Commands	19-5
Checking Connectivity	19-5
Displaying User Configuration.....	19-6
Setting the User Environment	19-6
Recording Console Input and Output.....	19-7
Using a Source File.....	19-7
Viewing the Version	19-7
Waiting for a Job to Complete.....	19-7
Viewing Help.....	19-8
Reviewing Commands	19-8
Displaying a Man Page.....	19-8
Ending a Mode	19-9
Disconnecting from Oracle Enterprise Manager Ops Center	19-9
Exiting the CLI.....	19-9
Universal Output Filters	19-9
Command Scripts	19-10
Mode Commands	19-10
Controllers.....	19-11
Viewing Controllers.....	19-11
Evacuating Assets From a Failed Proxy Controller	19-11
Credentials	19-12
Viewing Credentials	19-12
Modifying Credential Details.....	19-12

Deploy-Setup	19-13
Starting Deploy-Setup Mode.....	19-13
Selecting a Deployment Plan.....	19-13
Selecting a Target Asset	19-13
Deploy.....	19-13
Viewing the Plan	19-14
Setting the Attributes of a Deployment Plan.....	19-15
Adding Attributes to a Deployment Plan	19-17
Testing a Deployment Plan	19-18
Applying a Deployment Plan	19-18
Discover.....	19-18
Starting Discover Mode	19-18
Creating a Discovery Job	19-18
Discovering an OS.....	19-19
Discovering an ILOM Service Processor	19-20
Discovering an ALOM Service Processor.....	19-20
Viewing Discovery Profiles	19-21
Viewing Discovery Ranges.....	19-21
Viewing Driver Credentials.....	19-21
Viewing Discovery Profile Credentials	19-21
Viewing Discovery Ranges.....	19-21
Viewing Networks.....	19-21
Executing Discovery Profiles	19-21
FWImage	19-22
Starting Fwimage Mode.....	19-22
Viewing Firmware Images	19-22
Creating a Firmware Image.....	19-22
Updating a Firmware Image	19-23
Deleting a Firmware Image	19-23
FWProfile.....	19-23
Starting FW Profile Mode	19-23
Viewing Firmware Profiles.....	19-23
Checking the Status of a Firmware Profile.....	19-24
Creating a Firmware Profile	19-24
Updating a Firmware Profile	19-24
Deleting a Firmware Profile	19-24
Gear	19-24
Starting Gear Mode.....	19-25
Viewing Assets.....	19-25
Updating Asset Information	19-25
Viewing Jobs Associated With Assets.....	19-26
Managing an Asset	19-26
Unmanaging an Asset	19-27
Setting the Maintenance Mode of an Asset.....	19-27
Deleting an Asset	19-27
Reinstalling an Agent Controller on an Operating System	19-28
Rebooting an Operating System	19-28

Updating Firmware	19-29
Checking Firmware Compliance	19-29
Provisioning an OS	19-29
Changing Locator Lights	19-30
Enabling or Disabling Notifications	19-30
Powering Off a Server or Chassis	19-30
Powering On a Server or Chassis	19-31
Resetting a Server or Chassis	19-31
Refreshing a Server	19-31
Groups	19-32
Starting Groups Mode.....	19-32
Viewing Groups	19-32
Viewing Members of a Group.....	19-32
Viewing Group Memberships.....	19-32
Creating a Group	19-32
Updating Group Attributes.....	19-33
Adding an Asset To a Group	19-33
Removing an Asset From a Group	19-33
Moving an Asset From One Group To Another Group	19-33
Deleting a Group.....	19-33
Powering On the Assets in a Group.....	19-33
Powering Off the Assets in a Group	19-33
Resetting Servers in a Group.....	19-34
Incidents	19-34
Starting Incidents Mode.....	19-34
Viewing Incidents	19-34
Viewing Incident Details	19-34
Viewing Alerts Associated with Incidents	19-34
Acknowledging an Incident	19-34
Adding an Annotation to One or More Incidents.....	19-35
Assigning Incidents	19-35
Closing Incident	19-35
Jobs	19-35
Starting Jobs Mode.....	19-35
Viewing Jobs	19-35
Starting a Job.....	19-36
Repeating a Job on Failed Targets	19-36
Deleting a Job.....	19-36
Stopping a Job.....	19-36
Notifications.....	19-36
Starting Notifications Mode	19-36
Viewing Notifications.....	19-37
Deleting a Notification	19-37
OSImage	19-37
Starting OS Image Mode.....	19-37
Viewing OS Images	19-37
Importing an OS Image.....	19-38

Deleting an OS Image.....	19-38
OSProfile.....	19-38
Starting OS Profile Mode	19-38
Viewing OS Profiles.....	19-38
Importing a JET Template	19-38
Importing an OS Profile	19-38
Exporting an OS Profile	19-38
Copying an OS Profile.....	19-39
Deleting an OS Profile	19-39
ServerPools.....	19-39
Starting Serverpools Mode	19-39
Viewing Server Pools	19-39
Viewing Members of a Server Pool.....	19-39
Removing a Member From a Server Pool	19-39
Stats	19-39
Starting Stats Mode.....	19-39
Viewing Statistics	19-40
Update	19-40
Starting Update Mode.....	19-40
Viewing OS Update Policies.....	19-40
Viewing OS Update Profiles.....	19-40
Viewing OS Distributions	19-40
Viewing Snapshots	19-41
Searching Inventory.....	19-41
Checking Inventory	19-41
Checking Connectivity to the Knowledge Base	19-41
Querying Job History	19-42
Querying Job Status	19-42
Uploading a Knowledge Base Bundle	19-43
Checking the Connection Mode	19-43
Changing to Connected Mode	19-43
Changing to Disconnected Mode	19-43
Modifying an OS Asset	19-43
Applying an Update Profile to an OS Asset	19-44
Creating a Historical Snapshot	19-44
Deleting a Snapshot	19-44
Creating a Profile	19-45
Creating a Profile From Inventory	19-45
Deleting a Profile.....	19-45
Adding a File to a Distribution	19-46
Uploading Local Software in Bulk From a Directory	19-46
Adding a Local Action	19-46
Adding a Local Category.....	19-46
Deleting a Component	19-46
Adding a Configuration File	19-46
Setting a Component File.....	19-46
Finding Nodes	19-47

User	19-47
Starting User Mode.....	19-47
Viewing Users	19-47
Viewing Role Types.....	19-48
Viewing User Roles.....	19-48
Adding a User to the Enterprise Controller.....	19-48
Granting Roles to a User	19-48
Replicating User Roles	19-48
Removing Roles From a User.....	19-48
Adding a Directory Server.....	19-48
Viewing Directory Servers.....	19-49
Viewing Remote Users	19-49
Synchronizing All Remote Users.....	19-49
Synchronizing a Single User.....	19-49
Removing a Directory Server	19-49
Deleting a User	19-49
Virtualization	19-50
Starting Virtualization Mode	19-50
Viewing Hosts	19-50
Viewing Guests	19-50
Migrating a Non-Global Zone.....	19-50
Finding the Equivalent Browser User Interface Actions For the Mode Commands	19-50
Controllers Mode	19-51
Credentials Mode.....	19-51
Deploy-Setup Mode.....	19-51
Deploy Mode	19-52
Discover Mode	19-52
FWImage Mode.....	19-52
FWProfile Mode	19-53
Gear Mode.....	19-53
Groups Mode.....	19-54
Incidents Mode.....	19-55
Jobs Mode.....	19-55
Notifications Mode	19-55
OSImage Mode.....	19-55
OSProfile Mode	19-56
ServerPools Mode	19-56
Update Mode	19-56
User Mode.....	19-58
Virtualization Mode	19-58

Part VI Engineered Systems

20 Oracle Engineered Systems

Introduction to Oracle Engineered Systems	20-1
Datacenter Management	20-2

Finding Assets	20-2
Creating a Discovery Profile	20-3
Adding Assets Using a Discovery Profile	20-4
Viewing the Asset Content Pane	20-5
Relayed Incidents.....	20-6
Relayed Service Requests.....	20-6
Baseline Check	20-6
Oracle Exalogic Elastic Cloud	20-6
Oracle Enterprise Manager Ops Center Configuration.....	20-7
Understanding User Roles.....	20-7
Exalogic System Admin Role	20-8
Cloud Administrator	20-8
Cloud User	20-8
Adding a User.....	20-8
Viewing the Exalogic System	20-8
Dashboard Tab	20-10
Details Tab	20-11
Controls Tab.....	20-12
Networks Tab	20-12
Incidents Tab.....	20-13
Permitted Actions in the Exalogic System View	20-14
Viewing the Exalogic System Rack.....	20-14
Photorealistic View of the Rack	20-15
Permitted Actions in Rack View.....	20-16
Creating and Viewing Reports.....	20-16
SPARC SuperCluster T4-4.....	20-19
SPARC SuperCluster Configuration	20-19
Understanding User Roles.....	20-20
Viewing the SPARC SuperCluster System	20-20
Dashboard Tab	20-21
Details Tab	20-23
Networks Tab	20-24
Incidents Tab.....	20-24
Permitted Actions in the SPARC SuperCluster System View	20-25
Viewing the SPARC SuperCluster System Rack	20-25
Photorealistic View of the SPARC SuperCluster System Rack	20-26
Viewing the Exadata Storage Server	20-27
Discovering a SPARC SuperCluster Component.....	20-28
Creating and Viewing Reports.....	20-29
Related Resources	20-30

Part VII Reference

A Asset Attributes

Monitoring Attributes	A-1
Grouping Attributes	A-1
Time Zone Values	A-10

B Oracle Solaris Cluster

Upgrading an Oracle Solaris Cluster.....	B-1
Cluster Profiles	B-1
Obtaining the Cluster Profiles and Scripts	B-2
Uploading the Cluster Profiles and Scripts	B-3
Importing Cluster Profiles and Scripts	B-4
Editing the Core Profile for Provisioning	B-4

C Logs and Directories

Installation.....	C-1
General Events	C-1
High Availability.....	C-2
Software Updates	C-2
Agents	C-2
Local Database	C-2
Upgrade	C-3
Event Logs for Jobs	C-3
Controlling the Number of Common Agent Container Log Files	C-3

D JumpStart Enterprise Toolkit

JumpStart Enterprise Toolkit Configuration File Location	D-1
SUNWjet Parameters	D-1
Downloading Additional JET Packages	D-7

E Icons

Administration Icons	E-1
Badges.....	E-3
Boot Environment Icons.....	E-3
Connectivity Icons.....	E-4
Group Icons	E-5
Hardware Icons	E-6
Incident Icons	E-8
Job Icons	E-9
Kernel Icons.....	E-10
Library Icons.....	E-10
Message Center Icons	E-11
Monitoring Icons	E-11
Network Icons	E-13
Operating System Icons	E-13
Oracle VM Server Icons	E-14
Process Icons.....	E-15
Report Icons.....	E-15
SuperCluster Icons.....	E-17
Switch Icons	E-18
Universal Icons	E-19

Update Policy Icons	E-20
Update Profile Icons	E-20
vDC Icons.....	E-21

F Oracle Enterprise Manager Ops Center API

List of Public APIs.....	F-4
Public API Access.....	F-6
JMX.....	F-6
WS-Management.....	F-6
Public API Access Through JMX Clients.....	F-7
Public API Access Through WS-Management	F-7
Virtualization Management Stack	F-7
Presentation Tier	F-7
Service Tier.....	F-7
Domain Model.....	F-7

G Expression Query Language

Monitoring Attributes and Classes.....	G-1
Classes.....	G-1
AttributeType	G-2
Lexical Elements	G-2
Grammar	G-3
Semantics	G-4
Informal Examples	G-5

Glossary

Index

Preface

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

Audience

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

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information, see the [Oracle Enterprise Manager Ops Center documentation library](#).

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

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

Conventions

The following text conventions are used in this document:

Convention	Meaning
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, code in examples, text that appears on the screen, or text that you enter.

Part I

Getting Started

Part I contains the following chapters:

- [Chapter 1, "Introduction"](#)
- [Chapter 2, "Asset Management"](#)
- [Chapter 3, "Jobs"](#)

Introduction

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

Oracle Enterprise Manager Ops Center in the Data Center

After the assets have been discovered and brought under the management of the software as described in the *Oracle Enterprise Manager Ops Center Administration Guide*, you can do the following:

- Use a single console to monitor, configure, provision, and update the assets throughout their lifecycle, regardless of where an asset is located.
- See the status of all assets and drill down to see details of a specific asset.
- Track an individual asset's status and view the history of all the assets of that type.
- Track the progress of any incident investigation or service request.
- Deploy new assets in a manner consistent with the existing assets because you use profiles to set the attributes and plans to deploy the profiles to one target or many targets.

Roles and Connection Modes

The role of user account and the product's connection mode affects your ability to perform operations:

- All the actions that are included in your role are available in the Actions pane. When you must accomplish a task and a necessary action is not available, your administrator can add the role to your user account. When you log in again, the action is available.

The required roles for using a feature are listed in each chapter. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

- In Connected mode, deployment plans and profiles that rely on images and packages use the current images and packages downloaded from Oracle and vendor sites to the Enterprise Controller. If your site uses the product software in Disconnected mode, the images and packages in your local knowledge base do not change until your site acquires them.

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

User Preferences and Roles

You can change the way the user interface works:

- [Start In and Start Page Preferences](#)
- [Time Preferences](#)
- [Membership Graph Preferences](#)
- [User Preferences Summary](#)
- [Preferences By Role](#)

More information about the user interface is in the *Oracle Enterprise Manager Concepts Guide*.

Start In and Start Page Preferences

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

Time Preferences

You can set various time intervals to control when the software takes an action or performs an operation.

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

Membership Graph Preferences

You can change the default orientation of the Membership Graph to left, right, top, or bottom. You can choose the icon size as small or large and the level of depth for the assets to be displayed in Membership Graph.

User Preferences Summary

The User Preferences Summary displays a summarized history of the preferences that you have selected for your user preferences. You can view your set preferences for the start page, time intervals, asset default tab, and the Membership Graph.

Preferences By Role

As an administrator, you can set preferences for each role. You can set the display startup preferences, time interval settings, and an asset to be displayed. You can also set the preferences of the membership graph through selecting the depth, icon size, and orientation for the selected role.

How To Use This Collection

This document is divided into parts, not all of which are relevant to your site's activities or your role in the data center. Within each part, this document describes the capabilities of the product software, organized by feature. Where it is practical, all information about a feature is discussed in the same chapter. To avoid repetition, links have been added to other chapters in the document and to the How To library.

All the documentation for the Oracle Enterprise Manager Ops Center 12c Release 1 software is located at this site: <http://www.oracle.com/pls/topic/lookup?ctx=oc121> including end-to-end examples located on the How to tab.

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

Asset Management

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

The following features and topics are covered in this chapter:

- [Introduction to Asset Management](#)
- [Roles for Asset Management](#)
- [Actions for Asset Management](#)
- [Location of Asset Management Information in the User Interface](#)
- [Discovering and Managing Assets](#)
- [Using Tags](#)
- [Using Groups](#)
- [Related Resources for Asset Management](#)

Introduction to Asset Management

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

Roles for Asset Management

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

Table 2–1 Asset Management Tasks and Roles

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

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

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

Actions for Asset Management

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

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

Location of Asset Management Information in the User Interface

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

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

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

Discovering and Managing Assets

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

After assets have been discovered, they are automatically managed, giving Oracle Enterprise Manager Ops Center full access to the assets and enabling you to monitor, update, and provision them. For an asset to be managed, it must be accessible on a network that is associated with a Proxy Controller.

Hardware assets are managed using a set of credentials. Operating systems and virtualization software can be managed using an Agent Controller installed on the system or using only a set of credentials.

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

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

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

Declaring Servers for OS Provisioning

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

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

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

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

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

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

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

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

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

Example Discovery File

Here is an example of a discovery file.

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

The following variables can be used:

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

Declaring Servers for Service Processor Configuration

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

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

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

1. Select **All Assets** in the Navigation pane, then click **Add Assets** in the Actions pane.
2. Select **Declare an Unconfigured Hardware Asset**, then enter data for the servers to be declared:
 - **Number of Servers** – The total number of servers to be discovered.
 - **Model Categories** – The model category of the servers.
 - **Model** – The specific model of the servers.
 - **Server Names** – The names of the servers, including:
 - **Prefix** – A prefix that appears before each server name. This field is required.
 - **Starting Number** – The number of the first server. The number is increased by one for each additional server. This field is required.
 - **Suffix** – A suffix that appears after each server name.
 - **Network** – The network on which the server or servers is added.
 - **IP addresses** – The IP addresses to be used for the servers.
 - **MAC Addresses** – The MAC Addresses of the servers.

Figure 2–1, "Declare Unconfigured Hardware Assets" shows the Declare Unconfigured Hardware Assets Wizard.

Figure 2–1 Declare Unconfigured Hardware Assets

3. Click **Declare Asset**.

Adding Assets Using a Discovery Profile

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

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

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

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

Finding Assets

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

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

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

To find assets, perform the following steps:

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

Creating a Discovery Profile

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

To create a discovery profile, perform the following steps:

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

[Figure 2–2, "Create Discovery Profile"](#) shows the Identify Profile step of the Create Profile Wizard.

Figure 2–2 Create Discovery Profile

The screenshot shows the 'Identify Profile' step of the Create Profile Wizard. The window has a title bar 'Identify Profile' and a legend '* Indicates Required Field'. The form contains the following fields:

- Name:** A text input field with a red asterisk indicating it is required.
- Description:** A text input field.
- Asset Type:** A list box containing the following items:
 - Operating Systems
 - Server Hardware
 - Oracle Engineered Systems
 - Oracle VM
 - Storage
 - Networking
 - Datacenter Infrastructure
 - Cluster Products

4. Enter a name for the discovery profile and select a type of asset to discover. You can also add a description. Click **Next**.

5. (Optional.) Click the add icon to add one or more tags to discovered assets, then click **Next**.
6. (Optional.) Click the add icon to add one or more IP ranges. If you do not add IP ranges when creating a discovery profile, you can do so when you run the profile. Enter the following information for each IP range:
 - Name – A name for the IP range.
 - Description – A description of the IP range.
 - Network – Identify the managed network associated with the host names or IP addresses to route the discovery to the correct Proxy Controller. Select Automatic to route the job to the most appropriate Proxy Controller. The IP address of a target must resolve to only one known network for automatic routing to succeed.
 - IP ranges – An IP range or comma-separated list of IP addresses. If you want to target specific host names, you must enter them when you run the discovery profile.
7. Provide the discovery credentials information, then click **Next**.
 - Click **New** to add new discovery credentials for each protocol, or click **Select** to select an existing set of credentials.
 - Select **Deploy Agent Controller to deploy an Agent Controller on discovered operating systems**, or select **Manage Without Agent Controller to manage discovered operating systems without installing an Agent Controller**.
 - If the service tag parameters have been modified for the target assets, enter the Service Tag parameters:
 - Service tag passphrase – Necessary if a service tag has been configured to be encrypted.
 - Service tag port – Necessary if a service tag has been configured to use a port other than the default of 6481.
 - Service tag timeout – The default value is 20 seconds.
8. Review the summary and click **Finish** to create the discovery profile.

Editing a Discovery Profile

To edit a discovery profile, perform the following steps:

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

Copying a Discovery Profile

To copy a discovery profile, perform the following steps:

1. In the Navigation pane, click **Plan Management**.
2. Under Profiles and Policies, click **Discovery**.
3. Select the profile and click the Copy Profile icon.

4. Edit any of the information in the discovery profile, then click **Finish** to save the new discovery profile.

Deleting a Discovery Profile

To delete a discovery profile, perform the following steps:

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

Installing Agent Controllers From the Command Line

Use these procedures to install an Agent Controller and to register the target system.

Before You Begin

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

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

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

When you install an Agent Controller on a global zone, the Agent Controller installation installs, or upgrades to Java Runtime Environment (JRE) 1.6.0_31. Later versions of JRE are not affected.

To Manually Install and Configure an Agent Controller Using User Credentials

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

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

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

2. Identify the Agent Controller archive that is appropriate for the system where you intend to install the Agent Controller.
3. On the system where you want to install the Agent Controller (the target system), create a directory named `/var/tmp/OC`.

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

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

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

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

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

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

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

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

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


#

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

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

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

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

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

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

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

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

```
# /opt/SUNWxvmoc/bin/agentadm configure -u droot -p /var/tmp/OC/mypasswd -x
172.20.26.218
agentadm: Version 1.0.3 launched with args: configure -u droot -p
/var/tmp/OC/mypasswd -x 10.0.0.0
workaround configuration done.
Certificate:
Serial Number: 947973225
Version: 3
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_Agent Controller
Not valid before: Thu Jun 19 15:36:59 MDT 1969
Not valid after: Thu Apr 19 15:36:59 MDT 2029
Certificate:
Serial Number: 1176469424
Version: 3
Issuer: CN=flyfishing_scn-proxy_ca
Subject: CN=flyfishing_scn-proxy_ca
Not valid before: Thu Jun 19 15:36:56 MDT 1969
```

```
Not valid after: Thu Apr 19 15:36:56 MDT 2029
Accept server's certificate? (y|n)
Y
Connection registered successfully.
scn-Agent Controller configuration done.
Checking if UCE Agent Controller process is still running, it may take a
couple of minutes ...
Process is no longer running
UCE Agent Controller is stopped.
UCE Agent Controller is in [online] state.
Checking if UCE Agent Controller process is up and running ...
The process is up and running.
UCE Agent Controller is started.
Added the zone configuration automation successfully.
Added the service tags recreate script successfully.
#
```

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

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

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

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

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

```
# /opt/SUNWxvmoc/bin/agentadm unconfigure
agentadm: Version 1.0.3 launched with args: unconfigure
verified sc_console command is OK
End of validation
{output omitted}
End of configuration.
```

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

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

```
# sc-console list-connections
```

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

To Manually Install and Configure an Agent Controller Using a Token

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

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

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

2. Identify the Agent Controller archive that is appropriate for the system where you intend to install the Agent Controller.
3. On the system where you want to install the Agent Controller (the target system), create a directory named `/var/tmp/OC`.

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

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

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

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

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

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

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

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

```
# OpsCenterAgent/install -a
Installing Ops Center Agent Controller.
No need to install 120900-04.
No need to install 121133-02.
No need to install 119254-63.
No need to install 119042-09.
No need to install 121901-02.
No need to install 137321-01.
Installed SUNWjdmk-runtime.
```

```
Installed SUNWjdmk-runtime-jmx.
(output omitted)
6 patches skipped.
19 packages installed.
Installation complete.
Detailed installation log is at /var/scn/install/log.
Uninstall using /var/scn/install/uninstall.
#
```

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

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

(output omitted)

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

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

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

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

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

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

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

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

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

workaround configuration done.

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

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

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

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

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

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

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

{output omitted}
End of configuration.
```

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

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

```
# sc-console list-connections
scn-Agent Controller https://10.0.0.0:21165
```

urn:scn:clregid:abcdef12-6899-4bcc-9ac7-a6ebaf71c1f5:20090420171121805
#

Using Management Credentials

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

Upgrading Management Credentials From a Prior Version

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

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

To upgrade management credentials, perform the following steps:

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

Updating Management Credentials

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

To update management credentials, perform the following steps:

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

Creating Management Credentials

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

To create management credentials, perform the following steps:

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

Using Custom SSH Keys for OS Discovery

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

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

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

Figure 2–3 Create Credentials

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

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

Editing Management Credentials

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

To edit management credentials, perform the following steps:

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

Copying Management Credentials

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

To copy management credentials, perform the following steps:

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

Deleting Management Credentials

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

To delete management credentials, perform the following steps:

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

Editing Asset Attributes

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

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

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

To edit asset attributes, perform the following steps:

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

Using Access Points

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

The following are possible access points:

- The discovery credentials used to discover the asset.

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

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

To delete an access point, perform the following steps:

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

Deleting Assets

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

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

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

To Delete Assets

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

Special Discovery and Management Procedures

Most assets can be discovered and managed using the procedures in [Discovering and Managing Assets](#). However, some types of assets must be discovered or managed

using special procedures. Use these procedures when any of the following asset types is discovered or managed.

- Oracle Solaris 11 – You can discover and manage Oracle Solaris 11 operating systems only if the Enterprise Controller and Proxy Controller are installed on Oracle Solaris 11 and the Oracle Solaris 11 Update Library is configured.
- Microsoft Windows – You must enable Windows Management Instrumentation (WMI) on Microsoft Windows systems before discovering them. Once WMI is enabled, they can be discovered normally.
- Sun SPARC Enterprise M-Series Servers – You must ensure that user privileges and the status of each dynamic system domain are correct before discovering a Sun SPARC Enterprise M-Series server.
- Sun ZFS Storage Appliances – You must discover both the storage appliance and its service processor, and follow special procedures to manage them.
- Oracle Solaris Clusters – You must discover and manage Oracle Solaris Clusters in a specific order so that Oracle Enterprise Manager Ops Center can manage the entire cluster.

Windows Systems

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

To Enable WMI

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

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

To Allow WMI Through the Windows Firewall

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

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

```
netsh advfirewall firewall set rule group="remote administration" new enable=yes
```
 - On other Windows systems, use the following command:

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

Servers

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

Discovering a SPARC Enterprise M-Series Server

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

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

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

Perform the following tasks to discover this type of asset:

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

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

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

Sun ZFS Storage Appliance

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

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

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

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

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

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

The following is the general procedure for discovering a Sun ZFS Storage Appliance. Each step is described in the following sections.

- From the Sun ZFS Storage Appliance's User Interface:
 1. Set up the Sun ZFS Storage Appliance to use the Fibre Channel protocol and, to use dynamic storage libraries, the iSCSI protocol.
 2. Create an account for Oracle Enterprise Manager Ops Center
 3. If the Sun ZFS Storage Appliance uses a cluster configuration, verify that the administrative interfaces for the nodes are private.
- From Oracle Enterprise Manager Ops Center's User Interface:
 1. Create a discovery profile for the Sun ZFS Storage Appliance
 2. Discover the storage appliance, which creates a dynamic storage library.

Setting Up the Sun ZFS Storage Appliance's LUNs

These procedures specify the type of LUN, the access to the LUN, and the path to the LUN, using the Fibre Channel and iSCSI protocol. Fibre Channel has a physical connection between the storage device and the server that uses the storage. The iSCSI standard is a network connection. These procedures are performed on the Sun ZFS Storage Appliance and use its UI. These instructions configure the Sun ZFS Storage Appliance to expose LUNs as Fibre Channel and iSCSI LUNs. You can configure one type of LUN.

See the *Sun ZFS Storage Appliance Administration Guide* or the Sun ZFS Storage Appliance's on-line help for current and specific configuration procedures. The

following procedures are examples of the varied ways to configure the Sun ZFS Storage Appliance.

Configure Fibre Channel Ports

You must identify the ports through which the LUNs can be reached. These are called target ports. By default, all the ports are initiator ports. Change at least one port from an initiator port to a target port. Then create a target group and include the target ports in the target group.

Another way to protect access to LUNs is to specify which initiators can write to the LUN. Although the new storage library will be using the LUNs, the library has not yet been created and the Sun ZFS Storage Appliance has not yet been discovered by Oracle Enterprise Manager Ops Center so these entities cannot be included in the Initiators list.

1. Log into the Sun ZFS Storage Appliance.
2. Select Configuration and then select SAN.
3. Click the Fibre Channel Ports tab, if it is not selected.
4. Select at least one port and change its role from Initiator to Target.
5. Reboot the storage appliance.
6. (Optional.) Each port is identified by its WWN. To create an alternate name, select Configuration, then SAN, and then click the (i) icon to see Port Details. In the Alias field, enter a name for the port, for example a combination of host name and port number.
7. To create a new target group, select the port, drag it to the Target Group table, and drop it at the (+) icon. Click the Edit icon to name the new group. Drag more target ports to the group. Make a note of all the target groups you create because you will provide the names in the discovery procedure.

Configure iSCSI Targets

1. Log into the Sun ZFS Storage Appliance.
2. Select Configuration and then Services to verify that the iSCSI service is running and that the iSNS and Radius services are disabled.
3. Select Configuration and then select SAN.
4. Click the Targets tab and then iSCSI Targets.
5. Create a target with the following attributes:
 - Name: OC-iSCSI-Target
 - IQN: Auto-assign
 - Authentication: None
6. To create a new target group, select the new target, drag it to the Target Group table, and drop it at the (+) icon.
7. Click the Edit icon to rename the new group to OC-iSCSI

You will provide the name of the target group and target in the discovery procedure.

Create an Account for Oracle Enterprise Manager Ops Center

To create and maintain a SAN storage library, Oracle Enterprise Manager Ops Center must be able to log into the Sun ZFS Storage Appliance. Create a user account on the

appliance with the username `ocuser`, with the basic administration role, and with the following permissions and characteristics:

- (Required.) Permission to administer, configure and restart the SNMP service.
- Permission to configure alert filters and thresholds so that an Oracle Enterprise Manager Ops Center user with the appropriate role can set up the monitoring feature.
- Permission to administer the Phone Home Service so that an Oracle Enterprise Manager Ops Center user with the appropriate role can use the Enable Appliance Phone Home or Disable Appliance Phone Home actions.
- Permission to power off the appliance so that an Oracle Enterprise Manager Ops Center user with the appropriate role can power off the storage appliance remotely.

In addition, the `ocuser` account has these requirements:

- Must not be a kiosk type.
- Must not require session annotation.

[Figure 2–4](#) shows how the permissions are set. For some types of Sun ZFS Storage Appliances, the window uses multiple pages to display all options. See the Sun ZFS Storage Appliance documentation for information about creating user accounts. To view the Sun ZFS Storage 7000 System Administration Guide, log in to the Unified Storage System software interface and click Help in the top right corner of any screen.

Figure 2–4 User Account for Sun ZFS Storage Appliance

Edit Local User "ocuser" [CANCEL] [APPLY]

Properties

Username: ocuser

Full Name: ocuser

Password:

Confirm:

Require session annotation: ☐

Kiosk user: ☐

Kiosk screen: status/dashboard

Roles Exceptions [ADD]

Scope: Active Directory

Domain or workgroup name: *

☐ 2 Total

☐ domain Join an Active Directory domain

☐ workgroup Join a workgroup

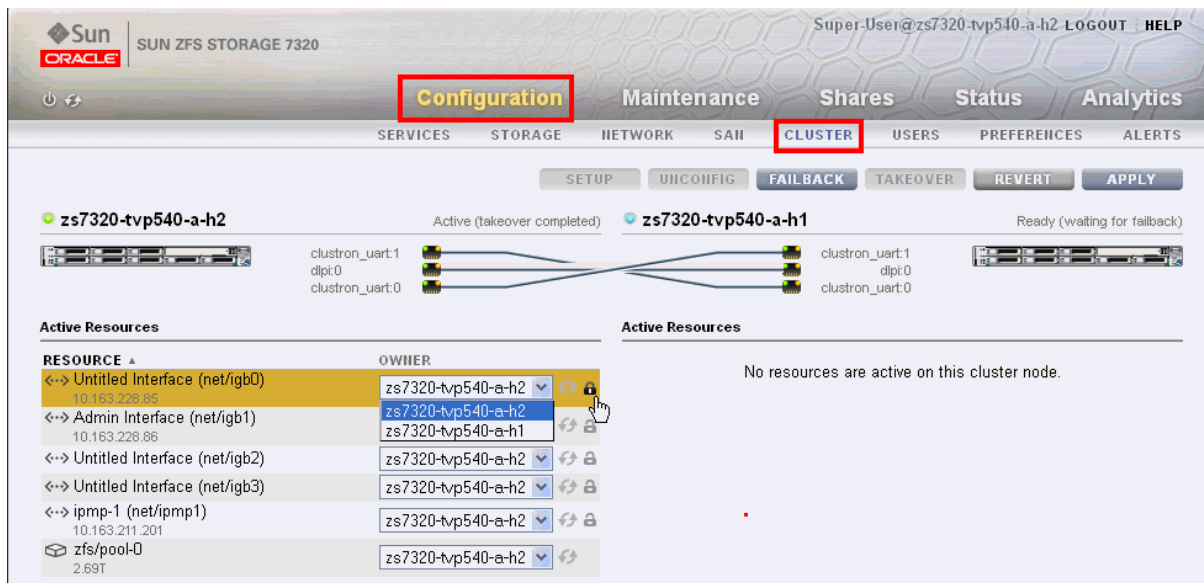
OBJECT	PERMISSIONS
alert	configure
appliance.*	powerOff
nas.*.*	changeGeneralProps, changeProtocolProps, changeSpaceProps, clone, createShare, destroy, rename, takeSnap
net	configure
stmf	configure
svc.scrk	administer, configure, restart
svc.snmp	administer, configure, restart
update	upload, update, delete
workflow.root.*	read

Verify the Administrative Interfaces

Use this procedure if the Sun ZFS Storage Appliance is a model with cluster configuration.

1. Log in to the Sun ZFS Storage Appliance.
2. Select Configuration and then Clusters.

Figure 2–5 Configuration Window for Sun ZFS Storage Appliance



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

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

Create a Discovery Profile

1. Expand Plan Management in the Navigation pane.
2. In the Profiles and Policies section, click **Discovery**.
3. Click **Create Profile** in the Actions pane.
4. Enter a name for the discovery profile.
5. For Asset Type, expand the Storage section and select Sun ZFS Storage Appliance. Click **Next**.
6. (Optional.) The Tags step of the wizard provides a way to add information for semantic tags. All assets discovered by this profile will have the same tags. Click **Next**.
7. In the Targets step, both the NFS and iSCSI targets are enabled. In the Plugin Specific Information field, enter the names of the target group and target that you configured on the Sun ZFS Storage Appliance.

Figure 2–6 Targets for Sun ZFS Storage Appliance

Targets * Indicates Required Field

Targets must be associated with existing managed networks in Ops Center in order to route the discovery to the appropriate Proxy Controller. Only networks associated with Proxies are shown. You may create a new managed network or modify existing network/proxy associations. If specified, the hostname of the storage asset must be resolvable on the network selected. If automatic is selected the system will attempt to route the job to the most appropriate Proxy Controller.

☒ **iSCSI**

* **Access Host:** 192.168.10.10

* **Access Port:** 3260

Plugin Specific Information: OC-iSCSI, OC-iSCSI-Target

Comma separated names of OVM target group and target previously configured on the storage appliance. Example: OVM-iSCSI,OVM-iSCSI-Target

☒ **NFS**

* **Access Host:** 192.168.10.10

* **Access Port:** 2049

8. Specify the credentials needed to get access to the storage appliance's service processor. Set the ssh credentials and, for iSCSI targets, the CHAP credentials. To create new credentials, you must have the Security Admin role.
9. Review the summary and click Finish to create the discovery profile.

Discovering a Sun ZFS Storage Appliance

Use the discovery profile to find the Sun ZFS Storage Appliance on the network, retrieve its configuration information, and make it a managed asset.

1. Expand **All Assets** in the navigation pane.
2. Click **Add Assets** in the Actions pane.
3. Select the option **Add and Manage Various Types of Assets via Discovery Probes**.
4. Select the new discovery profile.

Figure 2–7 Discovery Profile for Sun ZFS Storage Appliance

Discovery Profile - MySS7K

Profile Name	Target	Last Run
MySS7K	Sun ZFS Storage Appliance	10/15/2012 10:01:56 am M
Sample Solaris Discover...	Solaris, Linux OS	10/04/2012 11:56:15 pm M
discovery test	Solaris, Linux OS	10/05/2012 4:37:11 pm M

Specify the hostname of the storage asset that will be added. Targets must be associated with existing managed networks in Ops Center in order to route the discovery to the appropriate Proxy Controller. Only networks associated with Proxies are shown. You may create a new managed network or modify existing network/proxy associations. If specified, the hostname of the storage asset must be resolvable on the network selected. If automatic is selected the system will attempt to route the job to the most appropriate Proxy Controller.

Hostnames/IP Addresses:

Network:

Credentials

The following are credentials populated from the selected profile. You may change these values for the Add Assets job about to be run; the original profile will remain unchanged.

Discovery

Storage Administration : zfs_creds_stadm

CHAP : zfs_creds_chap

5. To direct the discovery probe, enter the host name of the storage appliance or its IP address.
6. For Network, select Automatic. This directs the discovery probe to locate the Proxy Controller closest to the storage appliance.
7. Because you added the credentials when you created the discovery profile, this procedure uses the same credentials. When you use this discovery profile in the future, you have the option of changing the credentials.
8. Click **Add Now** to submit the discovery job.

Oracle Enterprise Manager Ops Center displays the Sun ZFS Storage Appliance in the Storage section of the Assets tree. The LUNs defined in the storage appliance are now available for a new storage library. Because you discovered the storage appliance's iSCSI targets, a dynamic block storage library is created.

Oracle Solaris Cluster

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

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

To Discover an Oracle Solaris Cluster

1. Verify that all global nodes in the cluster are in cluster mode.
2. Discover all of the cluster's global nodes using a discovery profile. Create both ssh credentials and JMX credentials for the agent on each global node.
3. Discover the cluster using a discovery profile.

Using Tags

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


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

Adding Tags

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

1. On the Navigation pane, select an asset or group.
2. In the Actions pane, click **Edit Tags**.

Figure 2–8 Add Tags



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

Viewing Tags

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

Grouping Assets Using Tags

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

Deleting Tags

To delete tags, perform the following steps:

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

Using Groups

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

Types of Groups

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

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

System Groups

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

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

Oracle Enterprise Manager Ops Center creates these system groups:

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

User-Defined Groups

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

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

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

Viewing Group Data

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

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

- A group summary that shows the group's name, description, tags, location, and number of members.
- A membership graph showing the group's assets, any child groups, and any parent groups.
- A status summary showing the problems of the assets within the group.
- An asset summary showing basic data about the assets within the group.

Creating a Group

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

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

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

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

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

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

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

See [Appendix A](#) for a list of the asset attributes that can be used in group rules.

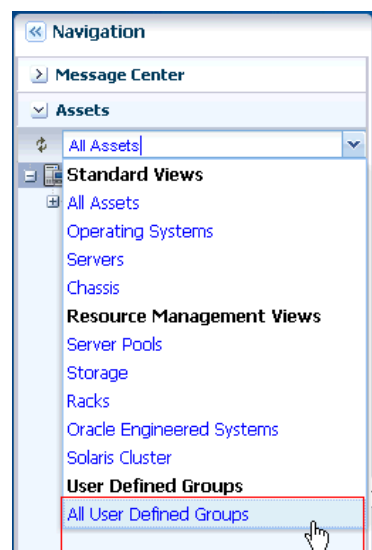
Editing a Group

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

To edit a group, perform the following steps:

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

Figure 2–9 All User Defined Groups



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

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

See [Appendix A](#) for a list of the asset attributes that can be used in group rules.

Adding Assets to a Group

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

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

Removing Assets from a Group

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

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

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

Moving Assets to a Group

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

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

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

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

Moving a Group

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

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

Deleting a Group

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

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

Related Resources for Asset Management

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

- [Chapter 4, "Monitoring Rules and Policies"](#)
- *Oracle Enterprise Manager Ops Center Administration Guide*

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

The following information is included:

- [Introduction to Jobs](#)
- [Roles for Job Management](#)
- [Job Management](#)
- [Viewing Jobs](#)
- [Taking Actions on a Job](#)
- [Changing the Job Properties](#)
- [Events for Job Management](#)
- [Related Resources for Job Management](#)

Introduction to Jobs

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

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

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

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

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

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

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

You can perform the following actions on jobs:

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

Roles for Job Management

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

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

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

Table 3–1 Job Management Tasks and Roles

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

Job Management

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

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

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

Viewing Jobs

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

Viewing All Jobs and Jobs With Specific Status

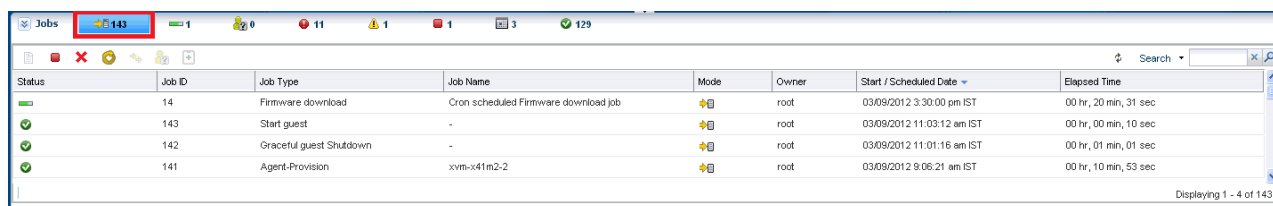
You can view a list of all jobs run by Oracle Enterprise Manager Ops Center.

To view all jobs, perform the following steps:

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

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

Figure 3–1 Jobs pane



Status	Job ID	Job Type	Job Name	Mode	Owner	Start / Scheduled Date	Elapsed Time
✓	14	Firmware download	Cron scheduled Firmware download job	✚	root	03/09/2012 3:30:00 pm IST	00 hr, 20 min, 31 sec
✓	143	Start guest	-	✚	root	03/09/2012 11:03:12 am IST	00 hr, 00 min, 10 sec
✓	142	Graceful guest Shutdown	-	✚	root	03/09/2012 11:01:16 am IST	00 hr, 01 min, 01 sec
✓	141	Agent-Provision	xvm-x41m2-2	✚	root	03/09/2012 9:06:21 am IST	00 hr, 10 min, 53 sec

Displaying 1 - 4 of 143

2. In the Jobs pane, click the All Jobs icon, highlighted in [Figure 3–1](#).

All the jobs are displayed.

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

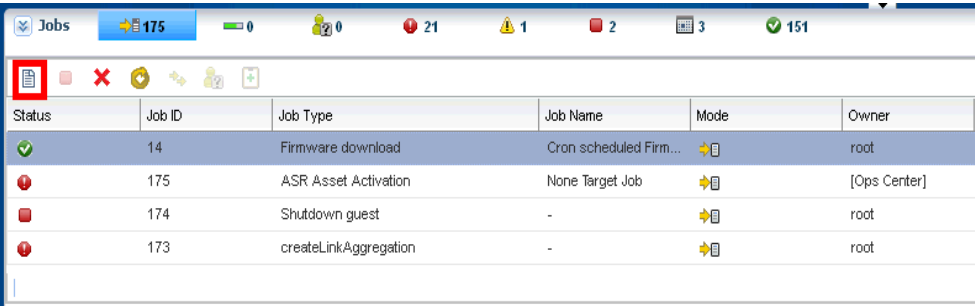
Viewing Job Details

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

To view job details, perform the following steps:

1. Select a job in the Jobs pane.

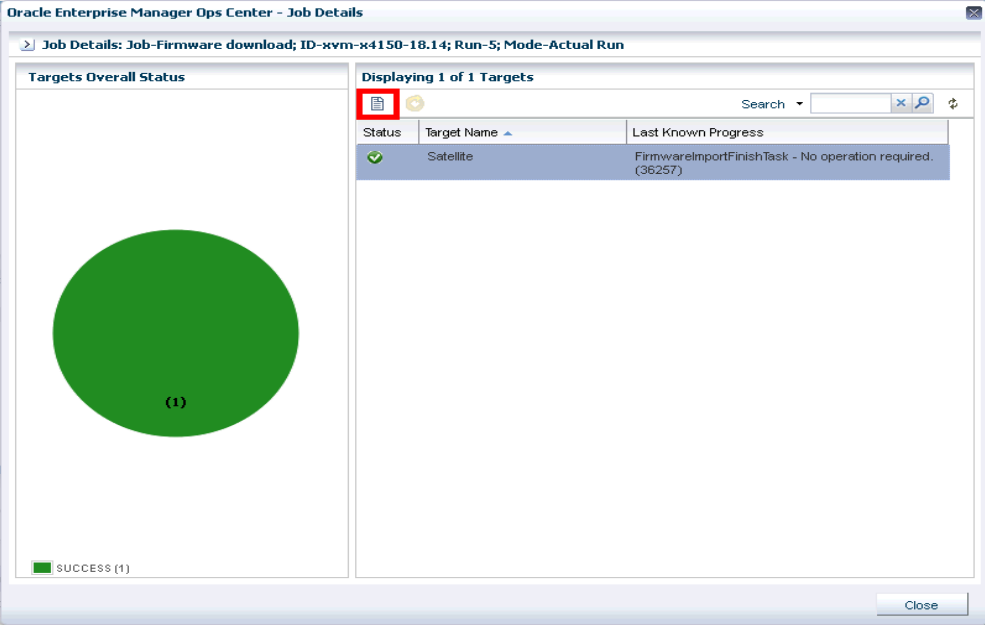
Figure 3–2 View Job Details



Status	Job ID	Job Type	Job Name	Mode	Owner
✓	14	Firmware download	Cron scheduled Firm...	✚	root
⚠	175	ASR Asset Activation	None Target Job	✚	[Ops Center]
✖	174	Shutdown guest	-	✚	root
⚠	173	createLinkAggregation	-	✚	root

2. Click the View Job Details icon, highlighted in [Figure 3–2](#) or double-click the job. The Job Details window is displayed.

Figure 3–3 Job Details Window



Status	Target Name	Last Known Progress
✓	Satellite	FirmwareImportFinishTask - No operation required. (36257)

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

3. Select a target and click **Display Selected Target Details** icon, highlighted in [Figure 3–3](#). You can also double-click the target to display the target details.
The Job Target Details window is displayed.

Figure 3–4 Job Target Details Window

Task	Target of the task	Result	Elapsed Time
✓ Satellite	Satellite	Flow execution is successful	00 hr, 21 min, 38 sec
✓ CompositeTierFreeTask			-
✓ Firmware Nightly Download Satellite->EC		No new firmware images are available for import. (36056)	00 hr, 21 min, 37 sec
✓ CompositeTierFreeTask			-
✓ SimpleCopyTask	Initial EC Library	Task completed successfully. (15029)	Less than a second
✓ ISOImageCRUDTask	Initial EC Library	Task completed successfully. (15029)	Less than a second
✓ FirmwareImportFinishTask	Satellite->EC	No operation required. (36257)	Less than a second

4. Select a target and click **Display Selected Task Details** icon. You can also double-click a task.
The Task Logs window is displayed with the Event Logs tab active, as shown in [Figure 3–5](#).

Figure 3–5 Task Logs Window

Oracle Enterprise Manager Ops Center - Task Logs: ConnectionModeTask-10.79.204.252-EC		
Event Logs		
Date	Severity	Message
04/26/2013 12:12:41 pm EDT	INFO	Setting connection mode.
04/26/2013 12:13:11 pm EDT	INFO	Setting connection mode to connected.

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

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

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

Monitoring Jobs for an Asset

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

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

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

Figure 3–6 Jobs pane in Assets View

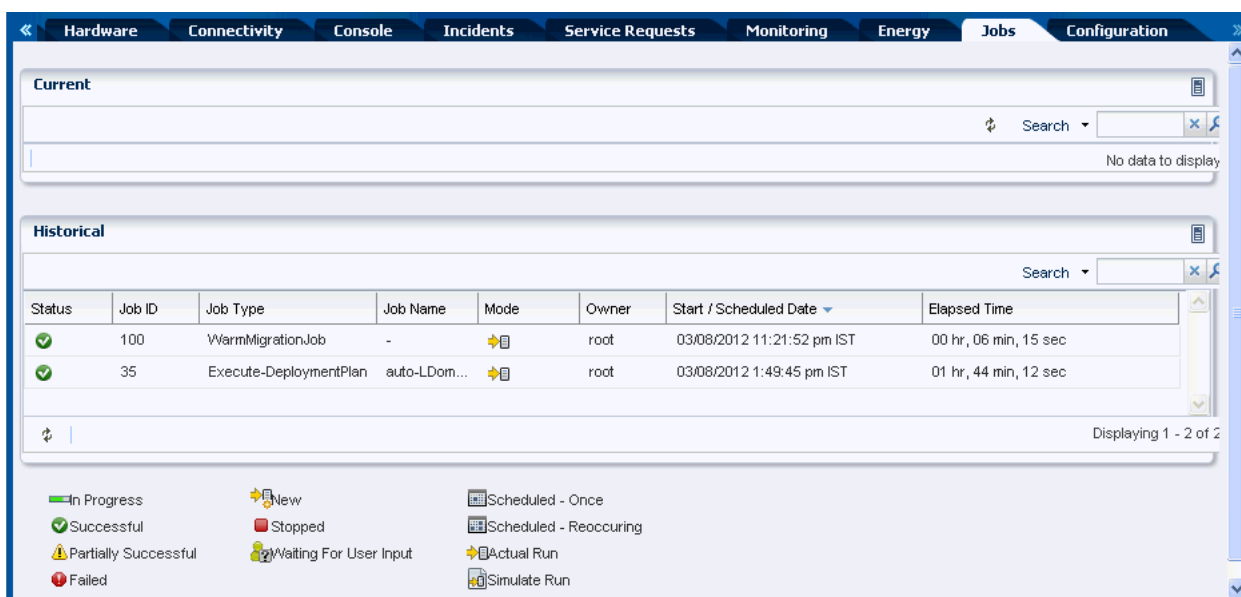


Figure 3–6 shows the Jobs pane for a selected asset. The Current and Historical tables show the jobs for the asset.

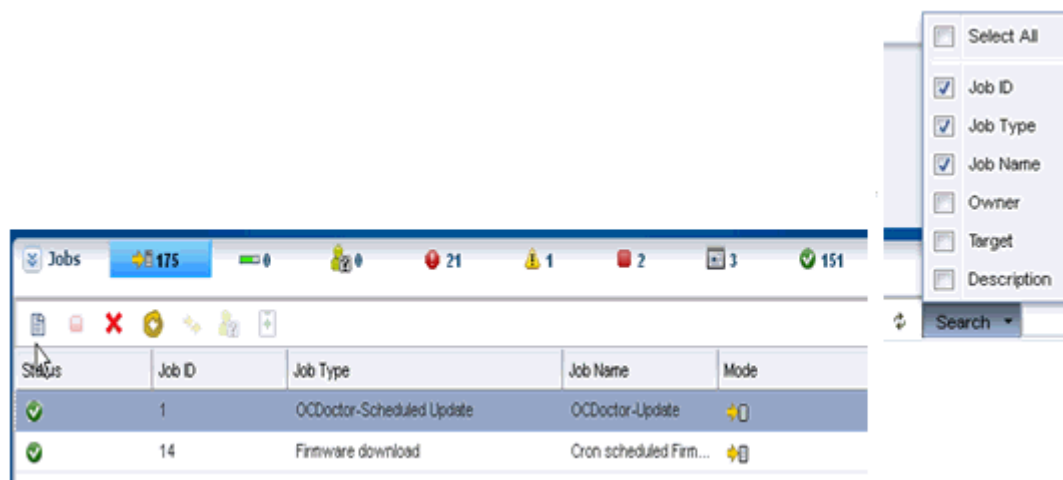
3. Select a job and double click to view the details of the job.

Search for Jobs

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

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

Figure 3–7 Jobs Search



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

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

Taking Actions on a Job

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

Answering Questions for a Job

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

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

Stopping a Job

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

To stop a job, perform the following steps:

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

Re-running a Completed Job

If a job is completed, you can repeat the job with a new Run ID. To re-run a job, perform the following steps:

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

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

Re-running a Job on Failed Targets

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

1. Select the job and view the job details.

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

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

Copying a Job

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

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

A new job wizard is displayed, using the information from the existing job.

2. Modify the job information, select a schedule, and click **Submit Job**.

The new job is submitted with a new Job ID.

Deleting a Job

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

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

To delete a job, perform the following steps:

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

Debugging a Job Using the OCDoctor

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

To debug a job, perform the following steps:

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

Changing the Job Properties

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

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

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

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

Events for Job Management

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

```
/var/cacao/instances/oem-ec/logs/cacao.*
```

Related Resources for Job Management

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

Oracle Enterprise Manager Ops Center Administration Guide

Part II

Configure

Part II contains the following chapters:

- [Chapter 4, "Monitoring Rules and Policies"](#)
- [Chapter 5, "Software Libraries"](#)
- [Chapter 6, "Storage Libraries"](#)
- [Chapter 7, "Networks"](#)
- [Chapter 8, "Plans and Profiles"](#)

Monitoring Rules and Policies

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

The following information is included:

- [Introduction to Monitoring](#)
- [Roles for Monitoring](#)
- [Actions for Monitoring Rules and Policies](#)
- [Location of Monitoring Rules and Policies in the User Interface](#)
- [Monitoring Rules](#)
- [Monitoring Policies](#)
- [Related Resources for Monitoring Rules and Policies](#)

Introduction to Monitoring

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

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

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

Monitoring typically uses the Agent Controller to gather information. When an operating system is an agentless managed system, the software monitors it remotely.

Roles for Monitoring

[Table 4–1](#) lists the roles required to complete monitoring tasks. You can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise*

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

Table 4–1 Monitoring Tasks and Roles

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

Actions for Monitoring Rules and Policies

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

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

Location of Monitoring Rules and Policies in the User Interface

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

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

Monitoring Rules

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

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

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

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

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

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

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

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

The following categories of monitoring rules are available:

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

User-Defined Rule Parameters

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

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

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

The following are some examples of editable parameters:

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

Enabled and Active Rules

Monitoring rules have two types of states:

- **Enabled or Disabled:** Disabling a rule removes that attribute from monitoring. You can disable and enable rules on a per asset or group basis.
- **Active or Inactive:** Reflects the system's state and indicates whether the software is monitoring the asset or group. When a rule is not enabled, monitoring is not active.

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

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

You can perform the following tasks:

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

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

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

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

To Edit a Monitoring Rule

Perform the following steps to edit a monitoring rule:

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

See [Appendix A, "Asset Attributes"](#) for more about the attributes that you can use in monitoring rules.

Using Historical Data to Determine Threshold Limits for a Specific Asset

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

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

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

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

Adding Monitoring Rules

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

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

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

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

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

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

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

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

- Simple: You can use the name, such as `SystemUpTime`
- Struct-like: You can drill-down into a field of the structure type. The fields always start with a lower-case letter, such as `SystemLoad.average1Minute`

- **Collection:** You can drill-down into a member of the collection. For Maps you do this by specifying the appropriate 'key'. When you set the key, specify the value for the 'name' field, to get the value of a single member. For example, use the following to check the 'enabled' value of the interface named *eth0*:
`InterfaceInfos.name=eth0.enabled`

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

You can perform a query which scans across all members of a collection by specifying the '*' wildcard value for the key or name. When you perform a query, you must use one of the following operators: *max*, *min*, or *like*. For the query syntax, see the *DomainQuery Java class* Javadoc. See [Appendix A, "Asset Attributes"](#) for details on the Javadocs.

See [Appendix G, "Expression Query Language"](#) for more information on the query language, grammar, lexical elements, and method details.

Monitoring Policies

A monitoring policy contains the information needed to monitor a managed asset with user-defined alert configurations, including defined thresholds and alert monitors.

A policy is for a specific type of resource, such as operating systems. A more specific policy might apply to a subset of operating systems, such as all Oracle Solaris operating systems. Each monitoring policy contains a number of alert monitors for a specific type of resource. Alert monitors watch the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

Each Monitoring policy contains a default set of rules. The default rule set and parameters depend on the managed asset subtype. You can apply any monitoring policy to a managed asset or group of assets to change the monitoring configuration on the subset of those resources of a compatible type.

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

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

You can use the following methods to create a policy:

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

You can view a policy from the asset view or the policy view. The asset view is

Extracting a Monitoring Policy is similar to copying a policy. You perform the copy policy action from the policy view and the extract policy action from the asset view. By extracting a policy from the asset view, you ensure that the subtypes are valid for that asset type. You can change the target subtype to a more specific or a more generic eligible target type. To be eligible, the policy must be a member of a more general policy for the specified target.

Default Monitoring Policies

When you discover and add an asset, the software applies the default monitoring policy for the asset's type immediately.

Some monitoring policies install probes or agents on managed resources, while other policies are designed to invoke arbitrary actions or scripts against the managed resource.

Available Monitoring Policies

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

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

- **User-Defined Policies:** Typically use generic rules and contain monitoring thresholds. A Operating System policy is an example of a user-defined policy that monitors the following generic operating system parameters: CPU usage, disk IO queue length and utilization percentage, file system usage, memory usage, network bandwidth utilization, and swap usage, and system load.
- **System Defined Policies:** A system-defined policy is read-only, you cannot disable or modify the rules defined in the policy. Typically use asset specific rules. The monitored details are determined by what is implemented on the specific asset. An example of a system-defined policy is the MSeriesChassis policy that monitors the Xsb Mode of a SPARC M-series chassis. The naming convention for system-defined policies is that the name always begins with the letters OC.

You cannot configure this type of rule, you can only turn the rule off and on.

The Navigation pane and center pane both display user-defined policies followed by system-defined policies. The icons and naming convention help to identify the type of policy in the Navigation pane. All system-defined profiles use the prefix OC. For example, in [Figure 4-2 Copy of OC - Chassis and OS Profile - Dev](#) are user-defined profiles and *OC - Chassis* is a system-defined profile. The policies appear in different sections in the center pane.

Figure 4–2 List of Monitoring Profiles


Monitoring Policy Details

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

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

Monitoring policy details vary, depending on the asset and associated resources being monitored. [Figure 4–3](#) is an example of a system-defined operating system monitoring policy, including the Policy Details. The Details section of the page contains information about the policy, including the name, description, and the type of policy. The policy details also describes the applicable types of assets or targets, whether the policy is a sub-type of another policy, when the policy was last modified, and whether the policy is the default policy for the target types.

Figure 4–3 System-Defined Operating System Monitoring Policy

Name:	OC - Operating System	Subtype:	Operating System
Description:	Generic Operating System	Nature:	System defined
Target Type:	Operating Systems	Last Modified:	03/19/2012 11:24:01 pm MDT
<input checked="" type="checkbox"/> Default monitoring policy for assets matching subtype and target type			
Policy Details			
Alert Monitoring Rules ▲		Alert Limits	Enabled ?
BooleanControl Monitoring Rules (1)			
File System Reachability		 Critical: false	Yes
Immediate Action: N/A			
Reachability Monitoring Rules (1)			
Operating System Reachability			Yes

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

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

- OC – Server: Monitors the following parameters on a generic server: Power status, server port status, CPU, NIC, fan, fan tray, memory, and power supply.
- OC – Server Pool: Monitors the total CPU allocation percentage, total cryptographic units allocation percentage, and the total memory usage percentage.
- OC – Solaris Cluster: Monitors the Oracle Solaris cluster install mode, reachability, and monitor.
- OC – Solaris Cluster Node: Monitors the online status, scalert.node reachability, and scalert.node alert monitor.
- OC – Solaris Cluster Zone Cluster Group: Monitors the scalert.zone cluster alert monitor.
- OC – Solaris Cluster Zone Cluster Node: Monitors the scalert.zone Cluster Node Alert Monitor
- OC – Storage: Monitors the following parameters for a generic storage device: operating system status, power status, server port status, and storage alert.
- OC – Switch: Monitors an Ethernet switch's power status, switch port status, and switch status.
- OC – Virtual Machine: Monitors the reachability of the virtual machine, if the operating system is running, and the migration status of a virtual machine.

Applying a Monitoring Policy to a Group

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

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

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

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

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

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

Creating a Group of Assets According to Monitoring Policy

When you create a user-defined group, you can specify a rule that filters out the assets based on the name of the associated monitoring policy. You can use this to group all

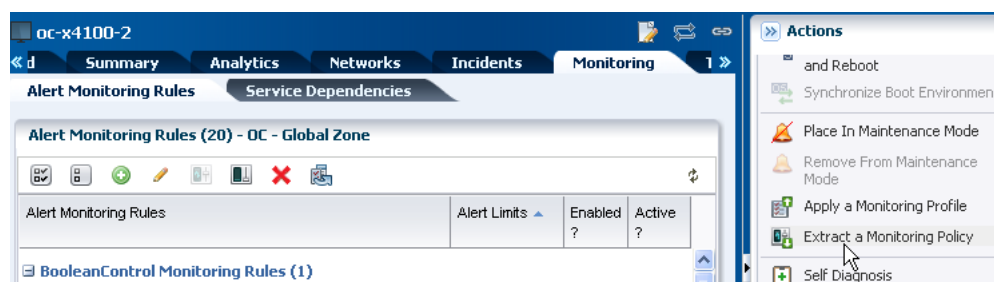
the assets associated with a given monitoring policy or to find assets that are not associated with a specific policy.

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

Extracting a Monitoring Policy

Extracting a monitoring policy is similar to copying a policy. Copying a policy is performed from the policy view, while extracting a policy is performed from the asset view, as shown in [Figure 4-4](#).

Figure 4-4 *Extract a Monitoring Policy Action*



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

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

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

Perform the following steps to extract a monitoring policy:

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

Figure 4–5 Extract Monitoring Policy

Extract Policy - Monitoring

Steps **Help**

1. Identify Policy
2. Summary

Identify Policy

* **Name:** S10 zone monitoring

Description: Use this policy to monitor Oracle Solaris 10 zones.

Subtype: Subtype
Global Zone

Target Type: Target Type
Operating Systems

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

Adding or Copying a Monitoring Policy

You can add a new monitoring policy or copy an existing monitoring policy.

Perform the following steps to add or copy a monitoring policy:

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

Deleting a Monitoring Policy

You can delete user-defined monitoring policies in the Plan Management section. Resources are associated with a monitoring policy. When you delete a user-defined

monitoring policy, any asset or group associated with that policy is automatically detached from the policy. However, the asset does retain the monitoring configuration that was defined by the policy.

You cannot delete system-defined policies.

Related Resources for Monitoring Rules and Policies

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

- See [Chapter 11, "Hardware"](#) for information about hardware monitoring.
- See [Chapter 12, "Operating Systems"](#) for information about monitoring operating systems.
- See [Chapter 9, "Incidents"](#) for information about managing incidents that result from monitoring.
- See [Appendix G, "Expression Query Language"](#) for details about the expression query language, grammar, lexical elements, and method details.
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm. The following How To documents relate to monitoring:
 - Understanding OS Performance and Capacity
 - Tuning Monitoring Rules and Policies

Software Libraries

The following information is included:

- [Introduction to Software Libraries and Repositories](#)
- [Roles for Software Libraries](#)
- [Actions for Software Libraries](#)
- [Location of Software Library Information in the User Interface](#)
- [Knowledge Base and Parent Repository](#)
- [Using Software Libraries](#)
- [Libraries for Oracle Solaris 11](#)
- [Libraries for Oracle Solaris 10, 9, 8 and Linux](#)
- [Images](#)
- [Local Content](#)
- [Backing Up Images and Local Content](#)
- [Related Resources for Software Libraries](#)

Introduction to Software Libraries and Repositories

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

A Software Library accepts the following types of images:

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

Another type of library is a storage library. See [Chapter 6, "Storage Libraries"](#) for information about this type of library.

Roles for Software Libraries

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

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

Table 5–1 Software Libraries Tasks and Roles

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

Actions for Software Libraries

Perform the following actions, depending on the requirements:

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

Location of Software Library Information in the User Interface

[Table 5–2](#) shows where to find information.

Table 5–2 Location of Library Information in the BUI

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

Knowledge Base and Parent Repository

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

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

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


EC Library for Oracle Solaris 10, 9, and 8 and Oracle Linux OS

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

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

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

Figure 5–1 Badge Identifying the Current Default Library



To Change EC Software Libraries

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

1. Create a software library, as described in [Creating a Software Library](#).
2. Expand Administration in the Navigation pane.
3. Click **Set Enterprise Controller Storage Library** in the Action pane. The window lists all libraries with the current library highlighted.
4. Click the new software library.
5. Click Apply.

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

Publishers and Parent of the Oracle Solaris 11 Repository

The Oracle Solaris 11 Package Repository resides at the Oracle site:

<http://pkg.oracle.com/solaris/release/>. This repository has the role of Publisher

and Parent for your local repository. The Oracle Solaris 11 Package repository is similar to the Knowledge Base but is only for Oracle Solaris 11 images and updates.

A component of an Oracle Solaris 11 update package is the IPS Package Manifest. The manifest consists of the following:

- Metadata for the package such as name, description, version, and category
- Files and directories that are included
- Dependencies on other packages
- Instructions to specify the services to restart on the system, to specify the aliases to update for a hardware driver, or to create the users and groups.

Other parent repositories, such as Oracle Solaris Cluster, are available. You can configure more than one parent repository and provide their credentials. For a list of available Oracle repositories and to download the key-certificate pair, see the `pkg-register.oracle.com` site.

Using Software Libraries

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

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

Viewing the Contents of a Software Library

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

To View the Contents of a Software Library

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

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

The Summary tab displays information about the entire software library:

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

The Library Contents table lists all the images in the library, organized by type, and includes the size and the date the image was modified. The other tables on the Summary tab describe the types of images: Service Processor Firmware, Component Firmware, and BIOS Configuration Snapshots

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

Creating a Software Library

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

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

Creating a Local Software Library

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

Creating a NAS Software Library

1. Expand Libraries in the Navigation pane.
2. Click **New NAS Software Library** in the Action pane.
3. Enter a unique name and a description.
4. Choose the type of service:
 - a. To use a unified storage appliance, select an appliance and specify the exported share to use.
 - b. To use an NFS server, enter the host name and path for the NAS appliance.
5. Click Create.

Libraries for Oracle Solaris 11

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

Oracle Solaris 11 Software Update Library

The Enterprise Controller can manage a local copy of a parent repository for the Oracle Solaris 11 Image Packaging System (IPS). The local repository, called the Oracle Solaris 11 Software Update Library, is used to provision the Oracle Solaris 11 operating system on assets. You configure the Oracle Solaris 11 Software Update Library when you install Oracle Enterprise Manager Ops Center, or you configure the library after installation. You assign storage to the Oracle Solaris 11 Software Update Library, and then populate it with a local copy of the IPS.

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

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

- Unconfigured – Oracle Solaris 11 Software Update Library was not configured when you installed and configured Oracle Enterprise Manager Ops Center.
- Configuring – The Oracle Solaris 11 Software Update Library is in the process of being configured.
- Syncing – The Oracle Solaris 11 Software Update Library is in the process of updating and synchronizing with the IPS or is in the process of being reconfigured. The Oracle Solaris 11 Software Update Library is locked and unavailable for use when in this state.
- OK – The Oracle Solaris 11 Software Update Library is ready to use.

If the Oracle Solaris 11 Software Update Library is in an unconfigured state, use the **Initialize Oracle Solaris 11 Software Update Library** action to populate the library with a local copy of the IPS packages. When complete, the new library appears in the Libraries section and contains three sections: Summary, Boot Image, and Packages.

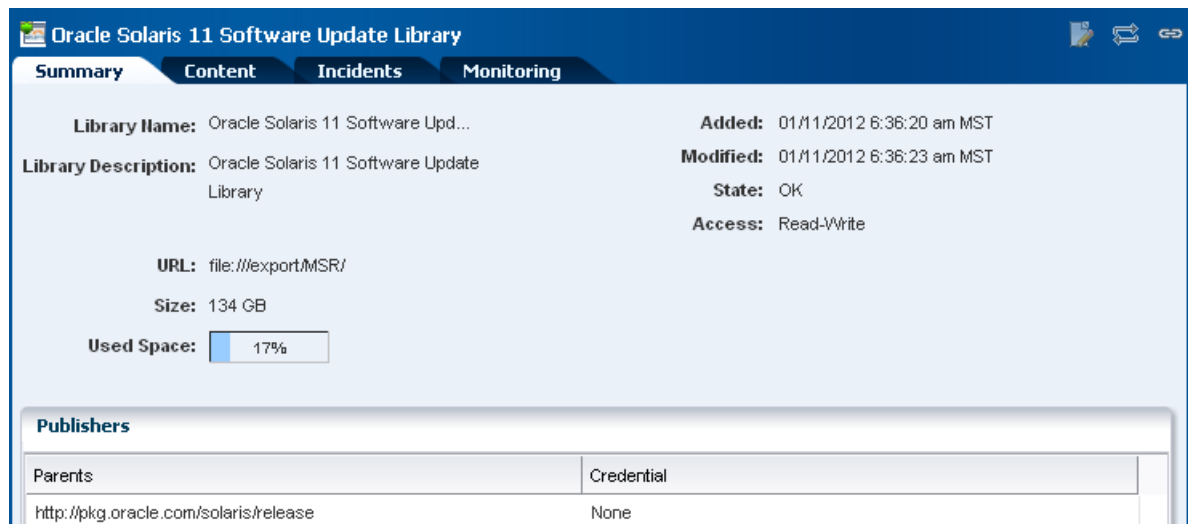
Note: The process of initializing the library takes several hours.

To keep the Oracle Solaris 11 Software Update Library up to date, synchronize it with the parent IPS repository. The Oracle Solaris 11 Software Update Library is locked during the process and you cannot perform OS provisioning or OS update tasks.

Summary of Oracle Solaris 11 Software Update Library

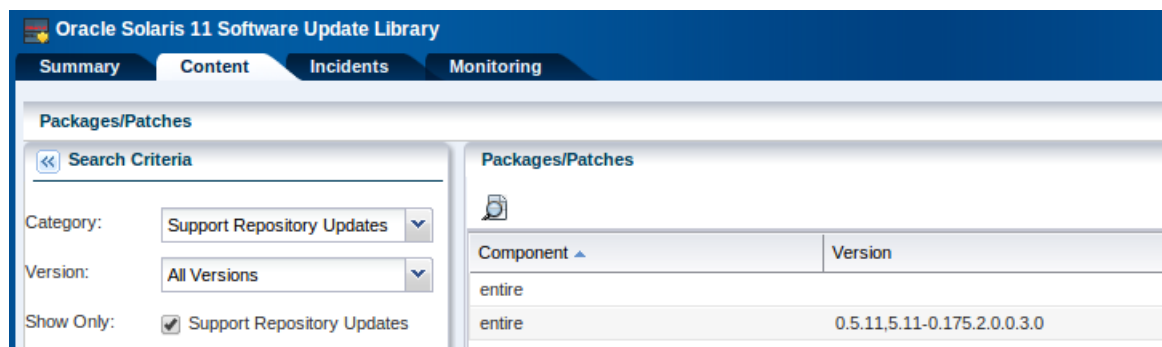
When the Oracle Solaris 11 Software Update Library is configured, its Summary tab, shown in [Figure 5–2](#), provides an overview of the health, status, and last update. The URL listed on the Summary is the location of the Oracle Solaris 11 Software Update Library. The size is the amount of space allocated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

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

Figure 5–2 Oracle Solaris 11 Software Update Library Summary

Content of the Oracle Solaris 11 Software Update Library

The Content tab, shown in [Figure 5–3](#), displays a list of the packages in the Oracle Solaris 11 Software Update Library with a brief description of each package. You can filter the list of packages by selecting a category or version. You can also search for a specific package.

Figure 5–3 Oracle Solaris 11 Software Update Library Content

Configuring Parent Repositories

You manage publishers for the parent repository and synchronize the Oracle Solaris 11 Software Update Library with the parent. You can create a recurring schedule to synchronize the Oracle Solaris 11 Software Update Library during non-peak hours.

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

Using an Alternate IPS Repository

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

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

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

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

1. Locate your site's Oracle Solaris 11 IPS repository or create one.
2. Expand **Libraries** in the Navigation pane.
3. Click **Software Libraries**.
4. Click **Initialize Solaris 11 Software Update Library** in the Action pane.
5. In the **Specify Library Location** field, enter the location of your IPS repository.
6. In the **URL** field for the parent repository, do not enter any text. Leave this field blank.
7. Click **OK**.

Adding Content

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

Deleting Oracle Solaris 11 Software Update Library

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

Libraries for Oracle Solaris 10, 9, 8 and Linux

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

You configure the Linux, Solaris 8-10 Software Update Library when you install Oracle Enterprise Manager Ops Center, or you can configure the library after installation using the **Create Update Library** action.

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

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

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

See [Images](#) for information about managing the content in the library.

See [Local Content](#) for information about managing the content of the local categories.

Images

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

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

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

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

Images for OS Provisioning

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

[Table 5–3](#) shows the location of OS images.

Table 5–3 Location of Images for Each Operating System

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

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

Requirements for OS Images

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


```
# mkisofs -o <name_of_OS.iso> -J -R /cdrom/<name_of_OS>
```
- The Oracle Enterprise Manager Ops Center software uploads or imports one ISO file per operation. If you are loading an ISO file from physical media and the file spans more than one CD, combine the content on one DVD.
- An ISO file cannot be made from Oracle Solaris installation CDs.

Images for Firmware Updates

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

The following firmware types are supported:

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

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

Example 5-1 Example of Firmware Metadata

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

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

```
Latest Sun System Firmware(6.1.2):
```

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

- To determine the models supported:

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

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

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

From this README file, identify the following metadata:

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

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

Uploading or Importing Images

To provision firmware or an OS, you use a deployment plan to direct Oracle Enterprise Manager Ops Center to retrieve the images from the appropriate software library and install them on the targeted assets. Each time you import an image, you also create a profile with the same name. You can use the default profile but, as a good practice, make a copy of the profile and rename the copy for use in deployment plans. Leave the original profile with its original name in the library.

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

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

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

Importing an Image

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

The Import Image window is displayed.

Figure 5–4 Import Image

Oracle Enterprise Manager Ops Center - Import Image

Import Image ? ORACLE

* Indicates Required Field

Image Type: ☒ ISO ☐ Storage Appliance Update ☐ Assembly Template
☐ FLAR ☐ Oracle VM Template ☐ Virtual Disk

Import Source: ☒ A directory that is accessible by Enterprise Controller
☐ Initial EC Library/blobs
☐ Local host/machine
☐ URL

Directory: /

Image to be Imported: **Available images in the selected directory**

Name
README.sol-11_1-upgrade-repo
sol-11_1-repo-full.iso-b
HEADER.html

* **Image Name:**

Description:

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

If the image file is located in an accessible location, enter the name of the directory or use the Browse button to navigate to the location.

7. Select the image you want to import.
8. Enter a name for the image and a description. Image names must be unique, can consist of up to 100 characters, and can include numbers, letters, and some special

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

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

Uploading an Image

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

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

The Upload ISO Image window is displayed.

Figure 5–5 Upload ISO Image

4. Click the name of the file in the Source of ISO File field or click Browse to navigate to the image.
5. Enter the name and description of the image.
6. Click Upload Image. The progress of the upload to the Enterprise Controller is displayed. After uploading, the image is in the software library.

Uploading Firmware Images

The Upload Firmware action can be directed to retrieve firmware images from either the Enterprise Controller's file system or a local file system. Use the procedures in *Keeping Your Firmware Up-to-Date* in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm to upload firmware.

Note: The UploadFirmware action is disabled for Exalogic systems.

Working with Firmware for Power Distribution Units

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

Uploading Firmware With Metadata

Use the procedures in *Keeping Your Firmware Up-to-Date* in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm to upload firmware. At the step for Select Firmware Components, select both firmware components:

- MKAPP_V`firmware_version`.DL
- HTML_V`firmware_version`.DL

Uploading Firmware Without Metadata

Use the procedures in *Keeping Your Firmware Up-to-Date* in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm to upload firmware. At the step for Select Firmware Components, select one firmware component, MKAPP_V`firmware_version`.DL and click Upload. To define the metadata for the firmware:

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

After the job is completed, repeat the procedure with the other firmware image, HTML_V`firmware_version`.DL with one difference: At Step 4, specify that this firmware image depends on the MKAPP_V`firmware_version`.DL image.

Creating a Firmware Profile for PDU Firmware Updates

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

Local Content

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

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

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

Local Categories

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

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

Adding a Local Category

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

Uploading a Local Action

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

- Pre-Actions – Script that runs on a managed host before the provisioning step starts. When you create the deployment plan, you select the Execute Pre-Install step and then select the profile that includes the script.
 - Post-Actions – Script that runs on a managed host after a job is completed. When you create the deployment plan, you select the Execute Post-Install step and then select the profile that includes the script.
 - Probes – Script that runs on a managed host to verify that a job can be performed.
 - Macros – Script that modifies a generic configuration file to make it specific for a managed host. Use macros to apply a single configuration file across multiple hosts by customizing the configuration file for each host's environment. The script outputs a single line that replaces a macro sign in a configuration file.
1. Expand Libraries in the Navigation pane.
 2. Click Linux and Oracle Solaris 8-10 Software Update Library.
 3. Click **Upload Local Action** in the Actions pane.

The Upload Local Action window is displayed in [Figure 5-6](#).

Figure 5–6 Local Upload Action

The screenshot shows a dialog box titled "Oracle Enterprise Manager Ops Center - Upload Local Action". It contains the following fields and controls:

- Action name:** A text input field.
- Description:** A larger text input area.
- Action type:** A dropdown menu currently showing "Pre-actions".
- Distribution:** A list box showing three options: "SOLARIS10_SPARC", "SOLARIS9_SPARC", and "SOLARIS10_X86".
- Parent Category:** A text input field showing "Pre-actions" and a "Browse..." button to its right.
- File to upload:** A text input field and a "Browse..." button to its right.
- Buttons:** "Upload" and "Cancel" buttons at the bottom right of the dialog.

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

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

Uploading a Local Software Package

You can upload software in the following formats:

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

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

1. Expand Libraries in the Navigation pane.
2. Click the operating system's Software Update Library.
3. Click **Upload Local Software Packages** in the Actions pane.

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

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

Uploading a Local Configuration File

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

Uploading a Local Configuration File

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

The Upload Local Configuration File window is displayed in [Figure 5-7](#).

Figure 5-7 Upload Local Configuration File

Oracle Enterprise Manager Ops Center - Upload Local Configuration File

Target path on server	<input type="text"/>	E.g. /etc/hosts
Version	<input type="text"/>	E.g. 2.1 or myversion
Description	<input type="text"/>	
Distribution	SOLARIS10_SPARC SOLARIS9_SPARC SOLARIS10_X86	
Parent Category	Configuration files	<input type="button" value="Browse..."/>
File to upload	<input type="text"/>	<input type="button" value="Browse..."/>

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

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

Uploading Software in Bulk

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

The files must be in the following formats:

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

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

Before You Begin

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

Uploading Local Software in Bulk

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

The Upload Packages, Patches, and RPMS window is displayed in [Figure 5–8](#).

Figure 5–8 Upload Local Software

Oracle Enterprise Manager Ops Center - Upload Patches, Packages and RPMs

Distribution: SOLARIS10_SPARC

☒ Upload from Directory Directory path on Server to load from:

☐ Upload from OS image

Select an OS image to mount and then browse.

Name	Description	Image Type	OS Type	Version	Parent
s10u7sparc	c424e7a7-f7e3-40b8-a825-cb08ee239005	iso	solaris	10sparc	none
oel-server-5-3	a0238d03-6a21-4b02-be59-fd1fe448b0df	iso	oracle	OEL5	none
rhel-server-5-3	7a5d61a6-8cc8-4286-a858-140b4a8517b2	iso	redhat	5	none

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

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

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

Viewing Results of a Bulk Upload Operation

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

To View Bulk Upload Results

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

Using Local Content

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

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

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

- Create the local configuration file and upload it.

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

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

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

Editing Local Content

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

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

Deleting Local Content

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

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

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

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

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

Backing Up Images and Local Content

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

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

Related Resources for Software Libraries

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

http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

- *Keeping Your Firmware Up-to-Date*

Storage Libraries

The following information is included:

- [Introduction to Storage Management](#)
- [Roles for Storage Management](#)
- [Actions for Storage Management](#)
- [Location of Storage Information in the User Interface](#)
- [Storage Libraries](#)
- [Types of Storage for Libraries](#)
- [Storage Hardware](#)
- [Storage Profiles](#)
- [High Availability for Storage Resources](#)
- [Related Resources for Storage Management](#)

Introduction to Storage Management

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

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

Roles for Storage Management

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

Table 6–1 Storage Tasks and Roles

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

Actions for Storage Management

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

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

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

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

Location of Storage Information in the User Interface

[Table 6–2](#) shows where to find information.

Table 6–2 Location of Library Information in the BUI

Object	Location
To see the storage libraries	Expand Libraries in the Navigation pane, then expand Storage Libraries.
To see storage hardware information	Expand All Assets in the Navigation pane, then scroll to the Storage section. Select a physical asset to see information about it in the center pane.
To see groups of storage arrays or file servers	Expand All Assets in the Navigation pane, then Storage in the Resource Management Views section.

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

Object	Location
To see storage profiles	Expand Plan Management in the Navigation pane, then click Profiles and Policies. The Discovery, RAID Controller, and Update Storage Appliance categories contain profiles.
To see the virtual host that is using a LUN	Expand Libraries in the Navigation pane, then expand Storage Libraries. Select the storage library to view the LUNs table in the center pane. The Allocated To field displays the virtual host's identifier.
To see the virtualization host that is using a storage library	Expand Libraries in the Navigation pane, then expand Storage Libraries. Select the storage library to view the Summary tab in the center pane. The IP address of the virtualization host is in the Allocated To field.

Storage Libraries


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

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


- For the metadata of its local virtual hosts, virtualization hosts (Oracle VM Server for SPARC, Oracle VM Server for x86, and global zones) can use a local storage library.
- For the metadata of all virtual hosts, virtualization hosts must use a Network Attached Storage (NAS) storage library.
- For data, Oracle VM Server for SPARC and Oracle VM Server for x86 can use either NAS shares or SAN LUNs for itself and for its virtual hosts. A global zone can use SAN LUNs for itself and for its non-global zones. The storage library must be associated with the virtualization host.

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

- SAN storage libraries expose data as virtual disks (LUNs), identified by their LUN GUIDs. A control domain makes raw partitions available to its logical domains using Fibre Channel or iSCSI, as described in [Block Storage](#)
- NAS storage libraries expose data as raw files and files in VMDK format. All types of virtualization hosts store metadata using NFS services, as described in [Network Attached Storage \(NAS\) Storage Libraries](#)

Figure 6–1 SAN and NAS Storage Libraries

The LUNs, raw files, and raw volumes store data and metadata for the virtual hosts. [Figure 6–2](#) expands [Figure 6–1](#) to show that the metadata for NAS virtual disks are stored in the NAS storage library. Metadata for SAN virtual disks is persisted in the SAN Storage Library.

Figure 6–2 Storage Library Metadata

Network Attached Storage (NAS) Storage Libraries

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

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

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


- If the storage library is used for the virtual hosts' metadata, a virtual host continues to function but Oracle Enterprise Manager Ops Center can no longer manage the virtual host. Because Oracle Enterprise Manager Ops Center relies on its interaction with the metadata in the storage library, jobs that must read or modify the metadata fail. You must manage the virtual host manually.
- If the storage library is used for NFS large files that support virtual disks, the virtual host does not function.
- If the boot disk is on the NFS share, the virtual host cannot be rebooted.

- The virtual host cannot be migrated.

Storage Libraries for Server Pools

You can group virtualization hosts with the same processor architecture to create a server pool. The virtualization hosts share all the storage and networks associated with the server pool. If you add a virtualization host to a server pool, the libraries associated with that virtualization host become available to all the other hosts in the server pool. A server pool must use a NAS storage library. [Figure 6–3](#) shows how virtualization hosts in a server pool get access to storage resources through the storage libraries.


Figure 6–3 External Storage



Storage Libraries for a Virtual Datacenter

The vDC inherits the storage resources allocated for the server pool. See [Storage Setup](#) in [Chapter 18](#).

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

Figure 6–4 Graph of Guest Storage Resource Allocation

Storage Libraries for Oracle Solaris Zones

A global zone provides storage resources to its non-global zones. The Oracle Solaris Zone must be associated with a storage library. See [Associating Storage Libraries](#) in [Chapter 14](#).

Storage Libraries for Oracle VM Server for SPARC

The control domain provides storage resources to its logical domain. The control domain must be associated with a storage library. See [Associate Storage Libraries](#) in [Chapter 15](#).

Storage Libraries and Repositories for Oracle VM Server for x86

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

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

Creating an Oracle VM Storage Repository

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

Uploading Templates and Assemblies to an Oracle VM Storage Repository

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

The Import Image window is displayed.

Figure 6–5 Import Image

Oracle Enterprise Manager Ops Center - Import Image

Import Image ? ORACLE

* Indicates Required Field

Image Type: ☒ ISO ☐ Storage Appliance Update ☐ Assembly Template
☐ FLAR ☐ Oracle VM Template
☐ Virtual Disk

Import Source: ☒ A directory that is accessible by Enterprise Controller
☐ Initial EC Library/blobs
☐ Local host/machine
☐ URL

Directory: /

Image to be Imported: **Available images in the selected directory**

Name
README.sol-11_1-upgrade-repo
sol-11_1-repo-full.iso-b
HEADER.html

* **Image Name:**

Description:

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

Types of Storage for Libraries

- [File Systems Libraries](#)
- [Block Storage](#)

File Systems Libraries

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

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

Viewing Local Libraries

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

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

Editing the Attributes of a Local Library

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

1. Expand Assets in the Navigation pane.
2. Select the asset.
3. Click the **Libraries** tab in the center pane. The asset's associated libraries and the guests that are stored in the libraries are listed.
4. Click the **Edit Local Library** icon.
5. In the Edit Local Library pane, enter the new name or description for the library.
6. Click the **Update** button. When the job is completed, the edited local library is listed in the Associated Libraries table.

Creating a Local Library

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

1. Create a file system with read/write permissions for only the root user.

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

Deleting a Local Library

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

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

Block Storage

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

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

- Disk – Physical storage media. A set of disks is a disk array.
- Volume – An aggregation of storage space provided by several disks.
- Slice – A partition of a volume that is exposed to the servers connected to the disk array.
- LUN (Logical Unit Number) – The representation of a slice.
- GUID – The Global Unique Identifier for a LUN.

If the storage array is a managed asset, the LUNs can provide dynamic storage.

[Table 6–3](#) compares static block storage libraries and dynamic block storage libraries.

Dynamic Block Storage

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

Static Block Storage

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

Table 6–3 Comparison of Static and Dynamic Libraries

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

Selecting LUNs For the Block Storage Library

Use this procedure in the following situations:

- You are creating a new SAN storage library and have associated the new library with a virtualization host or server pool immediately.
- You are increasing the storage capacity of an existing SAN storage library and chose to select LUNs from the available LUNs.

The Available LUNs table lists all of the LUNs that are accessible from either the associated library or from the pools, hosts, or zones you selected. For each LUN, the following information is displayed to help you identify the ones you want to add to the library.

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

- **Size (GB)** – Size of each LUN in Gigabytes.
- 1. Click one or more LUNs. Use the Search box to locate a specific LUN. When you finished, click Next or Add LUN.
- 2. If you are adding a LUN to an existing storage library, a new job starts. If you are creating a SAN storage library, review the details of the LUNs you have configured in the Summary pane.

Click Finish to close the library.

Adding Capacity to Dynamic Block Storage Libraries

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

1. Expand Libraries in the Navigation pane.
2. Click Block Storage in Storage Libraries.
3. Click one of the available storage libraries.
4. Click **Add LUN** in the Actions pane.
5. Choose the method for adding LUNs:
 - To add LUNs to the library manually, accept the default option: **Manually enter the GUID/WWN of the LUNs to be added**. See [To Add LUNs to a Block Storage Library Manually](#) to complete the procedure.
 - To select LUNs from the available LUNs, click the **Select from available LUNs** option. See [Selecting LUNs For the Block Storage Library](#) to complete the procedure.

To Add LUNs to a Block Storage Library Manually

You use this procedure in the following situations:

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

After you have accepted the default option to add the LUNs by name, the table shown in [Figure 6–6](#) is displayed:

Figure 6–6 Adding LUNs Manually

Library Name : **FCLib-11102010-153513**

Description :

NOTE: You can only select from available LUNs if the Fibre Channel library is already associated with one or more virtual pools, global containers or virtual servers.

☒ Manually enter the GUID/WWN of the LUNs to be added

☐ Add one or more LUNs to the library. Use Ctrl+Click and Shift+Click to select multiple LUNs

LUNs to be added to the library	
LUN GUID	LUN Name
	FCLib-11102010-153513-LUN1

Add LUNs **Cancel**

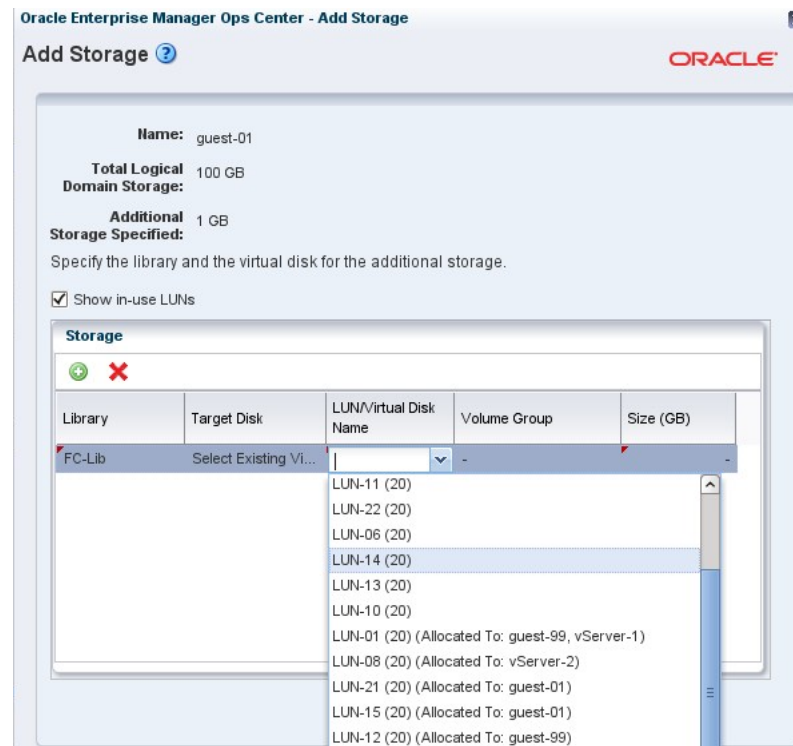
1. Click the icon to create an entry for the LUN.
2. Click in the GUID/WWN field and type the GUID or WWN for the LUN. The GUID is the Global Unique Identifier associated with each LUN, which is a hexadecimal number of 32 digits. If your site uses SCSI initiators and targets, you can enter the WWN for the LUN. You can edit the default name for each LUN. Click Next or Add LUN.
3. If you are creating the storage library, review the details in the Summary pane.
4. Click Finish to close the Fibre Channel library.

Adding Storage to a Logical Domain

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

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

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

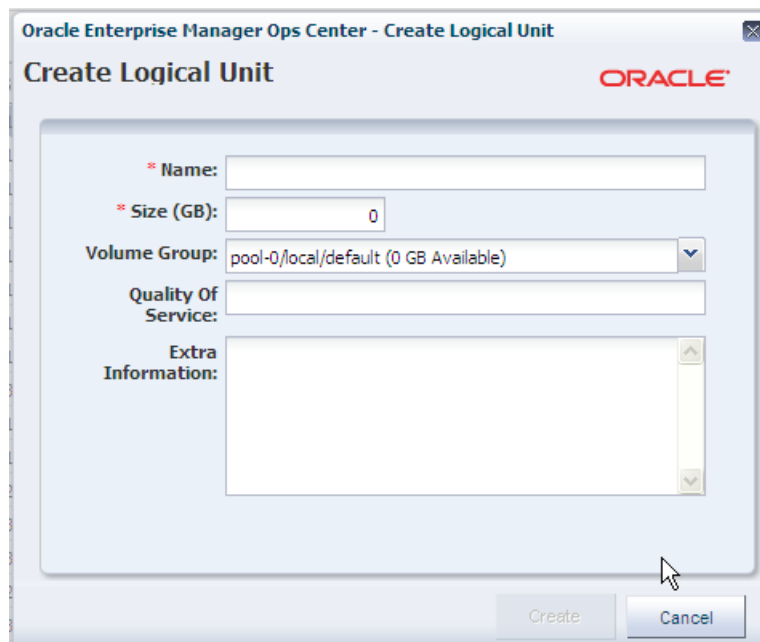
Figure 6–7 Adding LUNs

4. Select the library.
5. Select the LUN or virtual disk.
6. Click OK.

Create a LUN

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

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


Figure 6–8 Creating a LUNThe image shows a screenshot of the 'Create Logical Unit' dialog box in the Oracle Enterprise Manager Ops Center. The dialog has a title bar that says 'Oracle Enterprise Manager Ops Center - Create Logical Unit' and a close button. Below the title bar, the text 'Create Logical Unit' is displayed next to the Oracle logo. The main area of the dialog contains several input fields: a text field for 'Name' with a red asterisk, a text field for 'Size (GB)' with a red asterisk and the value '0', a dropdown menu for 'Volume Group' showing 'pool-0/local/default (0 GB Available)', a text field for 'Quality Of Service', and a large text area for 'Extra Information'. At the bottom right, there are two buttons: 'Create' and 'Cancel'. A mouse cursor is pointing at the 'Cancel' button.

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

Clone a LUN

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

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

Figure 6–9 Cloning a LUN

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

Oracle VM Storage Connect Plug-ins

Oracle Enterprise Manager Ops Center shares the capability of Oracle VM Manager to manage storage devices of various vendors. Oracle VM Storage Connect is an application programming interface (API) that exposes the storage device's features and attributes to Oracle Enterprise Manager Ops Center. When a storage vendor creates an Oracle VM Storage Connect plug-in, you install the plug-in so that you can manage the vendor's storage products. The storage devices are managed through the Oracle Enterprise Manager Ops Center's user interface. The plug-in software is available from the vendor.

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

1. Expand Administration in the Assets pane.
2. Select Enterprise Controller.
3. Click the **Proxy Controller** tab in the center pane.

The Installed StorageConnect Plugins section of the center pane lists the plug-in software and its details.

Storage Hardware

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

RAID Controller

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

- RAID Volumes
- RAID Levels
- Stripe size
- Number of disks

NAS Storage Appliance

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

The procedure for setting up the share for a library depends on several site-specific factors such as the version of NFS protocol and name service management. The example in this section provides one method of configuring the share on an NFS server running on the Oracle Solaris 10 operating system. See *Oracle Solaris Administration: Network Services* for the information about the Oracle Solaris 11 procedure.

Setting Up a Share on an NFS Server

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

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

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

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

See the `share_nfs(1M)` man page for information about NFS share options, and how to specify the access list.

3. Share the directory and verify that the directory is shared. For example:

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

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

Setting Up an NFS Client

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

```
# showmount -e <server-name>
export list for <server-name>:
```

/export/virtlib/lib0 (everyone)

Creating a NAS Storage Library

A NAS Storage Library is the same as a NAS Software Library but its purpose is different. The software library holds images for provisioning. The storage library holds metadata and data.

1. Expand Libraries in the Navigation pane.
2. Click Storage Libraries
3. Click **New NAS Storage** in the Action pane.

Figure 6–10 shows the Create New NAS library window.

Figure 6–10 Create New NAS Library

Oracle Enterprise Manager Ops Center - Create New NAS Library

New NAS Library

Library Name:

Library Description:

Server Pool Association: **Select the server pool you want to be associated with this library. Use Ctrl+Click and Shift+Click to select multiple server pools.**

Name	Description
NORM-VirtPool-0	
NORM-VirtPool-1	

Select the type of service:

☒ Exported share on a managed asset

Select a storage asset and then select the share that provides storage.

Asset:

Share:

☐ Other

Replace hostname, port, and path with the values for the storage resource.

URL:

4. Enter a name for the library and a description.
5. Select the server pool to be associated with this library.
6. For the Type, you have a choice of either entering the path of a shared directory or, if Oracle Enterprise Manager Ops Center is managing a unified storage asset, selecting the asset and specifying the exported share to use.

7. Click **New NAS Library** to create the library.

File Server

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

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

Storage Array

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

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

To use LUNs backed by Fibre Channel disks, you must enable multipathing on the Fibre Channel storage device or on its individual ports. Multipathed I/O (MPxIO) allows I/O devices to be accessed through multiple host controller interfaces. Multipathing is disabled by default for Fibre Channel devices on SPARC-based systems, but is enabled by default on x86-based systems. You use the `stmsboot (1m)` command to configure multipathing.

For systems using Oracle Solaris 11, see *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing*.

Storage Server: Sun ZFS Storage Appliance

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

This type of storage appliance contains a service processor, which must also be discovered. See [Sun ZFS Storage Appliance](#) for this procedure.

Managing a Sun ZFS Storage Appliance

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

- **Launch Appliance UI** opens a new browser window or tab for the main page.
- **Launch Detailed Dashboard** opens a new browser window or tab for the status page.
- **Launch Analytics** opens a new browser window or tab for the dynamic analysis page.

- **Manage Shares** opens a new browser window or tab for the share configuration page.
- **Manage Services** opens a new browser window or tab for the data services configuration page.

Displaying Volume Capacity

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

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

Auto Service Requests for the Sun ZFS Storage Appliance

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

Provisioning and Updating a Sun ZFS Storage Appliance

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

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

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

1. Expand Assets and then Storage in the Navigation pane.
2. Select the storage appliance.
3. Click the **System Details** tab in the center pane.
4. In the Appliance Updates table, select the update to remove.

5. Click the **Rollback Appliance Software** icon.
6. Confirm the action to submit the job.

Exadata Storage Server

The Oracle Exadata Storage Server is a fast, high capacity storage server. Each server comes has 12 disks connected to a storage controller. Exadata has three disk groups that span all of the disks so that each physical disk is in each disk group. Having all physical disks participating in data retrieval and storage maximizes performance.

Oracle Enterprise Manager Ops Center discovers the Exadata Storage Server and its disks when it manages a SuperCluster Engineered System. See [Chapter 20, "Oracle Engineered Systems"](#) for more information.

Storage Profiles

Oracle Enterprise Manager Ops Center provides default profiles for the following operations:

- Configuring RAID Controllers
- Updating the Sun ZFS Storage Appliance

High Availability for Storage Resources

Storage devices for an HA configuration must meet these requirements:

- Storage must be transferable between the primary and secondary Enterprise Controller systems. Do not attempt to use local file systems for high-availability storage.
- Storage must offer data redundancy capability, such as mirroring or RAID 5.
- Storage must offer performance that is sufficient to support operations.
- Storage must have the capacity to hold the data that the Oracle Enterprise Manager Ops Center software stores in the `/var/opt/sun/xvm` directory structure.

A variety of storage solutions meet these criteria, including NAS appliances, hardware RAID arrays and external JBODs. Storage can be attached directly to the Enterprise Controllers or through Storage Area Networks.

You must determine what storage solution offers the required capacity, performance, connectivity, and redundancy capabilities. Configuration procedures vary greatly among the available storage solutions, and among operating systems.

You must determine the specific failover procedures to use for the HA storage solution. Contact My Oracle Support to determine the procedures to use for your particular installation.

Related Resources for Storage Management

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

- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm
- For information about storage hardware, see [Chapter 11](#).

- To view the *Sun ZFS Storage 7000 System Administration Guide*, log in to the Unified Storage System software interface and click Help in the top right corner of any screen. You can also access this guide at the host name or IP address of the storage system:

- `https://hostname:215/wiki`

- `https://ipaddress:215/wiki`

See the Sun ZFS Storage Appliance Software product information page at <http://www.oracle.com/us/products/servers-storage/storage/unified-storage/sun-storage-7000-uss-103104.html> for links to more information.

The following information is included:

- [Introduction to Network Management](#)
- [Roles Required for Network Management](#)
- [Actions for a Network and Network Domain](#)
- [Location of Network Information in the User Interface](#)
- [Fabrics](#)
- [Network Domains](#)
- [Networks](#)
- [Properties of a Network](#)
- [Network Utilization](#)
- [Network Connectivity](#)
- [Network Hardware](#)
- [Network Profiles](#)
- [Oracle Enterprise Manager Ops Center's Networks](#)
- [Related Resources for Networks](#)

Introduction to Network Management

Oracle Enterprise Manager Ops Center manages network resources, from the physical to the virtual. Fabrics provide the physical infrastructure and network domains provide the logical infrastructure. Networks are created from the resources of a network domain.

Oracle Enterprise Manager Ops Center supports Ethernet and InfiniBand network protocols. While the Ethernet interconnect is the established and common interconnect, InfiniBand is popular in high-performance computing environments because it maximizes the speed of transactions using the short, multiple connections found in clusters and data centers.

- For an Ethernet network, both tagged and untagged VLANs are supported.
- For an InfiniBand network, partitions are supported.

Note: If you use an InfiniBand switch in an Ethernet network, the ports on the switch have Ethernet names.

Roles Required for Network Management

[Table 7–1](#) lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 7–1 Network Tasks and Roles

Task	Role
Edit Network Domain	Network Admin
Edit Network Attributes	Network Admin
Edit Network Services	Network Admin
Aggregate Links	Network Admin
Create IPMP Groups	Network Admin
Attach Network to Server Pool	Cloud Admin
Create Network	Network Admin
Create Private Network	Network Admin
Create Network Domain	Network Admin
Define Network	Network Admin
Delete Network Domain	Network Admin
Delete Network	Network Admin
Define Ethernet Fabric	Network Admin
Add Fabric	Network Admin
Remove Fabric	Network Admin
Assign Network	Network Admin
Connect Guests	Network Admin
Discover and Manage the Switches	Network Admin
Configure Network for Server Deployment	Server Deployment Admin

Actions for a Network and Network Domain

After a network is discovered or created, you can perform the following actions, depending on the requirements.

- Edit Network Attributes
- Edit Network Services
- Aggregate Links
- Create IPMP Groups
- Attach Network to Server Pool
- Define Network
- Delete Network
- Assign Network
- Connect Guests

- Discover and manage the switches

Location of Network Information in the User Interface

[Table 7–2](#) shows where to find information.

Table 7–2 Location of Network Information in the BUI

Object	Location
Fabric	Expand Networks in the Assets pane. Then select Fabrics.
Physical Fabric	Expand Networks in the Assets pane. Then select Fabrics and select Network Switches.
Network	To see all networks, regardless of type, expand Networks in the Assets pane. Then select Network Domains.
Services of a Network	Network Services tab: time server, WINS, DNS, and NIS. To modify these services, edit the network services. You cannot change the network's IP address or name.
Network Domain	Expand Networks in the Assets tree. The Default Network Domain is the first item.
Physical switch	Expand Assets and expand Network Switches. To see each port, click the Connectivity tab.

Fabrics

The fabric is the physical network infrastructure, such as switches, ports, host bus adapters, that provides network resources, through a network domain, to virtual assets.

When you use Oracle Enterprise Manager Ops Center to discover a physical switch or the host of switch, all the switching fabrics that the switch supports are also discovered. One physical fabric supports many fabrics, also called data links. The physical fabric is the collection of all switch ports, links, and physical interfaces or endpoints.

Each port on an Ethernet switch can support 128 fabrics through a VLAN ID. Each partition on an InfiniBand switch can support 32000 partition keys. The members of a partition are identified by partition keys or "P-keys," which act like Ethernet's VLAN tags. For example, if a server has two P-Keys, it participates in two different partitions.

Fabrics provide network resources in a way that depends on their type. [Table 7–4](#) shows the types of fabrics.

Table 7–3 Fabrics and the Network Domain

-	What Is Managed	Capability	Limits
Fully-managed switched fabric	The switch is a managed asset.	You can create a dynamic private network for each VLAN ID or partition key.	This type of fabric can be achieved on only the Sun Ethernet 10GbE Fabric switch or the Sun Datacenter InfiniBand switch and gateway.
Host-managed fabric	The host of the switch is a managed asset. The VLAN IDs or partition keys are available.	After you enable the VLAN IDs or partition keys on the switch ports connected to the host, you can create a dynamic private network for each VLAN ID or partition key.	This type of fabric is the most common type. To create a host-managed fabric from an existing Ethernet fabric, use the Define Ethernet Fabric action.
Unmanaged fabric	The switch is discovered or declared during the discovery of another asset but is not managed.	You can create static private networks.	This type of fabric cannot support dynamic private networks because the VLAN IDs or partition keys are not available.

Network Domains

A network domain is a container for fabrics, managed networks, and private networks. The network domain handles the relationship between the physical fabrics and the virtual assets, such as virtualization hosts or server pools. The fabrics provide data links and IP subnets to the network domain, which in turn, provides networks to the virtualization hosts and server pools.

Within the network domain, networks that have been discovered or specified are available for assignment. These are called public networks because their IP address space has been specified for their exclusive use. Another type of network is private, that is, the network is created using an IP address space that the network domain allocates to it.

A fabric can contribute to more than one network domain. When a network domain has more than one fabric, you designate one of the fabrics as the anchor fabric, which is the fabric from which new networks are created.

Public networks can be members of more than one network domain because their IP addresses are specific and dedicated. Private networks exist only within a specific network domain. This means a different network domain could construct a private network with the same IP address without a conflict.

In Oracle Enterprise Manager Ops Center, networks become part of a network domain in the following ways:

- An asset that has a network is discovered.
- A user creates a network.
- A network is created when it is required. This is a dynamic network.

Oracle Enterprise Manager Ops Center operates on more than one layer of the Open Systems Interconnection model, using the network domain. [Table 7–4](#) shows what the network domain manages in the physical to logical stack.

Table 7–4 Elements of a Network Domain

Layer	Asset	What Is Managed	Capability
Layer 3 Network: IP address	For Ethernet: fabric networks For InfiniBand: non-fabric networks	IP subnet and mask IP address range VLAN or Partition Services Routing	The network provides connectivity.
Layer 2 Data links	For a tagged Ethernet: VLAN For an untagged Ethernet: portID For InfiniBand: partition	VLAN IDs Partition keys (P-key)	A virtual host uses the virtual NIC and a virtual switch in a VLAN or partition.
Layer 1 Physical: switches, ports, host bus adapters	Fabrics	Varies, by type of fabric. See Table 7–3	Varies by type of fabric. See Table 7–3

Default Network Domain

The Oracle Enterprise Manager Ops Center software always has a Default Network Domain and all networks are members of it. If you have upgraded your product software from the previous release, the existing managed networks are now in the Default Network Domain. A new network becomes a member of the default network domain. If you direct the new network to a user-defined network domain, the network is also a member of that network domain.

User-Defined Network Domains

Like the default network domain, a user-defined network domain provides network resources to a server pool or virtualization host. You create a network domain to support the use of virtualization hosts, server pools, or a virtual datacenter. For example, a virtual datacenter uses server, storage, and network resources in a dynamic way, allocating and releasing resources whenever necessary. The network domain provides the network resources to the virtual datacenter.

Note: The Create Network Domain action is disabled for Exalogic systems because these systems use the Default Network Domain for all resources.

When you create a network domain, you set a limit on the number of networks that can be created in the network domain. Increase the number of networks when accounts in a virtual datacenter are not able to create vnets.

A new user-defined network domain includes the address space specified as private by the RFC 1918 specification. These addresses cannot be routed to the Internet and provide a way for organizations to create intranets. If your organization uses a portion of this private address space, reserve these IP addresses when you create a network domain so that the network domain does not use them.

Editing Attributes of a User-Defined Network Domain

You can change the name and description of the network domain and you can change the number of dynamic networks that are in use simultaneously.

To Edit Attributes of a Network Domain

1. Expand Libraries in the Navigation pane.
2. Expand Networks in the Navigation pane.
3. Select the network domain.
4. Click **Edit Attributes** in the Actions pane.

The Details tab is displayed in the center pane. The Name Description, and Number of Networks fields are now editable.

5. Edit the name or description or increase the number of networks.
6. Click Save.

Networks

In Oracle Enterprise Manager Ops Center, networks are the discovered and managed IP subnets. Oracle Enterprise Manager Ops Center manages network resources for its virtualization hosts.

Note: These networks are part of Oracle Enterprise Manager Ops Center's virtualization services. For a description of the networks that support the product, see [Oracle Enterprise Manager Ops Center's Networks](#).


Networks are associated with server pools, which contain virtualization hosts, or standalone virtualization hosts. When you assign a network to a server pool, the network is accessible from each virtualization hosts in the pool and every guest of each virtualization host.

You can use networks to do the following:

- Manage individual virtualization hosts
- Connect virtualization hosts to the Proxy Controller
- Allow guests to communicate with each other or with the Internet
- Connect remote JMX with the public API

A network depends on the physical network interface card (PNIC) that is available to the host. You can create one network for each physical network interface card. If one host has two PNICs, it is a good practice to create two networks: a management network and a data network. Then place all virtual hosts on the data network, keeping them separate from the management network. The management network is dedicated to giving access to internal resources of the data center.

[Figure 7-1](#) shows how two virtualization hosts participate in two networks. The actual network connection is made to the PNICs in the virtualization host. Network A is connected to PNIC 1 of both hosts and Network B is connected to PNIC 2 of the hosts.

Figure 7-1 Network with Virtual Hosts

Requirements for a Network

A network requires a physical network interface or a link aggregation and the following specifications:

- IP address and netmask or CIDR format
- If you use static IP addressing, the IP address of the management interface
If you use dynamic IP addressing, the range of allowed IP addresses and the gateway address

Public Networks and Private Networks

Networks are introduced into Oracle Enterprise Manager Ops Center in the followings ways:


- By discovering the fabric that supports the networks. All the attributes are discovered but, other than the name and description, they cannot be changed.
- By defining the specific attributes of the network from its fabric. This is the result of the Define Network action. In previous versions, this action was called Manage Network.
- By creating a network from the network domain. This is the result of the Create Network action when you assign the specific VLAN IDs or partition keys and the Create Private Network action when you allow Oracle Enterprise Manager Ops Center to fulfill the requirements from the network domain. For the procedures for creating these types of networks, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm

You can use public network for most purposes. If you are creating networks for the use of virtual datacenters, create a private network to ensure that the virtual datacenter has exclusive use of the IP address space that it gets from the network domain.

Networks for Server Pools

A server pool obtains its networks from a network domain. When you create a server pool, you associate a network domain with the server pool and you specify the physical network interfaces in the fabric that the virtualization hosts in the server pool will use. All virtualization hosts in the server pool are associated with the same set of networks. When you add a virtualization host to a server pool, the virtualization host has access to all the networks defined for the pool and can be an active member of the pool. [Figure 7-2](#) is an example of network connections to two virtualization hosts in a server pool. This server pool has two virtualization hosts and two network associations.

Figure 7-2 Network Connections for a Server Pool



All the virtual hosts of a virtualization host have the same network access. This allows you to migrate a virtual host from one virtualization host to another one within the pool.

Networks for Oracle Solaris Zones

You specify the networks that support a global zone. The network attached to a global zone also supports its non-global zones. You can attach one or more networks to a global zone using the **Attach Networks** action.

Although the fabric that supports the network does not change, you can specify the way the network works differently for each type of zone. A network that is assigned as shared on a global zone can be assigned as exclusive on another global zone.

- In Shared IP mode, more than one zone uses the network interface. You define the network interface when 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.

For a complete description, see [Managing Global Zone Networks](#) in [Chapter 14](#).

Networks for Oracle VM Servers

The Oracle VM Server Control Domain's network support is specified when the Oracle VM Server software is installed to provide networking facilities for the logical domains. For each network connection, a virtual switch is created.

For a complete description of networks for Oracle VM Server for SPARC, see [Attaching Networks](#) in [Chapter 15](#).

For a complete description of networks for Oracle VM Server for x86, see [Manage Networks](#) in [Chapter 16](#).

Networks for Virtual Datacenters

Each virtual datacenter uses server, storage, and network resources in a dynamic way, allocating and releasing resources when necessary. The virtual datacenter inherits its network resources from the network domain that supports the server pool. These networks form the public external networks for the virtual datacenter. These networks can then be assigned to the accounts in the virtual datacenter. When the user of an account creates a private vNet, either a dynamic private network is created or the static private network is made available for use in that account.

The public networks assigned to a virtual datacenter are shared among the accounts in a virtual datacenter. You can attach virtual servers from different accounts to the same network. When you create a virtual server, you specify the number of networks to which it can be associated.

Although the network domain provides the private network, the virtual datacenter creates virtual networks from each IP address. The size of the private network, that is the maximum number of members that you specify, is the maximum number of IP addresses or virtual networks and therefore, the maximum number of virtual servers that can connect to the private network. Virtual networks are displayed in the vDC Management section of the Navigation pane as shown in [Figure 18–1](#).

For a complete description of networks for virtual datacenters, see [Creating vNets](#) and [Network Setup](#) in [Chapter 18](#).

Bandwidth Management

A data link is a physical NIC, an aggregated link, or a virtual NIC. When a new data link is created, the operating system sets the default bandwidth flow. You cannot remove this flow. The flow is removed only when the physical link is removed.

In Oracle Solaris 11 operating system environments, you can manage the bandwidth flow of a data link, prioritizing the network traffic on the link and setting the maximum bandwidth limit.

Managing the Bandwidth Flows for a Data Link

1. Expand Assets in the Navigation pane.

2. Select an Oracle Solaris 11 operating system.
3. Click the Networks tab in the center pane.
4. Click the Bandwidth Management subtab in the center pane.
5. To modify a flow, click the Modify icon. To create a new link, click the Add icon, then specify a name for the flow and the physical network interface.

The name of flow must meet the following requirements:

- The first character must be alphabetic.
 - All characters must be alphanumeric: a-z, A-Z, 0-9, underscore ('_'), period ('.'), or hyphen ('-').
 - Maximum number of characters is 127.
6. Set the new bandwidth properties, as described in [Properties of Bandwidth Flow](#).

Properties of Bandwidth Flow

- Priority – Set the priority of the network traffic on the link as high, medium or low.
- Bandwidth Limit – Enable the bandwidth limit to allocate guaranteed bandwidth to the specified link. Enter the maximum value for bandwidth limit in Kbps, Mbps, or Gbps.
- Set attributes for the data flow to identify its network traffic:
 - Local and Remote IP – The source and destination IP address.
 - Transport – The Internet Protocol used such as TCP, UDP, SCTP, ICMP.
 - Ports – The source and destination ports for TCP, UDP, and SCTP.
 - DS Field – The type of service field in the IP packets' header.

IP Multipathing Groups

Using IP Multipathing (IPMP), two or more physical network interface cards (NIC) form a group that use one IP address. If one NIC fails, the other NIC in the group maintains network access.

A network interface can be a physical network interface card (NIC) or, for an Oracle Solaris OS asset, it can be an IPMP group or link aggregation. You can implement both methods on the same network because they work at different layers of the network stack.

For information about how IPMP groups work, see the Oracle Solaris 11 documentation at http://docs.oracle.com/cd/E23824_01/html/821-1458/gfkcy.html and the Oracle Solaris 10 documentation at http://docs.oracle.com/cd/E26505_01/html/E27061/mpoverview.html.

Note: IPMP groups are supported only for IPv4 protocol.

IPMP provides increased reliability, availability, and network performance for systems with multiple physical interfaces because IPMP detects a physical interface failure and migrates network access to another member transparently.

Using IPMP, you can configure two or more physical interfaces into an IPMP group. If an interface in the group fails or is removed for maintenance, IPMP migrates the failed

interface's IP addresses to another member of the group. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections.

The association between an IPMP group and a network must be unique; an IPMP group can be associated with only one network and a network can be associated with only one IPMP group or individual NICs.

In an IPMP group, you define whether each interface is a failover or a standby one. The actions of each type differ if the current network interface fails:

- Network access changes from the failed interface to the failover interface in the IPMP group and uses the failover interface data address. You must provide the data address for an interface that is defined as failover.
- Network access changes from the failed interface to the standby interface in the IPMP group but does not change its data address. The data address of the failed interface migrates to the standby interface.

Link-based failure detection in an IPMP group is always enabled if your interface supports this type of failure detection. You can also set up probe-based failure detection by providing a test address for each interface in the group.

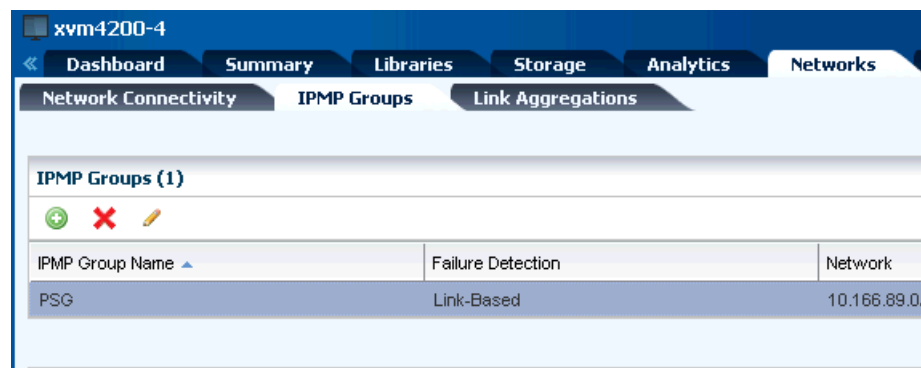
You can create a single IPMP group while provisioning an operating system. If you create IPMP groups manually, Oracle Enterprise Manager Ops Center identifies and displays the groups on the UI. See [Creating IPMP Groups](#) for information and procedures for creating IPMP groups.

Creating IPMP Groups

Note: For Oracle Solaris 11 OS, you cannot create IPMP groups. Instead, you can aggregate the links.

From the Network tabs, you have the option to create and manage the IPMP groups in the selected Oracle Solaris OS. [Figure 7-3](#) shows the options that are available to create and manage IPMP groups.

Figure 7-3 IPMP Group Option



To create an IPMP group, you must define the following parameters for the group:

- The active and the standby interfaces of the group. By default, an interface added to an IPMP group is active. You can configure as many standby interfaces as you want for the group.

- The link-based failure detection is enabled by default. You must select whether you want to enable Probe-Based failure detection and if so, you must provide the test address to track the interface status.
- You must assign the data addresses for the physical interfaces in the IPMP group. Data traffic flow use the data addresses that are hosted on the IPMP interface and flow through the active interfaces of that group.

To Create an IPMP Group

1. Select the Oracle Solaris OS in the Assets section.
2. Click the **Network** tab in the center pane.
3. Click the **IPMP Groups** subtab in the pane.
Existing IPMP groups in the OS are listed.
4. Click the **Create IPMP Group** icon to open the Create IPMP Group wizard.
5. Enter the following details for the IPMP group:
 - Provide a name for the IPMP group.
 - Select a network from the list of available network interfaces.
 - The Link-Based failure detection is always enabled by default. Select whether you want to enable Probe-Based failure detection.
 - Select the interfaces in the IPMP group.

Figure 7-4 shows the Specify IPMP Group Wizard step.

Figure 7-4 Specify IPMP Group Details

Click **Next** to specify the NIC settings.

6. When you have enabled probe-based failure detection, enter the test address for the NICs in the group.
7. Select the interfaces that are in standby mode.
You must have at least one active interface in the group. Click **Next**.

8. Enter the data address for the active interfaces of the group and select whether the interface has a failover and click **Next**.
9. Review the information and click **Finish** to create the IPMP group.

Link Aggregation or Bond

A network interface can be a physical network interface card (NIC) or, for an Oracle Solaris OS asset, it can be an IPMP group or link aggregation. You can implement both methods on the same network because they work at different layers of the network stack.

In an aggregated link, two or more NICs form a group and all members of the link aggregation provide network access at the same time. In addition to the high availability and load balancing that an IPMP group provides, an aggregated link can also provide increased throughput if the network ports are also aggregated.

When interfaces have been aggregated, they are treated as a single network interface. Oracle Enterprise Manager Ops Center includes any link aggregations in the list of available NICs as if the link aggregation were an individual interface. To assign a network with a link aggregation to an Oracle VM Server or global zone, select the link aggregation from the NIC list. You can view the link aggregation details on the Oracle VM Server's or global zone's Network tab as described in [Creating Link Aggregation](#).

Link aggregation is a standard defined in IEEE802.3ad. An aggregated link consists of several interfaces on a system configured as a single, logical unit. Link aggregation increases the speed and high availability of a connection between a server and a switch. The most common protocol used to manage link aggregation is LACP (Linked Aggregation Control Protocol). For information about how link aggregation works, see the Oracle Solaris 11 documentation at http://docs.oracle.com/cd/E23824_01/html/821-1458/fpjv1.html and the Oracle Solaris 10 documentation at http://docs.oracle.com/cd/E26505_01/html/E27061/fpjv1.html.

In Oracle Solaris 10 and by default in Oracle Solaris 11, the type of link aggregation you create is a trunk aggregation, which has these requirements:

- All the members of the aggregated link are connected to the same switch.
- The members of the aggregated link are of the same type. For example, NICs with the `e1000g` interface cannot be mixed with NICs that use the `bge` interface.
- The required driver is `GLDv3`.

Oracle Solaris 11 supports an alternative to trunk aggregation called Datalink Multipathing Aggregations (DLMP). This type of aggregation overcomes the limitations of trunk aggregation for network virtualization because DLMP aggregation works with more than one switch and provides the benefits of the link layer of the network stack to the aggregation.

In a trunk aggregation, every port is associated with every datalink in the link aggregation. In a DLMP aggregation, every port is associated with every datalink in the link aggregation and every port is associated with the primary network interface and any of its VNICs that are configured to use the link aggregation.

Note: In the current release, Oracle Enterprise Manager Ops Center can display the details of both trunk and DLMP aggregation and displays them for selection when attaching a network. However, it is not possible to create a DLMP link aggregation.

For a link aggregation created in Oracle Solaris 11 OS, the MTU size for one of the members of the aggregation must be at least 9216 bytes to allow Oracle VM Servers and logical domains to use VLAN tagged networks. To change the MTU size, see the instructions in [Maximum Transmission Unit \(MTU\)](#).

Creating Link Aggregation

To create a link aggregation for the Oracle Solaris OS, specify the following:

- Load balancing policy
- LACP mode and timer
- MAC address policy and if required, the MAC address

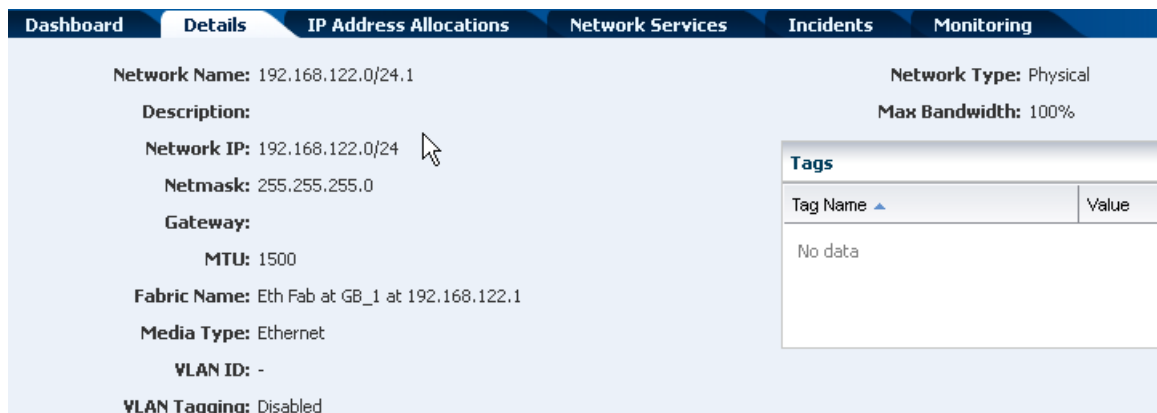
To Create a Link Aggregation

1. Click the **Network** tab.
2. Click the **Link Aggregation** subtab.
3. Click the **Create Link Aggregation** icon to open the wizard.
4. Enter the name of the link aggregation. By default, the name starts with **aggr**. Append a number to make the name unique.
5. Select the NICS for the link aggregation and then click **Next**.
6. Specify the following information:
 - Policy for load balancing
 - LACP mode and timer.
 - MAC address policy and MAC address if required.
 Click **Next** to view the summary.
7. Review the information and click **Finish** to create the link aggregation.

Properties of a Network

The characteristics of a network are displayed in its Network Details tab.

Figure 7–5 Network Details Tab



The image is a screen capture of the top section of the Network Details tab. There are 12 fields: Network Name, Description, Network IP, Netmask, Gateway, MTU, Fabric Name, Media Type, VLAN ID, VLAN Tagging, Network Type, and Max Bandwidth.

Name, Media Type, VLAN ID, VLAN Tagging, Network Type, and Max Bandwidth. It also contains a table for tags and their values but there are no entries.

The network IP address and its network type cannot be changed. Use the **Edit Network Attributes** action to change the network name and description, default gateway, MTU size, and to change the static IP routes. To change the MTU size, see the instructions in [Maximum Transmission Unit \(MTU\)](#).

Use the **Edit Managed IP Ranges** action to change the range of IP addresses that are available from the selected network. You specify the range with the starting IP address and the ending IP address. You have the option to exclude a specific IP address from the range. When you attach the network to a virtualization host, server pool, or virtual data center, this IP address is not available. You cannot exclude an IP address that is in use, which can be difficult to determine. For example, in a virtual data center, an account is assigned a range of IP addresses for its exclusive use. While the account exists, the IP addresses are in use, regardless of whether there is network activity.

IPv4 and IPv6 Protocols

Some environments have a mix of IPv4 and IPv6. Oracle Enterprise Manager Ops Center is "IPv6-aware." If an asset has an IPv6 network interface, Oracle Enterprise Manager Ops Center can read it and displays its information, but it cannot provision an IPv6 network or use IPv6 networks to discover, monitor, or provision assets.

Routing Mode

A virtual host uses the network assigned to it according to the host's routing mode. You specify a virtual host's routing mode during its initial configuration if you do not accept the default mode, Automatic Routing. Oracle Enterprise Manager Ops Center supports the following routing modes:

- Automatic Routing – This is the default routing mode. Applying the static routes depends on the following conditions:
 - If your site defined a default gateway or static route or retrieved one from the DHCP server, this route is used and dynamic routing is disabled.
 - If no default gateway or static route is available, dynamic routing is enabled.
- Dynamic Routing Off – The virtual host uses the default gateway and any static routes configured for the network. The default gateway is retrieved from the DHCP server.
- Dynamic Routing On – The virtual host uses routes provided by the dynamic routing service. The default gateway and any static routes configured for the network are ignored.

Static Route for the Network

Static routes specify the route for external access. Although you define a default gateway for a network, it might not reach a particular subnet. In this case, you must also provide a static route for the subnet.

When you create a network, you can specify the static route. To add static routes after the network has been created, use the following procedure.

To Add a Static Route for the Network

1. Click Managed Networks in the Navigation pane.
2. Select a network from the list of networks.
3. Click **Edit Network Attributes** in the Actions pane.
4. Click the Add icon in the Static Routes table. A row is added to the table.
5. Enter the values for destination IP, netmask, and gateway.
6. Click Finish.

You can delete a static route and change the order of the routes using the icons in the Static Routes table.

Address Allocation Method

When you define a new network, you specify how its IP address is assigned:

- Static IP: You enter a specific IP address.
- Automatic: An available IP address is assigned to the network.
- Do not allocate IP: No IP address is assigned at this time.

Maximum Transmission Unit (MTU)

The default size for the network's Maximum Transmission Unit (MTU) is 1500 bytes. If your network interface card is one of the following types, you can change the size of the MTU to a size between 576 and 9216 bytes. However, to assign the network to a logical domain, the minimum MTU size is 1500 bytes.

- e1000g
- ce
- nxge
- nge
- bge
- xge
- hme
- ixgbe
- hxge
- ipge
- igb

When you specify a size greater than 1500 bytes, Oracle Enterprise Manager Ops Center modifies the network interface card's MTU size. For other types of network interface cards, the MTU is changed when the card's driver firmware is updated to support the new MTU size. However, to change the MTU value for an IPMP group, you must edit the MTU value manually.

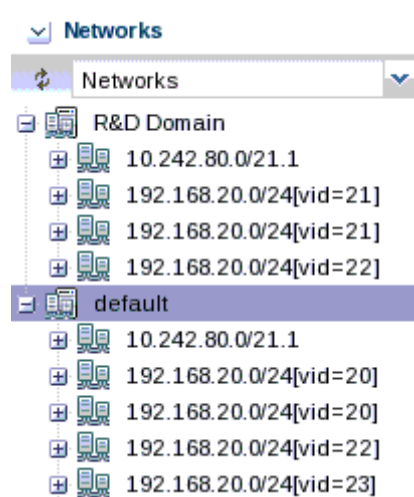
Note: When you provision an operating system, the MTU size resets to the default value. You must change the MTU again after you provision the system.

VLAN and VLAN Tags

For fabrics based on Ethernet protocol, the ability to use VLAN tags is an attribute of each network. Use the **Edit Network Attributes** action to add or change the VLAN capability. Use the **Define Network** action to specify a network as tagged, untagged, or both.

When one CIDR supports both tagged and untagged networks, you can distinguish them by the default User Friendly Name (UFN), as shown in [Figure 7-6](#). Oracle Enterprise Manager Ops Center appends the VLAN ID or tag to the UFN. For an untagged network with no VLAN ID, the string [UNTAG] is appended.

Figure 7-6 VLAN Tags



Network Utilization

Oracle Enterprise Manager Ops Center collects information every five minutes on every managed asset and displays the last hour of data on the asset. To see utilization data for a network over longer periods of time, up to six months, create a Network Utilization chart, which includes operating system, operating system for a virtual machine, virtual host, and server pool. You can also create a network utilization chart for an OS group or host group.

Network Connectivity

Connectivity is the network interface of the system. You can view information about a hardware asset's Network Interface Card (NIC) on the Connectivity tab of the asset's dashboard, including name, connection status, MAC address, and the corresponding IP address.

For switch hardware, the Connectivity tab shows information about each port.

For an Oracle Solaris OS, the Connectivity tab includes IPMP groups and aggregated links.

- The IPMP Groups subtab shows the group's name, its assigned network, and the type of failure detection, either link-based, probe-based, or both. For each IPMP group, the details include the state of the connection for each NIC, whether it is in standby mode or failover mode, and the IP address the NIC supports.

- The Link Aggregation subtab shows the aggregation's name, its MAC address, and its attributes. For each aggregated link, the subtab shows the state of the connection for each NIC, whether it is in standby mode or failover mode, and the IP address the NIC supports.

When you attach or assign networks or when you create virtual hosts, [Figure 7-7](#) shows an example of a step in the wizard where you configure the network connection.

Figure 7-7 Configure Interfaces

Specify Configuration Settings for each Network Connection						
Hostname	Domain	SR-IOV	Network	NIC	Address Allocation Method	IP Address
smt4v2-1	primary	<input type="checkbox"/>	1.1.1.0/24.1(v...	net0	Use System...	System
smt4v2-1	MyRD0	<input type="checkbox"/>	10.166.85.0/2...	net0	Do not Alloc...	Not Alloc

Network Hardware

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches and Sun Datacenter InfiniBand switches. These switches reside in the system or blade system and provide the switch fabric.

The InfiniBand Gateway switch can expose the ports of a server that resides on an InfiniBand partition to an Ethernet network. To create an Ethernet on InfiniBand (EoIB) interface on the switch, you associate the switch's external port (eport) with the InfiniBand partition where the server resides, creating a virtual NIC (vNIC). The server's ports are displayed on the Switch Connectivity tab in the center pane.

For more information about these switches, see [Switch Details](#) or see [Related Resources for Networks](#) for links to the switch documentation.

Network Interface Card (NIC)

The Network Interface Card (NIC) is the physical connection between a network switch and a network. When you create a network or attach an asset to a network, you select the NIC. You can create one network for each physical network interface card. To see the NICs for a server, select the server and then click the Connectivity tab. The Server Processor Connectivity table lists all of the NICs.

Network Switches

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches and Sun Datacenter InfiniBand switches. These switches reside in the system or blade system and provide the switch fabric.

For the Ethernet switches, both tagged and untagged VLANs are supported.

If you use an InfiniBand switch in an Ethernet network, the ports on the switch have Ethernet names.

For more information about these switches, see the product documentation:

- For the Sun Ethernet 10GbE Fabric switch, see <http://docs.oracle.com/cd/E19934-01/index.html>

- For Sun Datacenter InfiniBand switch, see <http://docs.oracle.com/cd/E19654-01/index.html>

Network Profiles

Oracle Enterprise Manager Ops Center provides default profiles for the following operations:

- Monitor Network hardware – Reports Cisco switch's connection to assets on the Switch Connectivity tab.
- Discover a switch – Use a discovery profile with Cisco iOS credentials.

Oracle Enterprise Manager Ops Center's Networks

This section describes the requirements for the networks that Oracle Enterprise Manager Ops Center uses. This section does not discuss the networks that support virtual hosts and server pools.

You can implement Oracle Enterprise Manager Ops Center's network connections using any combination of VLANs and switches. However, each network, whether management, provisioning, or data, must be assigned to separate VLANs.

Network Switch Configuration

Use these guidelines to configure a network switch for a system running the Oracle Enterprise Manager Ops Center software.

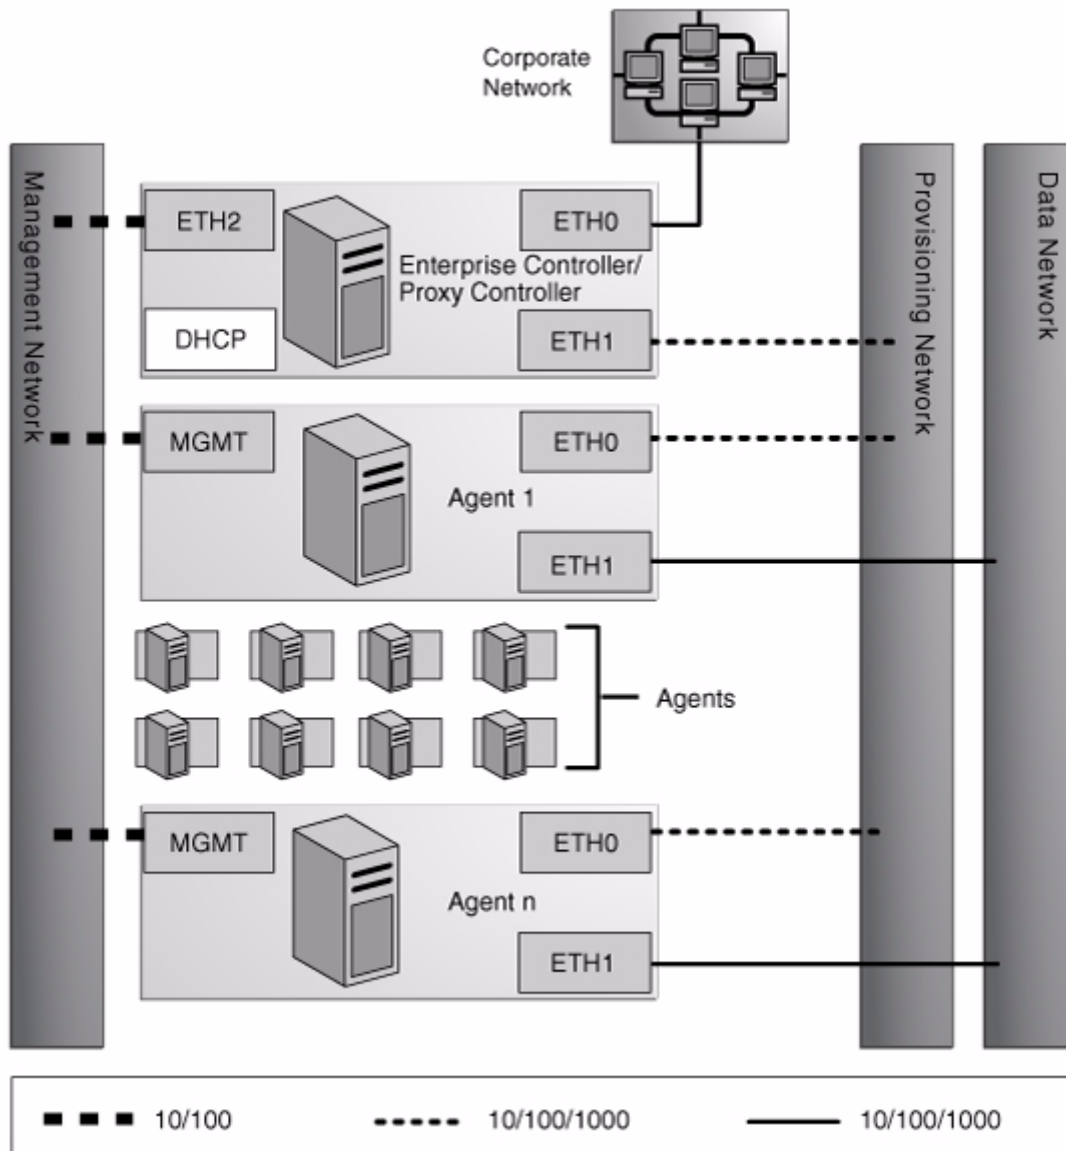
- Use an Virtual LAN (VLAN)-capable switch.
- Discover and manage the switch.
- Create a separate VLAN for Oracle Enterprise Manager Ops Center management and provisioning networks.

For Ethernet connectivity:

- The management network must be a 10/100 connection.
- The provisioning and data networks must be at least a 10/100/1000 (1 GB) connection.

Separate Management, Provisioning, and Data Networks

- Separate networks, as shown in [Figure 7-8](#), provide the highest security and the lowest number of points of failure.
- Additional NICs are needed to support this configuration.

Figure 7–8 Separate Management, Provisioning, and Data Networks

A configuration with separate management, provisioning, and data networks has these requirements:

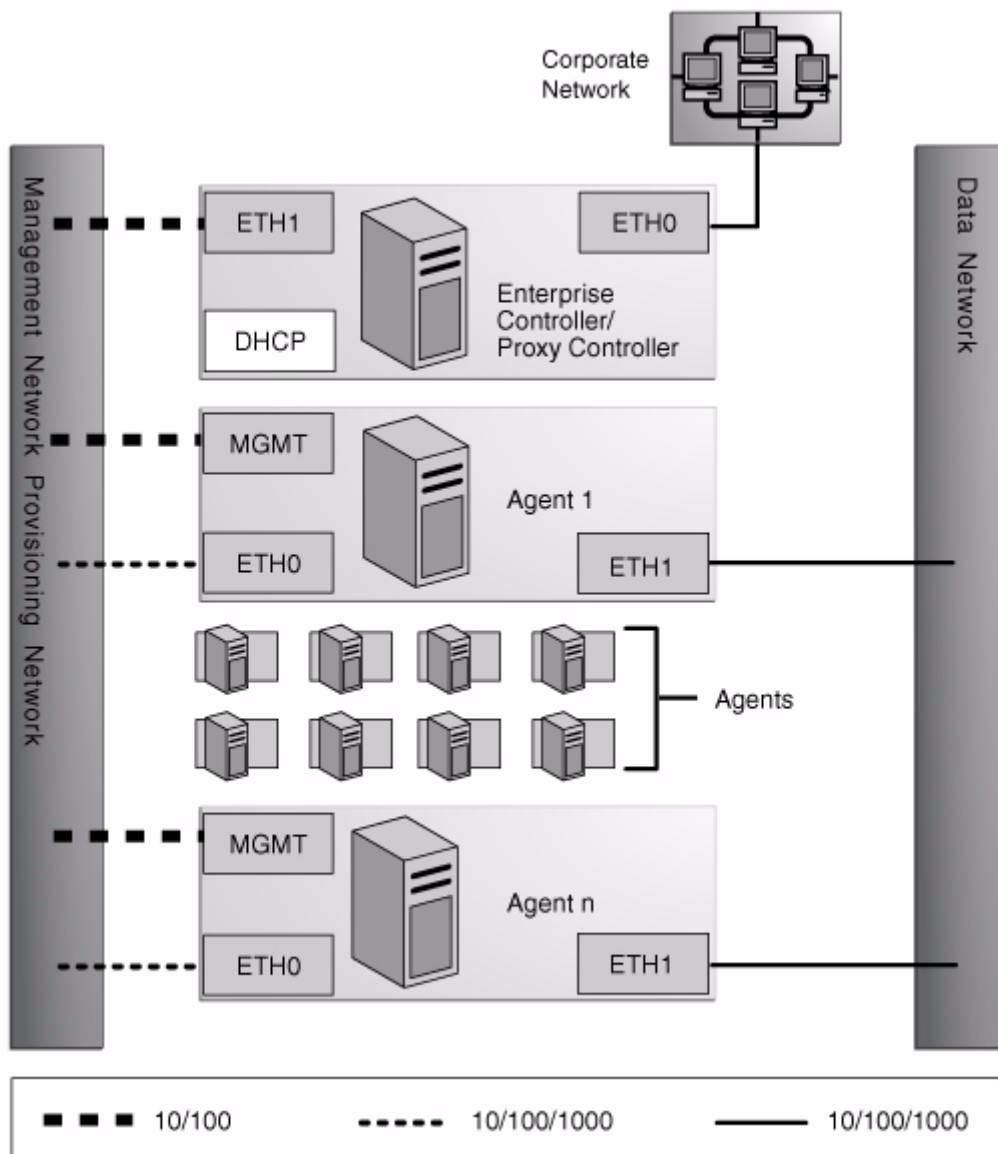
- Enterprise Controller/Proxy Controller
 - ETH0 connects the Enterprise Controller/Proxy Controller to the corporate network for external access. Configure the ETH0 IP address, netmask, and gateway to meet corporate connectivity requirements.
 - ETH1 connects the Enterprise Controller/Proxy Controller to the provisioning network and must be on the same network as the ETH0 connections of the Agent Controllers. Only the Enterprise Controller/Proxy Controller and the Agent Controllers must reside on the provisioning network. ETH1 must be a 1 Gb NIC interface.
 - ETH2 connects the Enterprise Controller/Proxy Controller to the management network and must be on the same network as the management port connections of the Agent Controllers. Configure the ETH2 IP address,

netmask, and gateway to enable connectivity to the Agent Controllers' management port IP addresses. ETH2 must be a 100 Mb NIC interface.

- The DHCP service allocates IP addresses to the Agent Controllers for provisioning operating systems.
- Agent Controllers
 - Each Agent Controller's management port connects the Agent Controller to the management network and must be on the same network as the ETH2 connection of the Enterprise Controller/Proxy Controller. The management port must be a 100 Mb connection.
 - ETH0 connects the Agent Controller to the provisioning network and must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. ETH0 must be a 1 GB connection.
 - ETH1 connects the Agent Controller to the data network through the switch to provide corporate network access to the Agent Controller. ETH1 must be a 1 GB connection.

Combined Management and Provisioning Network and a Separate Data Network

- Combining the management and provisioning networks, as shown in [Figure 7-9](#), reduces system and network security.
- No additional NIC is needed on the Enterprise Controller or Proxy Controller.

Figure 7–9 Combined Management and Provisioning Network and a Separate Data Network

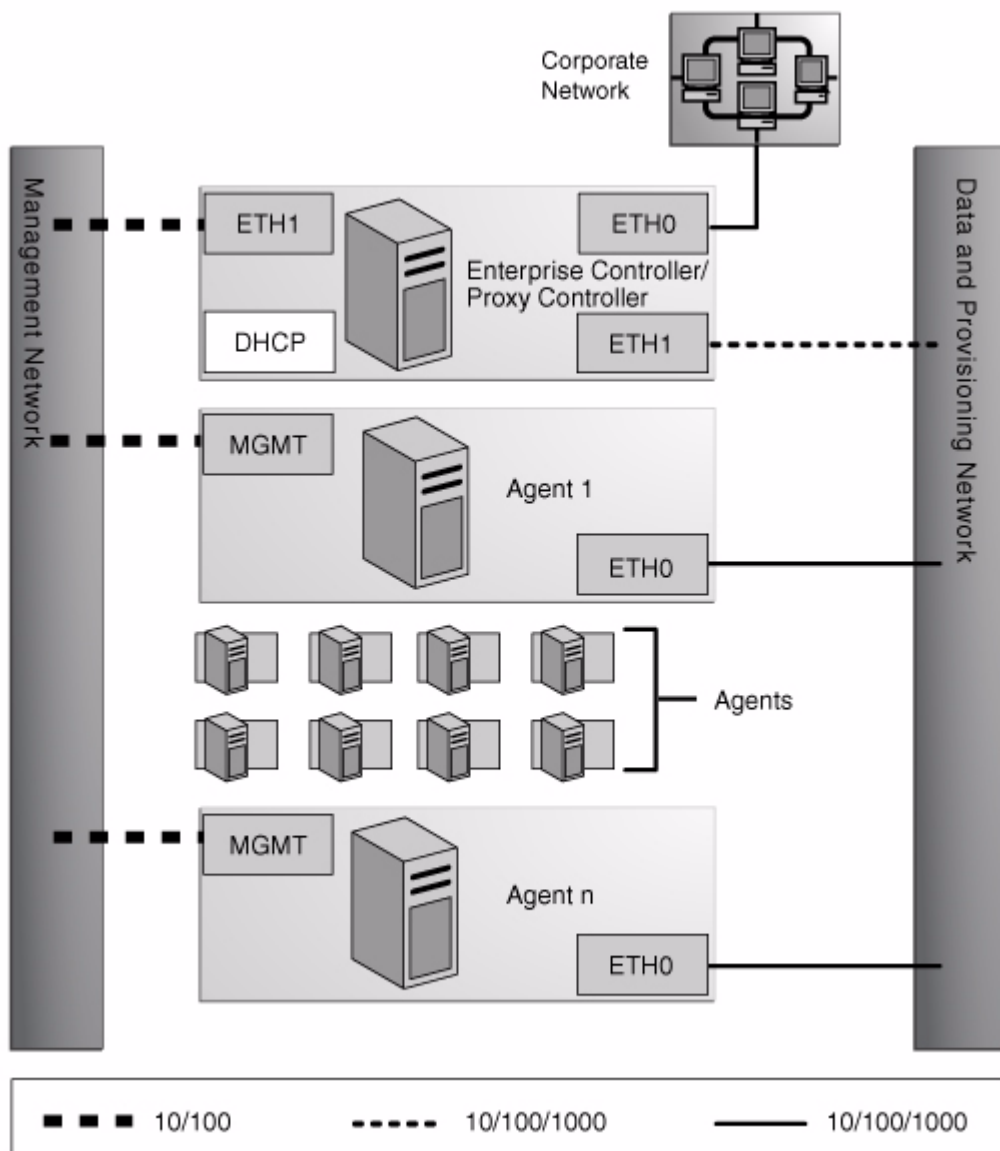
- Enterprise Controller/Proxy Controller
 - ETH0 connects the Enterprise Controller/Proxy Controller to the corporate network to provide external access. Configure the ETH0 IP address, netmask, and gateway to meet corporate connectivity requirements.
 - ETH1 connects the Enterprise Controller/Proxy Controller to the management and provisioning network and must be on the same network as the MGMT and ETH0 connections of the Agent Controllers. Only the Enterprise Controller/Proxy Controller and the Agent Controllers must reside on the management and provisioning network. The ETH1 IP address, netmask, and gateway must be configured to enable connectivity to the Agent Controller's management port IP addresses. ETH1 must be a 1 Gb NIC interface.
 - The DHCP service allocates IP addresses to the Agent Controllers for provisioning operating systems.
- Agent Controllers

- Each Agent Controller's management port connects the Agent Controller to the management and provisioning network and must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. The management port must be a 100 Mb connection.
- ETH0 connects the Agent Controller to the management and provisioning network and must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. ETH0 must be a 1 GB connection.
- ETH1 connects the Agent Controller to the data network through the switch to provide corporate network access to the Agent Controller. ETH1 must be a 1-GB connection.

Combined Provisioning and Data Network and a Separate Management Network

Figure 7–10 shows a configuration with a network for both provisioning and data.

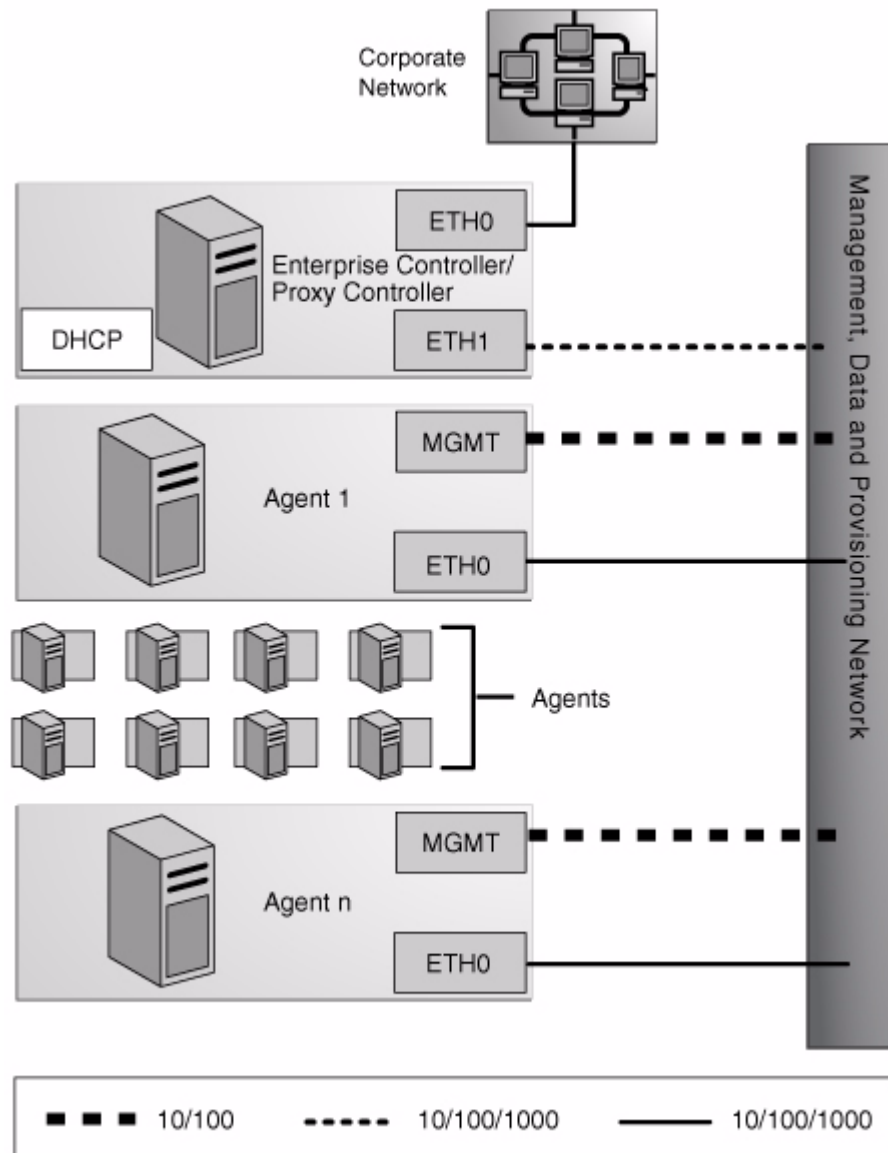
Figure 7–10 Combined Provisioning and Data Network and a Separate Management Network



- Enterprise Controller/Proxy Controller
 - ETH0 connects the Enterprise Controller/Proxy Controller to the corporate network to provide external access. Configure the ETH0 IP address, netmask, and gateway to meet corporate connectivity requirements.
 - ETH1 connects the Enterprise Controller/Proxy Controller to the provisioning and data network and must be on the same network as the ETH0 connections of the Agent Controllers. Only the Enterprise Controller/Proxy Controller and the Agent Controllers must reside on the data and provisioning network. ETH1 must be a 1 Gb NIC interface.
 - ETH2 connects the Enterprise Controller/Proxy Controller to the management network and must be on the same network as the management port connections of the Agent Controllers. Configure the ETH2 IP address, netmask, and gateway to enable connectivity to the Agent Controller's management port IP addresses. ETH2 must be a 100 Mb NIC interface.
 - The DHCP service allocates IP addresses to the Agent Controllers for provisioning operating systems.
- Agent Controllers
 - The management port connects the Agent Controller to the management network and must be on the same network as the ETH2 connection of the Enterprise Controller/Proxy Controller. The management port must be a 100 Mb connection.
 - ETH0 connects the Agent Controller to the data and provisioning network to provide corporate network access to the Agent Controller. ETH0 connection must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. ETH0 must be a 1 GB connection.

Combined Provisioning, Data, and Management Network

- [Figure 7-11](#) shows the least secure system and network.
- No additional NIC is needed for the Enterprise Controller/Proxy Controller.

Figure 7-11 Combined Provisioning, Data, and Management Network

- Enterprise Controller/Proxy Controller
 - ETH0 connects the Enterprise Controller/Proxy Controller to the corporate network to provide external access. Configure the ETH0 IP address, netmask, and gateway to meet corporate connectivity requirements.
 - ETH1 connects the Enterprise Controller/Proxy Controller to the combined management, provisioning, and data network and must be on the same network as the MGMT and ETH0 connections of the Agent Controllers. Only the Enterprise Controller/Proxy Controller and the Agent Controllers must reside on the combined network. ETH1 must be a 1 GB NIC interface.
 - The DHCP service allocates IP addresses to the Agent Controllers for provisioning operating systems.
- Agent Controllers

- Each Agent Controller's management port connects the Agent Controller to the management, provisioning, and data network and must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. The management port must be a 100 MB connection.
- ETH0 connects the Agent Controller to the management, provisioning, and data network, and must be on the same network as the ETH1 connection of the Enterprise Controller/Proxy Controller. ETH0 also connects the Agent Controller to the data network through the switch to provide external corporate network access to the Agent Controller. ETH0 must be a 1 GB connection.

Firewall Rules

The Enterprise Controller must reach some external sites. If you have explicit firewall rules enabled to allow access to these services from your Enterprise Controller, you must update these rules to allow access to the addresses in [Table 7–5](#).

Table 7–5 IP Address and Port Requirements

Site	IP Address	Port	Purpose
https://java.net/projects/oc-doctor/downloads	192.9.164.103	Port 443	Updates to OCDoctor utility
https://java.net/projects/oc-cluster-profiles	192.9.164.103	Port 443	Access to Oracle Solaris Cluster profiles and scripts.
login.oracle.com	141.146.8.119	Port 443	Logging into Oracle sites
updates.oracle.com	141.146.44.51	Port 443	Access to Oracle Knowledge Base for OS updates
inv-cs.oracle.com	192.18.110.10	Port 443	Product registration
hs-ws1.oracle.com	192.18.110.11	Port 443	Product registration
support.oracle.com	141.146.54.16	Port 443	My Oracle Support
www.oracle.com	96.17.111.33 96.17.111.49	Port 80	-
aru-akam.oracle.com	na	Port 80	Provides local IP addresses to optimize download speed. Use <code>nslookup</code> to resolve the IP address, add the address to the <code>/etc/hosts</code> file, and open the firewall for the address.
a248.e.akamai.net	na	Port 443	Provides local IP addresses to optimize download speed. Use <code>nslookup</code> to resolve the IP address, add the address to the <code>/etc/hosts</code> file, and open the firewall for the address.

Network Port Requirements and Protocols

The Enterprise Controller's default port is 443. If port 443 is in use, the Enterprise Controller uses Port 11165. [Table 7–6](#) describes all the required ports and their protocols.

Table 7–6 Required Ports and Protocols

Communication Direction	Protocol and Port	Purpose
Enterprise Controller	Port 443, then Port 11165 Port 8005	Enterprise Controller in Disconnected mode
Enterprise Controller	Port 443, then Port 11165	Enterprise Controller in Connected mode
Browser to Enterprise Controller	HTTP, TCP: Port 80	Redirects to port 9443
Browser to Enterprise Controller	HTTPS, TCP: Port 9443	Web interface
Enterprise Controller to Local Database	Port 11176	Oracle Listener port
Enterprise Controller to Proxy Controller	SSH, TCP: Port 22 ICMP ping: Type 8 Code 0 (echo request)	Enterprise Controller installs or upgrades a Proxy Controller through the UI.
Proxy Controllers to Enterprise Controller	HTTPS, TCP: Port 443	Proxy Controller pushes data about assets to Enterprise Controller. Proxy Controller pulls data for jobs, updates, Agent Controllers, and OS images from the Enterprise Controller.
Proxy Controllers to Enterprise Controller	HTTP: Port 8004	WAN Boot traffic
Proxy Controllers to Enterprise Controller	ICMP ping: Type 0 Code 0 (echo reply)	During upgrades, Proxy Controllers use ICMP ping.
Remote Proxy Controller to Enterprise Control through an SSH Tunnel	SSH, Port 21161	When a Proxy Controller is deployed on a network outside of the firewall, the SSH Tunnel and Port 21161 change the direction of communication so that the remote Proxy Controller does not initiate communication with the Enterprise Controller.
Proxy Controller to ALOM or XCSF Service Processors	SSH, TCP: Port 22 or Telnet, TCP: Port 23 SNMP, UDP: Port 161 TCP: Port 6481 (for discovery by service tags)	Proxy Controller discovers, manages, and monitors the service processor.
Proxy Controller to ILOM Service Processors	SSH, TCP: Port 22 SNMP, UDP: Port 161 IPMI, TCP, UDP: Port 623 TCP: Port 6481 (for discovery by service tags)	Proxy Controller discovers, manages, and monitors the service processor.
Proxy Controller to ALOM or XCSF Service Processor	FTP, TCP: Port 21	Proxy Controller provisions firmware on an ALOM service processor. Port 21 transfers the firmware image. A transient random port is opened for the duration of the operation.
Proxy Controller to ILOM Service Processor	TFTP, UDP: Port 69	Proxy Controller provisions firmware on an ILOM service processor. Port 69 transfers the firmware image. A transient random port is opened for the duration of the operation.

Table 7–6 (Cont.) Required Ports and Protocols

Communication Direction	Protocol and Port	Purpose
Service Processor to Proxy Controller	SNMP, UDP: Port 162 ICMP ping: Type 0 (echo reply)	For monitoring hardware, the service processor sends SNMP traps to the Proxy Controller. For a failed connection, Proxy Controller receives ICMP ping Type 3 (destination unreachable).
Proxy Controller to OS Host	SSH, TCP: Port 22 or Telnet, TCP: Port 23 TCP: Port 6481 (for discovery and monitoring by service tags) ICMP, Type 0 Code 0 (echo reply)	Proxy Controller discovers, manages, and monitors an asset.
Proxy Controller to OS Host	DHCP, UDP: Port 67	Proxy Controller provisions an OS.
OS Host to Proxy Controller	HTTP, TCP: Port 8004 Oracle Solaris 11 Automated Installer Web Server: Port 5555 to accept requests from the OS Host during provisioning <ul style="list-style-type: none"> For provisioning by DHCP, the opened port is on the Proxy Controller. For provisioning by WAN Boot, the opened port is on the Enterprise Controller or a Proxy Controller. 	OS Host reports status of OS updates and status of Agent Controller installation. OS Host downloads Agent Controller archive file.
OS Host to Proxy Controller	DHCP, UDP: Port 68 TFTP, UDP: Port 69 TCP+UDP: Port 37 HTTP, TCP: Port 8004	OS Host responds to Proxy Controller inquiries during bare-metal OS provisioning
Agent Controller to Proxy Controller	HTTPS, TCP: Port 21165	Agent Controllers push asset data to Proxy Controller. Agent Controllers pull data for jobs.
Agent Controller to Proxy Controller	HTTPS, TCP: Port 8002	Agent Controllers pull updates from Proxy Controller.
Agent Controller on Oracle Solaris OS or on Oracle hardware to co-located Proxy Controller	SNMP: Port 1162, or a port in the range of 1100 through 1200	For monitoring assets, the Agent Controller sends trap notifications and fault management alerts (FMA) to the Proxy Controller as local traffic. Because the Proxy Controller is using Port 162, a co-located Agent Controller uses Port 1162, if it is available, or a port in the range of Ports 1100 through 1200.
Java client to public APIs	TLS: Port 11172	JMX access from clients

Table 7–6 (Cont.) Required Ports and Protocols

Communication Direction	Protocol and Port	Purpose
WMI client on Proxy Controller to Agent Controller	Port 11162	WMI client resides on the Proxy Controller and communicates with the WMI server on the Agent Controller. The Proxy Controller uses the DCOM protocol to monitor a Windows system. The Proxy Controller opens a TCP connection to the Windows DCOM registry port, TCP 135, which provides a lookup service to the WMI scripting DCOM object. The Proxy Controller connects to the DCOM object. The port number for this connection is allocated by the Windows system.
Proxy Controller to NFS server	Use an NFS server that is on the same side of the firewall as the Proxy Controller. Refer to your OS documentation to set up the NFS server.	Proxy Controller pulls provisioning images from NAS Library
Global Zones or Oracle VM Servers to NFS server	Use an NFS server that is on the same side of the firewall as the Proxy Controller. Refer to your OS documentation to set up the NFS server.	Global Zones and Oracle VM Servers push their metadata and virtual host images to NAS Library
OCDoctor to java.net	HTTPS, TCP: Port 80	Acquires product updates.

Related Resources for Networks

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

- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm
- For the Sun Ethernet 10GbE Fabric switch, see <http://docs.oracle.com/cd/E19934-01/index.html>
- For Sun Datacenter InfiniBand switch and gateway, see <http://docs.oracle.com/cd/E19654-01/index.html>

Plans and Profiles

This chapter provides an overview of the concepts of profiles, policies, and plans. Detailed information is covered in the corresponding feature chapters.

The following information is included:

- [Introduction to Plans and Profiles](#)
- [Roles for Plans and Profiles](#)
- [Actions for Deployment Plans](#)
- [Location of Profiles and Plans in the User Interface](#)
- [Operational Profiles and Plans](#)
- [Profiles and Policies](#)
- [Deployment Plans](#)
- [Related Resources for Plan Management](#)

Introduction to Plans and Profiles

Oracle Enterprise Manager Ops Center uses a combination of profiles and plans to reduce complexity and increase consistency when you perform standard management and operational activities. Templates are available in the Plan Management section of the UI. You use templates to create customized profiles and plans for your environment.

The software uses profiles and plans to create a reusable set of procedures to perform tasks, such as provisioning and updating. You create, manage, and access all profiles and plans in the Plan Management section of the user interface.

A profile defines how a task is performed and enables you to define what is allowed, and not allowed, to be installed on a system. A plan provides the steps to complete an operation, defines the targets upon which the tasks are performed, and enables you to deploy the profile to the targets. Together, plans and the associated profiles enable you to define operations. You create your own profiles and plans from a set of templates.

Plan Management contains three basic components:

- **Operational Profiles and Plans:** Perform one or more operational activities, such as scripts and utilities to fix common problems, monitoring rule configurations and thresholds, and a knowledge base that you can create based on the incidents that occur in your environment.
- **Profiles and Policies:** Profiles define the configuration of components for a specific type of system and task, such as the naming schema and configuration options to

use when creating a new Oracle Solaris 10 zone. Update policies define the level of interaction you want when applying patches and packages.

- **Deployment Plans:** Perform standard management activities, such as provisioning and creating Oracle Solaris Zones.

To see all profiles and plans, expand Plan Management in the Navigation pane. Select the profile or plan to see the available actions. For each plan, you can define the course of action to take when a step cannot be completed on a target. You can choose to stop the job at the first failure, or attempt to complete as much as possible.

Deployment plans are all based on a set of pre-defined templates. You can use the available templates to create customized profiles and plans for your data center. Some templates are designed for simple tasks and other, more complex plans, are designed to perform a series of functions. You can build complex plans by combining existing plans, for example, adding an operational profile and plan to the end of a deployment plan.

Roles for Plans and Profiles

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

Table 8–1 Plan Management Roles and Permissions

Task	Role
View a profile or plan	Read
Create a profile or plan	Profile Plan Administrator
Edit a profile or plan	Profile Plan Administrator
Copy a profile or plan	Profile Plan Administrator
Delete a profile or plan	Profile Plan Administrator
Create a profile or plan	Profile Plan Administrator

Actions for Deployment Plans

After you create a profile and plan, you can perform the following actions:

- [Editing a Profile](#)
- [Editing a Deployment Plan](#)
- [Copying a Deployment Plan](#)
- [Deploying a Plan](#)
- [Deleting a Deployment Plan](#)

Editing a Profile

Editing a profile enables you to change the profile configuration. Changing values other than the name or description will create a new version. When a profile is referenced by one or more deployment plans, the default action is to update the referring plans to the new version.

Perform the following steps to edit a profile:

1. Expand **Plan Management** in the Navigation pane.
2. Expand the **Profiles and Policies** navigation tree, then select a profile.
3. Click **Edit Profile**.
4. Complete the changes, then click **Finish**.

Editing a Deployment Plan

You can edit the deployment plan details, alter the plan configuration by skipping steps in the plan, change the profile or plan bound to each step, or save the plan under a different name to create a new plan.

Note: When you edit a deployment plan that is referred to by another plan, for example, in a nested plan, the referring plan is not automatically updated to refer to the edited plan's version. You must manually modify the referring plan if you want it to use the modified version.

Perform the following steps to edit a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select the **Edit Deployment Plan** option:
 - Method 1: Select the deployment type from the tree and select a plan from the list to enable the icon. Click the **Edit Deployment Plan** icon.
 - Method 2: Expand the selected deployment type and select a plan from the list. The plan details appear. Select **Edit Deployment Plan** from the Actions pane.
3. Edit the following details of the plan:
 - Plan Name: Edit the name to create a new plan. You create a new plan when you modify the plan name.
 - Description: Provide a description of the plan.
 - Failure Policy: Select whether you want the plan execution to stop at failure or complete as much as possible.
4. Configure a step of the plan by setting or changing the associated profile, or by creating a new profile.
5. (Optional) Edit the plan by replicating the steps and associate targets depending on the type of plan selected.
6. Click **Save** to save any changes made to the plan. When you have changed the name, a new plan is saved with the version v1.

Copying a Deployment Plan

You can copy an existing deployment plan, rename it, and create a new plan.

Perform the following steps to copy a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select **Copy Deployment Plan**:

- Method 1: Select the deployment type from the tree and select a plan from the list to enable the Copy Deployment Plan icon. Click the **Copy Deployment Plan** icon.
 - Method 2: Expand the selected deployment type and select a plan from the list. Select **Copy Deployment Plan** from the Actions pane.
3. Edit the following details of the plan:
 - Description: Provide a description of the plan.
 - Plan Name: By default, the plan name is Copy of *<plan name>*. For example, Copy of Firmware Update. You can modify the name.
 - Failure Policy: Select whether the plan execution is to stop at failure or complete as much as possible.
 4. Configure a step of the plan by setting or changing the associated profile or by creating a new profile.
 5. Edit the plan by replicating the steps and associate targets, depending on the type of plan selected.
 6. Click **Save the new plan**. A new plan is saved with the version v1.

Deploying a Plan

When you deploy a plan, you select the target assets against which the plan is executed. After you select the targets, you have the option to temporarily override the profile configuration for specific assets before you deploy the plan.

Many actions use deployment plans. See the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm for end-to-end examples.

Deleting a Deployment Plan

You can delete a deployment plan or only a version of the plan. When the selected deployment plan is not referenced by any other plans, you can confirm deleting the plan or its version. When the plan is used in other plans, the Delete Deployment Plan option is not enabled.

Perform the following steps to delete a plan:

1. Select **Plan Management** from the Navigation pane.
2. Use one of the following methods to select Delete Deployment Plan option:
 - Method 1: Select the deployment type from the tree. The plans of that type are listed in the center pane. Select a plan from the list. The Delete Deployment Plan and Delete Version icon is enabled. Click **Delete Deployment Plan** or **Delete Version**.
 - Method 2: Expand the selected deployment type and select a plan from the list. Select **Delete Deployment Plan** or **Delete Version** from the Actions pane.
3. Click **Delete** to confirm the delete action.

Location of Profiles and Plans in the User Interface

Table 8–2 *Location of Profile and Plan Information in the UI*

Object	Location
Operational Plans	Expand the Plan Management section of the Navigation pane, then scroll down to Operational Plans.
Operational Profiles	Expand the Plan Management section of the Navigation pane, then scroll down to Operational Plans. Operational Profiles are in the Operational Plans section. Click Operational Profiles for a list of available profiles.
Deployment Plans	Expand the Plan Management section of the Navigation pane, then scroll down to Deployment Plans. Click Deployment Plans for a list of existing plans. To locate a specific plan, expand the folder for the plan you want, such as Configure RAID.
Profiles	Expand the Plan Management section of the Navigation pane, then scroll down to Profiles and Policies. Click Profiles and Policies for a list of existing profiles. Expand the folder for the type of profile you want, such as RAID Controller.

Overview of Version Control

Profiles and plans versions are numbered sequentially. When you create a profile or plan, the version number is one. When you edit a profile or plan, you create a version that is referenced by a new number.

When you create a version, you have the option to automatically change the related plan during the update. When you do not choose that option, the operational and deployment plans that are using the profile are not updated with the new version. For example, you have a firmware profile that is using firmware image A, the profile version is one. You want to update the profile to use the latest firmware image, B. When you update the profile to change the image to B, version two of the profile is created. When you have five plans that are using version one of the firmware profile, all of those plans continue to use version one and image A. When you want the plans to use image B, you must manually update the plans to reference version two of the profile.

The same is true for plans. When you edit a deployment plan that is referred to by another plan in a complex plan, the referring plan is not automatically updated to refer to the edited plan's version unless you choose the option when you edit the plan.

You can view version details by highlighting the plan, then clicking the View Version Details icon. Use the arrows to view the different versions. When you want a plan to use the new version of the profile, edit the plan and associate it with the correct version of the profile. You can edit the following fields in an Operational Plan: Description, Failure Policy, and Associated Profile. You cannot edit the name of a profile or plan. When you want a different name, you must copy or create a profile and plan.

You can delete a version of a profile or plan. When you delete a version and more than one version exists, the previous version becomes the default. When only one version of the plan exists, both the profile and plan are deleted. When the plan is used in other plans, then the Delete Deployment Plan option is not enabled.

Operational Profiles and Plans

Operational profiles are designed to assist you in the day-to-day operation of your data center.

An Operational Profile performs one or more tasks needed to operate your environment. The profile contains a single shell script and can include asset attributes as environmental variables. For each profile, you can choose one of the following types of shell scripts:

- **EC Shell:** The script runs only on the Enterprise Controller and is executed with the logged-in user's credentials.
- **Remote Shell:** The script can run on any managed system that contains a remote agent and is executed with root permissions.

You have the option of adding two types of variables to operational profiles:

- **System defined variables:** Use these variables to define specific information in the script. The following are software-specific variables:
 - \$OC_TARGET_NAME
 - \$OC_JOB_ID
 - \$OC_SASL_FILE
 - \$OC_UFN
 - \$OC_TARGET_TYPE
- **Additional Variables:** Use these variables to add information to the script and enable users to change that information when they execute the operational plan. Additional variables are user-defined.

A list of supported asset attributes is available in [Appendix A, "Asset Attributes"](#).

You must have an Operational Plan to execute an Operational Profile. When you create an Operational Profile, the default action also creates an Operational Plan. You use the plan to execute the profile on a managed resource or group of resources, such as performing state changing actions.

You can create a simple plan, such as disabling print capabilities, and deploy the plan across a group of resources in your data center. A more sophisticated example is to create several Operational Plans and add them as steps in a complex type Deployment Plan. For example, you can create an Operational Plan to shut down all logical domains and another Operational Plan to shut down an Oracle VM Server for SPARC. You can then add these plans as steps in a complex Deployment Plan.

Actions for Operational Profiles and Plans

You can perform the following actions, depending on the role:

- [Creating an Operational Profile and Plan](#)
- [Editing an Operational Profile or Plan](#)
- [Copying an Operational Profile](#)
- [Deleting an Operational Profile or Plan](#)

Creating an Operational Profile and Plan

An Operational Plan defines how an Operational Profile is deployed, and against which targets. By default, creating a profile also creates an Operational Plan.

You can save a shell script on the Enterprise Controller and download it into the plan, or you can enter the script in a field when you create the plan. Both types of shell scripts are executed by the user. They differ in the location, either on the Enterprise

Controller or on the remote Agent, and the user credentials needed to execute the profile.

An operational plan defines the targets and failure policy for an operational profile. The profile defines one or more operations that are to be performed on a managed resource or group of resources. For example, deploying thresholds onto a managed resource, or performing state changing actions such as shutting down all logical domains and then shutting down an Oracle VM Server for SPARC.

The profile uses a shell script to define the operations. When you do not have an operational profile, you can create a profile when you create the plan. When you create the profile, either download a shell script that is on the Enterprise Controller (EC Shell script), or enter the shell script (Remote Shell script) in the profile. The EC Shell and Remote Shell are both shell scripts that are executed by the user. They differ in the location (on the Enterprise Controller or on the remote Agent) and manner (user credentials) of the execution. The EC Shell is executed with the credentials of the user that is logged in. The Remote Shell can run on any managed system that contains a remote agent and is executed with root permissions.

Perform the following steps to create an operational profile and plan:

1. Expand **Plan Management** in the Navigation pane, then click **Operational Profiles**.
2. Click **Create Profile**.
3. Name the new profile and add a description. Select a subtype from the list to identify the type of target for this profile. Click **Next**.
4. Define the script, then click **Next**.
 - a. Define the type, either **EC Shell** or **Remote Shell**.
 - b. Define the time out parameters.
 - c. Enter the script in the Script field, or click **Load** if you are loading an EC Shell script that is saved on the Enterprise Controller. Click **View System Variables** to see the available software-specific variables.
5. (Optional) Add user-defined variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Editing an Operational Profile or Plan

When you edit a profile or plan, you create a version. When you create a version, you have the option to automatically change the related plans during the update. See [Overview of Version Control](#) for how versions are maintained.

You can view version details by highlighting the plan, then clicking the View Version Details icon. Use the arrows to view the different versions. When you want a plan to use the new version of the profile, edit the plan and associate it with the correct version of the profile. You can edit the following fields in an Operational Plan: Description, Failure Policy, and Associated Profile. You cannot edit the name of an Operational profile or plan. When you want a different name, you must copy or create a Operational Profile and Plan.

Copying an Operational Profile

You can copy an existing operational profile, rename it, and create a new profile and plan.

1. Expand **Plan Management** in the Navigation pane, then click **Operational Profiles**.
2. Select the profile in the center pane that you want to copy, then click the **Copy Profile** icon. Or you can double-click the profile in the center pane to display actions in the Actions pane.
3. Rename the new profile you are creating and revise the description, then click **Next**.
4. (Optional) Edit the script. Click **Next**.
5. (Optional) Edit the variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Copying an Operational Plan

You can copy an existing operational plan, rename it, and create a new profile and plan.

1. Expand **Plan Management** in the Navigation pane, then click **Operational Plans**.
2. Select the plan in the center pane that you want to copy, then click the **Copy Plan** icon. Alternatively, double-click the plan in the center pane to display actions in the Actions pane.
3. Rename the new plan you are creating and revise the description. You can choose to change the Failure Policy for the new plan and change the Operational Plan steps. You can change the associated profile or plan for each step. To change the profile or plan, click the associated profile or plan to select from a list of available options. To add additional steps, click the **Replicate Step** icon, then select the profile or plan from the list to associate with that step. Click **Save**.
4. (Optional) Edit the script. Click **Next**.
5. (Optional) Edit the variables in the Specify Additional Variables page. Click **Next**.
6. Review, then click **Finish**.

Deleting an Operational Profile or Plan

You can delete a version of a profile, or you can delete all versions of the profile and the associated plan.

You cannot delete a version of an operational plan version unless you created the version. Deleting a version of a plan might impact the Incident Knowledge Base or deployment plans that reference the version. Before deleting a version, verify that the version is not being used.

Note: When you delete a version and more than one version exists, the previous version becomes the default. When only one version of the plan exists, the operational profile and plan are deleted.

Profiles and Policies

Profiles and policies define how a job is performed and the level of interaction.

About Policies

Policies are lists of actions that are explicitly approved or denied. Policies are specific to OS update and define the amount of user interaction you want when applying OS patches and packages. For example, you might want to automate applying fixes without user interaction, but you might want to pause a job and require user approval before performing a downgrade or uninstall. You can also deny certain actions, such as installing uncertified patches or packages.

Policy settings are hierarchical. When you have not defined a policy for a component, the policy for that component's parent applies. For example, it is possible to create a policy that allows the system to install a given component, but prohibits installing specific versions of that component.

For more information about update policies, see ["Creating Update Policies"](#).

About Profiles

Profiles define how a standard management tasks are performed. With the help of wizards, you create a customized set of profiles. Once created, authorized users can use the profiles to perform tasks, such as discovery and provisioning. The software has pre-defined profiles for some common OS tasks, such as reboot and check for installed security patches. You can choose to use the pre-defined profiles for those actions or create your own.

Profiles are organized in the UI by category. The profiles are discussed in more detail in various sections of the documentation. The following are the profile categories and where you can find more specific information about that task and profile:

- Discovery: See [Chapter 2, "Asset Management"](#)
- Service Processor: See [Chapter 11, "Hardware"](#)
- RAID Controller: See [Chapter 11, "Hardware"](#)
- Firmware: See [Chapter 11, "Hardware"](#)
- Dynamic System Domain: See [Chapter 11, "Hardware"](#)
- OS Provisioning: See [Chapter 12, "Operating Systems"](#)
- Logical Domain: See [Chapter 15, "Oracle VM Server for SPARC"](#)
- Solaris Zone: See [Chapter 14, "Oracle Solaris Zones"](#)
- Virtual Machine: See [Chapter 16, "Oracle VM Server for x86"](#)
- BIOS Configuration: See [Chapter 11, "Hardware"](#)
- Boot Environment: See [Chapter 12, "Operating Systems"](#)
- Storage Appliance Update: See [Chapter 11, "Hardware"](#)
- Update Policies: See [Chapter 13, "Operating System Updates"](#)
- Update Profiles: See [Chapter 13, "Operating System Updates"](#)

With the exception of the Discovery profile, these types of management profiles require a deployment plan to execute the tasks on specific targets. When you create a profile, the default setting is to create a corresponding deployment plan. See [Deployment Plans](#) for more information on these types of plans.

Deployment Plans

Deployment Plans use profiles to perform standard management activities in a consistent and repeatable manner. A deployment plan defines the sequence of operations or steps that must be carried out on an asset to deploy it together with the specification or profile that each step applies and the resources that are required to apply it such as network addresses and system names.

A comprehensive set of deployment templates is available for you to use to create a variety of deployment plans. Each template is an unbound deployment plan which defines the steps of execution, but not the profiles and assets. A deployment template enable you to define a task and the associated resources (such as images and network addresses), that are needed to complete the task.

The software provides several plans that contain more than one step. In some cases, you can associate steps in a plan with another plan. The associated plan is referred to as a nested plan. You can use nested plans as shared building blocks, much in the same way as profiles are used. Configuring a single nested plan once and reusing it in many other plans reduces the number of individual operations that you must complete in the UI. The templates appear in the user interface in alphabetical order, not by type. However, it is useful to categorize the plans as simple, multi-step, and complex. See the different categories for a description of the plans.

The following lists the plans that are located in each category:

- [About Simple Deployment Plans](#)
 - Configure RAID
 - Configure Service Processor
 - Provision OS
 - Update Firmware
 - Update BIOS Configuration
 - Update Storage Appliances
 - Create Oracle Solaris Zones
 - Create Virtual Machines
- [About Multi-Step Deployment Plans](#)
 - Install Server
 - Create Dynamic System Domain
 - Create Logical Domains
 - Software Deployment/Update
 - Update Solaris 11 OS
 - Create Boot Environment
 - Update Firmware and Install Oracle VM Server for SPARC
- [About Complex Deployment Plans](#)
 - Configure M-Series Hardware, Create, and Install Domain
 - Configure Server Hardware and Install OS
 - Configure and Install Dynamic System Domain
 - Configure and Install Logical Domains

- Configure and Install Oracle Solaris Zones
- Configure and Install Virtual Machines

Use the templates to create your own deployment plans and configure the plans using profiles. The settings and values in the profiles bound to each step are defaults. You can modify the plan before it is actually applied. You can further constrain the profile settings and values by the target systems to which the plan is applied.

With the exception of complex plans, the templates do not allow you to add steps to a plan. You can use only those steps that are defined in the template from which the plan is derived. Complex plans enable you to add one or more deployment and operational plans inside a complex type plan.

About Simple Deployment Plans

The software provides you with the ability to create, configure, manage and execute deployment plans which drive the hardware, firmware and software provisioning activities in a repeatable fashion.

You create plans from defined templates. Each plan defines the sequence of steps that must be carried out for configuration or provisioning of a system. Plans might contain a single step or a sequence of multi-steps. Each step in the plan is configured by associating a profile or another plan.

A simple plan contains a single step with a single image. You can define the image used by the plan, but you cannot add more than one image or add steps.

The following plans are simple deployment plans:

- **Configure RAID:** Use this plan to configure the RAID controller on a server. See [Configuring a RAID Controller](#) for how to use the profile and plan.
- **Configure Service Processor:** Use this plan to configure the service processor on a chassis. See [Configuring the Service Processor](#) for how to use the profile and plan.
- **Provision OS:** Use this plan to provision an operating system. See [Provisioning Operating Systems](#) for how to use the profile and plan.
- **Update Firmware:** Use this plan to update firmware. See [Firmware Provisioning](#) for how to use the profile and plan.
- **Update BIOS Configuration:** Use this plan to update the BIOS configuration of servers. See [Creating a BIOS Configuration Profile and Plan](#) for how to use the profile and plan.
- **Update Storage Appliances:** Use this plan to update storage appliance software. See [Storage Libraries](#) for how to use the profile and plan.
- **Create Oracle Solaris Zones:** Use this plan to create zones. See [Creating a Zone Profile](#) and [Creating and Deploying Zone Plans](#) for how to use the profile and plan.
- **Create Virtual Machines:** Use this plan to create virtual machines for Oracle VM Server for x86. See [Create Virtual Machines](#) and [Deploying Virtual Machine Plan](#) for how to use the profile and plan.
- **Create Dynamic System Domain:** Use this plan to create dynamic system domains. See [Configuring a Dynamic System Domain](#) for how to use the profile and plan.
- **Create Logical Domains:** Use this plan to create logical domains. See [Creating a Logical Domain Profile](#) for how to use the profile and plan.

In most cases, when you create a profile, the default action also creates a simple deployment plan.

About Multi-Step Deployment Plans

Several deployment templates with multi-step sequences are available. These plans are designed to provide you with a customized and repeatable way to perform common operations with a single click.

The following plans are multi-step deployment plans:

- **Install Server:** Use this plan to provision the server and update the OS.
- **Software Deployment and Update:** Use this plan to apply script based update profiles. See [Updating an Operating System From a Deployment Plan](#) for how to use the profile and plan.
- **Update Solaris 11 OS:** Use this plan to update Oracle Solaris 11 operating systems. See [Updating Oracle Solaris 11](#) for how to use the profile and plan.
- **Create Boot Environment:** Use this plan to create Oracle Solaris boot environments. See [Creating an Oracle Solaris 11 Boot Environment](#) and [Creating an Oracle Solaris 10 Boot Environment](#) for how to use the profile and plan.
- **Update Firmware and Install Oracle VM Server for SPARC:** Use this plan to update firmware and then install Oracle VM Server for SPARC on the system. See [Installing Oracle VM Server for SPARC](#) for how to use the profile and plan.

The steps in a multi-step plan are defined in the template. You cannot skip steps or add steps.

About Complex Deployment Plans

You can use a combination of profiles, deployment plans, and operational plans to create a complex deployment plan that enables you to automate a variety of detailed workflows into a single plan. Complex plans provide flexibility to structure plans that meet your local requirements, increasing consistency and allowing for a greater level of automation.

When you create complex deployment plans, you can choose to skip a step in the plan. Skipped steps are not processed when the plan is applied. You can replicate certain steps to perform the same operation but using a different profile or plan. You can also add one or more deployment plans and operational plans.

The following plans are complex deployment plans:

- **Configure M-Series Hardware, Create and Install Domain:** Use this plan to configure an M-Series server, create dynamic system domains, provision OS on the domains, and update the domains.
- **Configure Server Hardware and Install OS:** Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- **Configure and Install Dynamic System Domain:** Use this plan to create dynamic system domains, provision and update OS on the domains.
- **Configure and Install Logical Domains:** Use this plan to create logical domains and provision OS on the logical domains.

Related Resources for Plan Management

For profile and plan details, and how to use individual profiles and plans, go to one of the following resources.

- [Chapter 2, "Asset Management"](#)
- [Chapter 12, "Operating Systems"](#)
- [Chapter 11, "Hardware"](#)
- [Chapter 13, "Operating System Updates"](#)
- [Chapter 14, "Oracle Solaris Zones"](#)
- [Chapter 15, "Oracle VM Server for SPARC"](#)
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

Part III

Operate and Maintain

Part III contains the following chapters:

- [Chapter 9, "Incidents"](#)
- [Chapter 10, "Reports"](#)
- [Chapter 11, "Hardware"](#)
- [Chapter 12, "Operating Systems"](#)
- [Chapter 13, "Operating System Updates"](#)

Incidents

This chapter describes how you can use the software to identify, assign, and resolve incidents.

The following information is included:

- [Introduction to Incident Lifecycle Management](#)
- [Roles for Incident Management](#)
- [Actions for Incident Management](#)
- [Location of Incident and Service Request Information in the User Interface](#)
- [Incidents](#)
- [Message Center](#)
- [Annotations](#)
- [Incidents Knowledge Base](#)
- [Annotations in the Incidents Knowledge Base](#)
- [Incident Lifecycle Management](#)
- [Oracle Services and Service Requests](#)
- [Related Resources for Incident Management](#)

Introduction to Incident Lifecycle Management

When an asset is not operating within the parameters defined in the monitoring rules and policies, Oracle Enterprise Manager Ops Center generates an alert and an incident.

An alert indicates that a monitored asset is not performing as expected. The monitoring rule parameters determine when an alert is triggered and the severity: Informational (info), Warning, or Critical.

An incident is raised when the monitoring rule is asserted by the raising of one or more alerts that the monitoring rule requires. New alerts will update an open incident. One or more subsequent alerts trigger the monitoring rule, which notifies the incident management system. The incident management system detects that there is already an open incident for the monitoring rule for that asset and correlates the alerts under the open incident and the worst severity level is associated with the incident. For example, when the incident is at a Critical severity level and a new Warning alert is added to the incident, the incident severity remains at the Critical level.

Oracle Enterprise Manager Ops Center uses a help desk approach to manage the incidents in your data center. All open incidents appear in the Message Center. You can assign incidents to others for resolution, add comments, provide a list of possible causes and impacts, provide recommendations, add utilities or scripts to resolve an issue, view progress, and open a service desk ticket.

The following are the main components that help you to track and manage known issues on your monitored assets:

- **Message Center:** Central location for details on incidents, notifications, and service requests.
- **Annotations:** Comments, suggested actions, and operational plans that enable you to effectively manage an incident.
- **Incidents Knowledge Base:** Database of your annotations and actions for specific types of incidents and severity levels.
- **Oracle Services and Service Requests:** Service request details on all requests submitted to My Oracle Support through Automated Service Request or the Oracle Enterprise Manager Ops Center UI. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about Automated Service Requests and how to enable them.

When an incident appears, you can assign it to a user for resolution and use annotations to add comments and suggested actions. You can build an Incident Knowledge Base that contains your annotations from specific incidents, or add suggested fixes or automated fixes for a specific type of incident.

If you cannot resolve an incident, you can open a service request with My Oracle Support inside the Oracle Enterprise Manager Ops Center UI. The information gathered by the software is automatically populated into the service request. You can track the status of any service request opened from within Oracle Enterprise Manager Ops Center, whether they belong to you or one of your co-workers.

Roles for Incident Management

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

Table 9–1 Incident Management Tasks and Roles

Task	Role
View Incidents	All
Assign Incidents	Fault Administrator
Add Annotation to incidents	Fault Administrator
Acknowledge incidents	Fault Administrator
Take Actions on Incidents	Fault Administrator
Mark Incidents as Repaired	Fault Administrator
Close Incidents	Fault Administrator

Table 9–1 (Cont.) Incident Management Tasks and Roles

Task	Role
Take Actions on Notification	Fault Administrator
Delete Notifications	Fault Administrator

Actions for Incident Management

Incident management is automatically enabled after an asset is discovered and managed. Alerts and incidents are generated based on the monitoring profiles and policies.

Actions for Incidents

You can perform the following actions for a specific incident:

- **View Alerts:** View all alerts that are generated when the state is outside the defined monitoring parameters.
- **View Annotations:** View the scripts or suggested actions that you or others in your organization have associated with an incident.
- **View Possible Impacts and Causes:** View the possible impacts and causes for an incident.
- **View Comments:** View comments that you or others in your organization have associated with an incident.
- **View Suggested Actions:** View suggested actions that you or others in your organization have associated with an incident.
- **Add Annotation to the Incident:** Add a script, suggested action, or comment to a specific incident.
- **Assign the Incident:** Assign an incident to a user.
- **Acknowledge the Incident:** Acknowledge, or accept, an incident that has been assigned to you.
- **Take Actions on the Incident:** Take action to resolve an incident. Options include executing a suggested action, executing a script that is part of an operational profile, executing a command, or executing a script.
- **Mark Incident as Repaired:** Identifies the incident as being fixed.
- **Close the Incident:** Closes the incident and removes it from the Message Center.
- **Open a Service Request:** Open a service request for the incident with My Oracle Support.

Actions for Service Requests

When the Enterprise Controller is connected to My Oracle Support and service requests are enabled, you can perform the following service request actions:

- **View all open service requests:** View all open service requests that were filed with the Oracle Enterprise Manager Ops Center software.
- **View your service requests:** View open service requests that you filed with the Oracle Enterprise Manager Ops Center software.

Actions for Incidents Knowledge Base

The Incidents Knowledge Base is located in the Operational Plans section of the Plan Management section. You can perform the following actions for the Incidents Knowledge Base:

- View and sort the annotations by type: The Incident Knowledge Base appears as a table with the name, description, subtype, target type, version and date the annotation was last updated.
- Add an annotation for an Incident type: Add a comment, suggested action, or automated action for a type of incident and monitored attribute. You define the severity levels that the annotation is applicable.
- Edit an annotation for an incident type: Edit an annotation for a type of incident.
- Delete an annotation for an incident type: Delete an annotation for a type of incident.

Location of Incident and Service Request Information in the User Interface

Incident and service request information is available in the Message Center in the Navigation pane. Use the Message Center to view and manage view notifications, service requests and warranty information for an asset. You can also view incident details by expanding the asset and clicking the Incident tab.

Table 9–2 Location of Incidents and Service Request Information in the UI

Object	Location
Unassigned Incidents	Expand the Message Center in the Navigation pane, then click Unassigned Incidents. Or, click the asset in the Assets section, then click the Incidents tab.
My Incidents	Expand the Message Center in the Navigation pane., then click My Incidents. Or, click the asset in the Assets section, then click the Incidents tab.
Incidents Assigned to Others	Expand the Message Center in the Navigation pane, then click Incidents Assigned to Others. Or, click the asset in the Assets section, then click the Incidents tab.
Open Service Requests	Expand the Message Center in the Navigation pane, then click Open Service Requests.
My Service Requests	Expand the Message Center in the Navigation pane., then click My Service Requests.
Notifications	Expand the Message Center in the Navigation pane, then click Notifications.
Relayed Incidents	Expand the Message Center in the Navigation pane, then click Relayed Incidents.
Relayed Service Requests	Expand the Message Center in the Navigation pane, then click Relayed Incident Requests.
Incident Knowledge Base	Expand Plan Management in the Navigation pane, then click Incident Knowledge Base.

Incidents

Each time a monitored attribute does not meet its monitoring rule, a new alert is generated. The first alert raises an incident, which appears in the Message Center. Subsequent alerts for the same rule and asset are correlated with the same open incident.

When a value for an attribute exceeds its monitoring rule and then later meets the rule, the alert is cleared automatically. The change in the attribute value does not automatically clear the incident. When an attribute's value is moving in and out of its monitoring rule's parameters, alerts are generated and cleared continuously. A new incident is only generated when the original incident is closed. When an incident is not yet closed, the new alerts are aggregated into the existing incident. You must close an incident to clear it from the user interface.

About Incident Severity Badges

Incident severity badges are an option that you can enable to display an incident severity badge next to an asset show the incident status. Badges also appear in the dashboard membership graph to show the status of an asset. See [Badges in Appendix E](#) for a description of each badge.


To Display Incident Severity Badges

Perform the following to display the badges in the Navigation pane:


1. Click **Setup** in the upper right corner of the UI.
2. Click **My Preferences**, then click **User Interface Preferences**.
3. Click the check box to make the badges visible in the user interface.

When an asset has an incident, the severity badge appears next to the asset icon in the Asset hierarchy. When it is the highest severity incident in the membership of a group, the severity badge also appears next to the parent assets. In [Figure 9-1](#), the operating system for xvmbrrm-t5140-2 has a critical incident. The critical incident badge also appears on the system and service processor. Any group that this OS is a member of, such as All Assets and Operating Systems, also display the badge.

Figure 9-1 Critical Incident Badge



When the incident is the only critical incident, the badge is removed when the incident is acknowledged, marked repaired, or closed. When open incidents are still present, the next highest severity badge displays. For example, when both a Critical and a Warning incident is detected and the Critical incident is acknowledged, the Critical badge is replaced with the Warning badge because that is now the highest level unacknowledged incident.

Figure 9–2 Warning Incident Badge

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

See [Chapter 4, "Monitoring Rules and Policies"](#) for more information about creating and maintaining monitoring rules and profiles.

An incident is one or more alerts on the same monitored attribute and asset. A ticker at the top of the UI contains the number of unassigned critical incidents, unassigned Warning incidents, number of relayed incidents, the number of critical incidents and the number of warning that are assigned to you. You can click an icon to display the incidents dashboard for that category.

Figure 9–3 Incidents Ticker

Message Center

Incidents, notifications, and service requests that are generated by managed assets appear in the Message Center. Your assigned user role determines the actions that you can perform in the Message Center.

Incidents, notifications and service requests appear in the following categories in the Message Center:

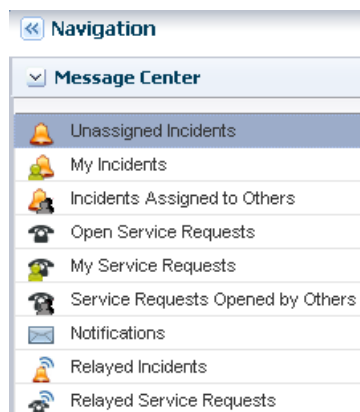
- **Unassigned Incidents:** Newly created incidents and those that have not been assigned an owner.
- **My Incidents:** All incidents that are assigned to you. You can perform additional actions to these incidents to manage their status, such as: Take Action, mark as being repaired, acknowledge, and open a service request for the incident.
- **Incidents Assigned to Others:** Incidents that are currently assigned to other users. You can view these incidents, but you cannot perform specific actions on them.
- **Open Service Requests:** All Service Requests that have been filed on assets that are submitted through Oracle Enterprise Manager Ops Center.
- **My Service Requests:** All Service Requests that you submitted on assets that are managed by Oracle Enterprise Manager Ops Center.
- **Service Requests Opened by Others:** All Service Requests that have been submitted on assets that are managed by Oracle Enterprise Manager Ops Center.
- **Notifications:** Information that is automatically generated by services running in the backend while monitoring the assets. Typically, notifications are less important than incidents.
- **Relayed Incidents:** All incidents reported from any discovered Oracle Engineered System. You must log in to the Oracle Enterprise Manager Ops Center instance

that manages each Oracle Engineered system to fix any incidents related to its assets.

- Relayed Service Requests: All open Service Requests for any discovered Oracle Engineered System.

Closed incidents and service requests do not appear in the Message Center.

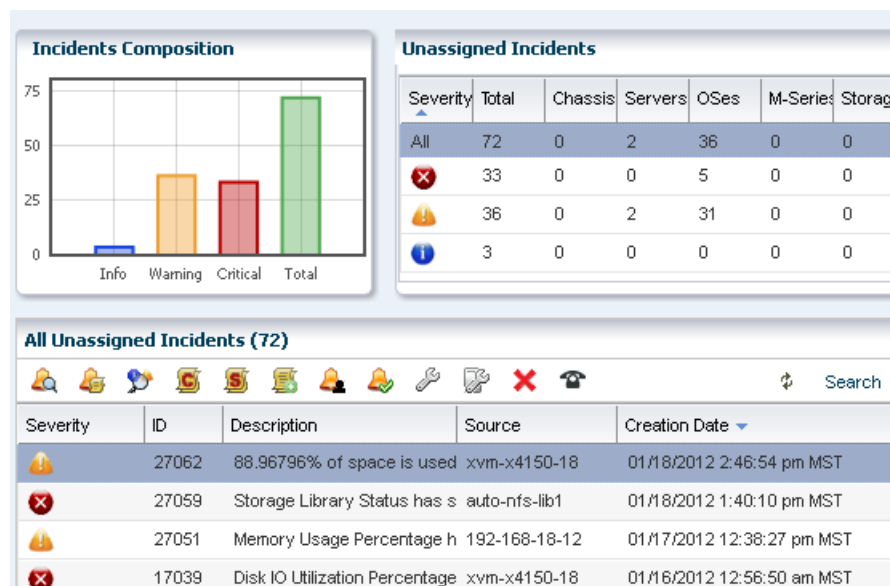
Figure 9–4 Message Center



A bar chart on the page visually displays the number of new incidents by severity and the total number of unassigned incidents. The page also contains a table that categorizes the incidents by severity and asset type. Select a row in the table to display all unassigned incidents for the selected severity category. A third table, at the bottom of the dashboard provides details and action icons.

The All Unassigned Incidents table shows all open unassigned incidents. Each incident receives an ID to help you to track the issue. The table includes a description, source, Creation Date, and URL field for each incident. Hover over the URL icon for an incident to display a pop-up window. The default sort is by the creation date; however, you can sort by any column.

Figure 9–5 Unassigned Incidents Dashboard



The following information is available in the Unassigned Incidents table:

- Severity: The severity icon shows the severity level, either informational (info), warning, or critical.
- ID: A generated ID assigned to the incident to help track the issue.
- Description: A brief description of the incident.
- Source: The asset that is generating the incident.
- Creation Date: The date and time that the incident generated.
- URL: Contains incident details. Hover over the URL icon to display the duration of the incident, when the incident was assigned, to whom the incident was assigned, any suggested actions, the source of the incident, and a larger description field.

A row of icons at the top of the Unassigned Incidents table provides actions that are available to you, based on your user role. The following action icons are available for incidents, the actions are in the order that the icons appear in [Figure 9–6](#):

- View Alerts
- View Annotations
- View Possible Impacts and Causes
- View Comments
- View Suggested Actions
- Add Annotation to Incidents
- Assign Incidents
- Acknowledge Incidents
- Take Actions on Incident
- Mark Incidents as Repaired
- Close Incidents

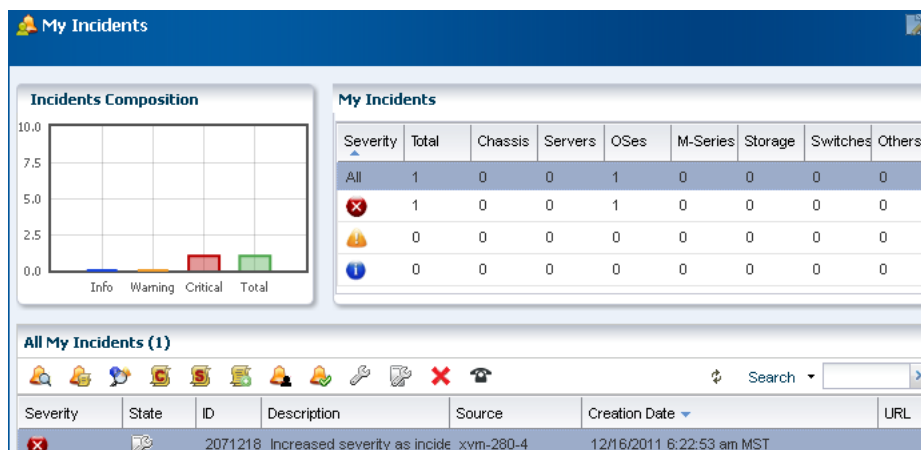
Figure 9–6 Incident Actions



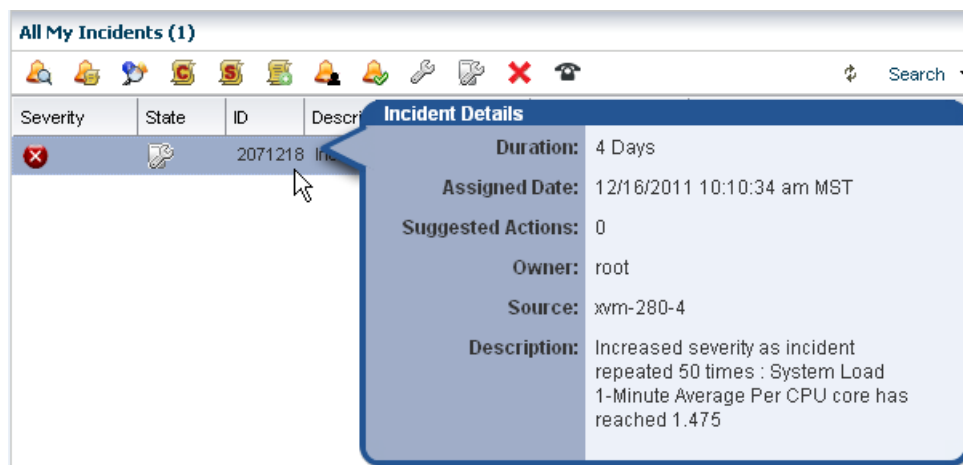
Incidents Dashboards in the Message Center

Incidents assigned to you appear in the My Incidents dashboard in the Message Center. Anybody with the appropriate permission level can see the incidents in your queue, can reassign incidents to another user, and can assign you new incidents. Unassigned Incidents and Incidents assigned to others have their own categories in the Message Center. Each Incident dashboard is laid out the same.

A bar chart on the page visually displays the number of new incidents by severity. The page also contains a table that categorizes the incidents by severity and asset type.

Figure 9–7 My Incidents

Select a row in the table to display all incidents for the selected severity category and to drill down for more details. The software assigns each incident an ID for tracking purposes. Hovering your mouse over any field in the incident row displays details.

Figure 9–8 Incident Details

The icons in the center pane enable you to perform actions on a specific incident or view the alerts that make up an incident in the center pane. Each icon has text to define the action.

Figure 9–9 Incident Icons

The following actions, as described in [Actions for Incident Management](#), are available:

- View Alerts
- View Annotations
- View Possible Impacts and Causes

- View Comments
- View Suggested Actions
- Add Annotation to Incidents
- Assign Incidents
- Acknowledge Incidents
- Take Actions on Incident
- Mark Incidents as Repaired
- Close Incidents
- Open a Service Request

You can configure the software to send you e-mail or pager notification of incidents for critical assets or severity levels. The Notifications appear in the Message Center.

Annotations

Annotations are comments, suggestions, or automated operations with associated scripts that you can use to document an incident. They are defined by an asset type, asset resource type, an attribute, or an incident type.

Annotations are as powerful as you want to make them. You can View Annotations, View Possible Impacts and Causes, View Comments, View Suggested Actions, and Add Annotations that are associated with an incident in the Message Center or from the Incidents tab in the Asset view, as shown in [Figure 9–10](#).


Figure 9–10 Annotation Icons



You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, and when you close an incident.

Figure 9–11 Add Comment Type Annotation

Add Annotation

Severity	ID	Description	Source
	1653	Memory Usage Percentage has reached 82.60498	ocbrm-ijpgs13

* Annotation Type:

Associated Operational Plan:

* Synopsis:

Note: Investigating why memory usage percentage is high. |

☐ Save this annotation in the Incidents Knowledge Base to associate the annotation with all incidents of this type and severity.

You can associate annotations with an incident instance or an asset type. The Incidents Knowledge Base contains your annotations, by asset type, and stores the information on the Enterprise Controller. You can view this annotation by browsing the Incidents Knowledge Base.

For more robust incident management, associate an annotation with an incident and store the annotation in the Incidents Knowledge Base to provide an automated solution or a suggested action when a similar incident is detected. To provide an automated solution, create an operational plan that contains a shell script. When a specific incident occurs, the script is executed automatically. To provide a suggested fix or course of action for a specific incident, you can create a text only annotation that provides a suggested course of action, or you can include a shell script.

When a rule is triggered, and an alert or incident is identified, the software checks the Incidents Knowledge Base and Incident Profiles for the Incident Type and any associated annotations. The software generates an incident of the defined type and severity level and attaches any annotations. When an automated operation annotation is associated with the incident, the script is executed. When a suggested action annotation is associated with the incident, the text and script (when available) appear in the incident details.

About Annotation Views

Annotations are available from several different views within the UI:

- In the Message Center, when an annotation is associated with an incident
- In the Asset view, when an annotation is associated with an incident
- In the Incident Knowledge Base (in the Plan Management section), when an annotation is associated with an asset type

When an annotation is associated with an incident, you can click the View Annotations icon to see the associated annotation.

You can create annotations by an asset type, asset resource type, an attribute, or an incident type. The annotations are located in the Incidents Knowledge Base (KB). Annotations are automated operations with associated operational profile, suggested fixes or actions, or text only comments. You can associate Automated Action and Suggested Action annotations with an operational profile, which can contain a script.

You can also add an annotation to the Incidents KB when you add an annotation to a specific incident.

A comment is a type of annotation. You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, and when you close an incident. You can also use annotations to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions. To add a comment, see [Adding an Annotation](#).

Annotations in the Incidents Knowledge Base (KB) are associated with types of incidents. You can delete an annotation in the Plan Management section of the UI.

Incidents Knowledge Base

Annotations enable you to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions.

You can use your annotations to build an Incidents Knowledge Base that is based on your company policies, procedures, and incidents.

You can add annotations to the Incidents Knowledge Base in Plan Management or from a specific incident. Either way, annotations enable you to provide solutions or recommended actions for specific incidents.

The following types of Annotations are available:

- **Automated Operation:** You can create an annotation that references an Operational Profile and automatically executes the profile when a specific incident occurs.
- **Suggested Action:** You can provide a suggested fix or course of action for a specific incident. The suggested course of action can be text only, or you can reference an Operational Profile.
- **Comment:** You can write notes about the situation, such as the current status of an incident. This is a text only field.

You can build an Incident Knowledge Base that contains your annotations from specific incidents, or add suggested fixes or automated fixes for a specific type of incident. You can add possible causes and impacts or you can add an annotation. An annotation is either a comment or a suggested action. Annotations enable you to associate an operational plan that contains a utility or script to correct a specific issue. You can choose to associate the annotation with all incidents of the same type and severity and save it in the Incidents Knowledge Base. The next time the incident occurs, you can access the possible causes and impacts and annotations identified in the previous incident and resolve the issue more quickly.

When an incident is detected, Oracle Enterprise Manager Ops Center checks the Incidents Knowledge Base and Incident Profiles for the asset type and corresponding incident type. Any associated annotations are added to the Incident and Operational Profiles referenced in Automated Operation annotations are executed against the asset on which this incident was open. You can create, update, and delete annotations in the Incidents Knowledge Base.

Annotations in the Incidents Knowledge Base

For more robust incident management, you can associate an annotation with a rule and store the annotation in the Incidents Knowledge Base to provide an automated solution or a suggested action when a specific incident is detected. To provide an automated solution, create an operational plan that contains a shell script. When a specific incident occurs, the script is executed automatically. To provide a suggested fix or course of action for a specific incident, you can create a text only annotation that provides a suggested course of action, or you can include a shell script.

When a rule is triggered, and an alert or incident is identified, the software checks the Incidents Knowledge Base and Incident Profiles for the Incident Type and any associated annotations. The software generates an incident of the defined type and severity level and attaches any annotations. When an automated operation annotation is associated with the incident, the script is executed. When a suggested action annotation is associated with the incident, the text and script (when available) appear in the incident details.

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

1. Open the incident from the Message Center or Assets view.
 - Message Center View: Click **Message Center**, then click an Incident category: **Unassigned Incidents**, **My Incidents**, or **Incidents Assigned to Others**.
 - Assets View: Click the asset in the Assets section of the Navigation pane, then click the **Incidents** tab in the center pane.
2. Click the incident in the center pane.
3. Click the **View Annotations** icon.

Example 9–1 Example of Using Annotations

The CPU usage on a Sun Fire x4150 host is exceeded and an incident is generated. You assign the incident to Lee. Lee is concerned because these systems are often used to host Oracle Solaris Zones. Lee adds the following comment to the incident: "This asset is not powerful enough and cannot cope with the load". Lee also wants to associate an annotation with the Global Zone asset type to recommend checking for processes that are consuming excessive CPU usage on the Global Zone. Lee adds the following annotation to the asset type: "Run the 'prstat 1 1' command to check which processes are taking CPU." The annotation is saved in the Incidents Knowledge Base and appears the next time CPU usage is exceeded on a global zone asset type.

Incident Lifecycle Management

Incident management in Oracle Enterprise Manager Ops Center consists of several components that are designed to work together to simplify managing incidents for a large number of assets. The components include monitoring rules, suggested actions, and methods for automating incident identification and resolution.

Monitoring includes a standard set of monitoring rules, consisting of an asset's attribute and the threshold value for that attribute. When Oracle Enterprise Manager Ops Center performs monitoring, it generates alerts, which connect to both the incident management and notification features.

When an asset is not operating within the parameters defined in the monitoring rules and policies, Oracle Enterprise Manager Ops Center generates an incident and displays the information in the Unassigned Incidents section. Incidents appear as Informational (Info), Warning, or Critical severity. All incidents appear in the Message Center. When an incident is first detected, it appears in the Unassigned category. When you assign an incident to yourself, it moves from the Unassigned Incidents to My Incidents. When an incident is assigned to someone else, it appears in Assigned Incidents.

For example, the CPU usage on a Sun Fire x4150 host is exceeded and an incident is generated. You assign the incident to Bob. Bob is concerned because these systems are often used to host Oracle Solaris Zones.

Bob reviews the incident and adds the following comment to the incident: This asset is not powerful enough and cannot cope with the load. Bob also wants to associate an annotation with the Global Zone asset type. He wants to add a recommended action annotation to the asset type to check for processes that are consuming excessive CPU usage on the Global Zone. He adds the following annotation to the asset type: Run the `prstat 1 1` command to check which processes are taking CPU. The annotation is saved in the Incidents Knowledge Base and displays the next time CPU usage is exceeded on a global zone asset type.

Methods of Incident Management

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

- Message Center: View the status of incidents and assign incidents.
- Annotations: Add notes and change status. Use annotation options to provide recommended actions or fixes, or add custom scripts to provide an automated response to an incident.
- Operational Plans: Deploy a shell script against a specific asset, or asset sub-type to automate incident resolution.
- Incidents Knowledge Base: Collect comments and suggested actions for known issues for future use.

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

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

You can manage incidents from either the Message Center or from the Asset view. You can view and add comments and annotations, take action on an incident, and close incidents.

- The Message Center provides a list of all incidents. Select an incident to see its details and activity.
- From the Asset tree in the Navigation pane, select the asset and then click the Incidents tab to see a list of incidents for that asset.

When you have the Manage or Admin role for the asset, you can take action on the incident. The person assigned to the incident must also have the Manage or Admin role. When the icon is not active, you do not have the appropriate role.

Viewing Unresolved Incidents

You can view unresolved incidents for a specific asset or by incident:

- To view unresolved incidents for a specific asset, click the asset in the Navigation pane, then click the Incidents tab in the center pane.
- To view unresolved incidents from the Message Center, click one of the following:
 - Unassigned Incidents
 - My Incidents
 - Incidents Assigned to Others

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

View high-level details by hovering your mouse over the incident or clicking the incident in the Unresolved Incidents table. You can drill down to view the alerts that make up the incident by clicking the incident, then clicking the Alerts icon in the center pane.

Viewing Incident Details

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

- How long the incident has been open, or the duration
- When the incident was assigned
- The number of suggested actions for the incident
- Who is assigned to the incident
- Which resource is affected, or the cause
- A description of the incident

Perform the following steps to view incident details:

1. Select **Assets** in the Navigation pane.
2. Select an asset that has an incident badge next to the icon. The Dashboard page displays with the status of the asset.
3. Click the **Incidents** tab.
4. Hover over the incident to display the incident details.
5. To display the alerts that are associated with the incident, click the **Alerts** sub-tab or click the **Alerts** icon in the center pane. The alerts that make up the incident are displayed, including the current and highest alert status, and the alert history.

Assigning an Incident

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

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

1. To display an incident from the Message Center, click Message Center, then click **Unassigned Incidents** in the navigation pane.

To display an incident from the asset view, click the asset in the Navigation pane, then click the **Incidents** tab.


2. Select one or more incidents in the center pane, then click the **Assign Incidents** icon.
3. Select a user name from the **Assign To** list, which is the list of users who have either the Manage or Admin role for the asset. To relocate an assigned incident back to the Unassigned Incidents queue, select **No One** from the list.
4. (Optional) Add a note in the text field.
5. Click **Assign Incidents**.

Acknowledging Incidents

Acknowledging an incident indicates that you are investigating the issue. You can acknowledge an incident when you have the Admin or Manage role for the asset on which the incident is identified.


When severity badges are enabled and an incident occurs on an asset, a severity badge appears next to the asset in the Asset hierarchy. When it is the highest severity incident in the membership, it also appears next to the parent assets. In [Figure 9–12](#), the operating system for xvmbrm-t5140-2 has a critical incident. The critical incident badge also appears on the system and service processor. The badge appears on any group that this operating system is a member of, such as All Assets and Operating Systems.

Figure 9–12 Critical Incident Badge on Asset and Parent Assets



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

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

Figure 9–13 Effect of Acknowledging a Critical Incident

1. Open the incident from the Message Center or Assets view.
 - Message Center View: Click Message Center, then click an Incident category: Unassigned Incidents, My Incidents, or Incidents Assigned to Others
 - Assets View: Click the asset in the Assets section of the Navigation pane, then click the Incidents tab in the center pane.
2. Select one or more incidents, then click the **Acknowledge Incidents** icon in the center pane.

Adding an Annotation

Annotations are defined by the asset type. Annotations are comments, a suggested action, or a reference to an operational profile. Any user can add an annotation to a Incident. To add an entry to the Incidents Knowledge Base requires Oracle Enterprise Manager Ops Center Admin permissions.

1. Open the incident from the Message Center or Assets view.
 - Message Center View: Click **Message Center**, then click an Incident category: **Unassigned Incidents**, **My Incidents**, or **Incidents Assigned to Others**
 - Assets View: Click the asset in the **Assets** section of the Navigation pane, then click the **Incidents** tab in the center pane.
2. Select the incident, then click the **Add Annotations** icon in the center pane.
3. Select one of the following types from the Annotation Type from the drop-down list:
 - **Comment**: Text only option that is designed to be used to add a note or editorial comment.
 - **Suggested Action**: Text required and a script is optional.
4. Select an operational plan from the drop-down list of operational profiles defined for the type of asset on which this incident is open.
5. The Synopsis field is completed based on the annotation type. Edit the synopsis, as needed. The UI does not have a character limit, but the API allows for 80 characters.

Note: When you enter more than 80 characters, the synopsis is truncated to the first 80 characters when viewed in the annotation.

6. Type a description or instructions in the **Note** field. There is no character limit.

7. To add the annotation to the Incidents Knowledge Base and include the annotation for every incident of this type and severity, click the check box.

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

8. Click **Save and Execute** or click **Save**.

Displaying Annotations

You can display annotations for an asset type in the Incidents Knowledge Base.

1. Click **Plan Management**.
2. Expand Incidents Knowledge Base in the Navigation pane, then select the asset type. The annotations associated with the asset type appear in the center pane.

Viewing Comments

A comment is a type of annotation. You can add informational comments and notes to an issue while you are working on a resolution, when you mark an incident as fixed, or when you close an incident. You can also use annotations to build a Incidents Knowledge Base that contains a mixture of comments, suggested actions, and automated actions. To add a comment, see [Adding an Annotation](#).

1. Expand the **Message Center**, then click one of the following:
 - **Unassigned Incidents**
 - **My Incidents**
 - **Incidents Assigned to Others**
2. Click the incident in the center pane, then click the **View Comments** icon.

Taking Action on a Incident

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

The incident management functionality provides varying levels of control. You can build the Incidents Knowledge Base with annotations that contain a combination of automated actions, suggested actions, and comments. You can execute one or more of the suggested actions associated with the Incident or use an operational plan to correct the incident. When an action is not available in a suggested action or Operational Profile, you can execute a command or a custom script that is stored on the managed asset or on the Enterprise Controller.

See [Adding an Annotation](#) for information on adding a suggested action to a Incident.

See [Incidents Knowledge Base](#) to add a suggested action or automated action in the Incidents Knowledge Base. See [Creating an Operational Profile and Plan](#) for information about creating an operational profile.

1. Open the incident from the Message Center or Assets view.

- Message Center View: Click Message Center, then click an Incident category: Unassigned Incidents, My Incidents, or Incidents Assigned to Others
 - Assets View: Click the asset in the Assets section of the Navigation pane, then click the Incidents tab in the center pane.
2. Select the incident.
 3. Click the **Take Actions** on a Incident icon in the center pane.
 4. Select the action to perform:
 - When the Incidents Knowledge Base has provided a suggested action for the incident, select **Execute the Selected Suggested Action** option and then select the action from the table.
 - When an operational plan has a suggested action, select the **Execute an Operational Plan** option, then select the plan from the drop-down list.
 - To run a script or command that is not part of a suggested action or operational plan, select the **Execute a Command or Script File** option.
 - To execute a command, enter the command in the field.
 - To browse for a script, click **Browse** and then select the script from the **File Chooser** popup.
 5. Select where to run the script, on the managed asset where the incident is open, or on the Enterprise Controller.
 6. Define the time out period for the action, in minutes, hours, or days.
 7. (Optional) Add a note describing the action taken.
 8. Click **Execute Selected Action**.

Maintenance Mode

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

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

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

When 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.

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 an Incident Repaired

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

When the incident was previously in an Unassigned state or Assigned to someone, the severity was taken into account in the computation of the highest severity to propagate up to antecedent assets in the navigation pane. After marking this incident as repaired, its severity badge does not appear in the assets list in the navigation pane.

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

- Unassigned Incidents
- My Incidents
- Incidents Assigned to Others

or

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

2. Select one or more incidents, then click the **Mark Incidents as Repaired** icon in the center pane.
3. (Optional) Select the incident, then add a Note.
4. Click **Tag Incidents as Being Repaired**.

Closing an Incident

The incident stays open until you close it, even if the alerting condition is cleared. To remove an incident from the Message Center and the asset view, you must close the incident or take no action on it for seven (7) days. When any action is taken on an incident, such as adding an annotation, the counter is reset.

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

When an incident is closed, its status changes to Closed, the incident is deleted from the list of active incidents, and the incident is no longer displayed in the UI. You can retrieve information about a closed incident for 60 days by using the public API. After 60 days, closed incidents are permanently deleted. To edit the time limit, you must edit the value in the public API. You can disable the time limit by setting the value for the number of days to 0.

Note: When the monitoring condition is still true after the incident is closed, a new alert is raised and a new incident is created.

To Close an Incident

You can close an incident from the asset view or from the following categories in the Message Center:

- Unassigned Incidents

- My Incidents
- Incidents Assigned to Others

Perform the following steps to close an incident from the asset view:

1. Click the asset in the Assets section of the Navigation pane, then click the Incidents tab.
2. Select one or more incidents, then click the **Close Incidents** icon in the center pane.
3. (Optional) Select the incident, then add a Note.
4. (Optional) To temporarily disable the monitoring rule that identified the incident, click the **Action** check box, then define when to enable the monitors.

This action does not disable the monitoring rule for all assets. The action disables the monitoring rule for only the assets that were related to the incident to avoid raising a similar incident on the same assets.

5. Click **Close Incidents**.

Oracle Services and Service Requests

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

Use the Oracle Services feature to view the contract or warranty information and any service requests for a specific asset. You can also view service requests that were the result of an alert or incident in Oracle Enterprise Manager Ops Center, view service request details, and file a service request.

- Contracts and Warranties

Maintaining a valid inventory of the assets in your data center, including contracts and warranties, can be a time-consuming and labor-intensive process. Use Oracle Enterprise Manager Ops Center to display current contract and warranty information for a specific asset, or view the entitlements associated with your Oracle online account. When a contract or warranty is about to expire, Enterprise Manager Ops Center generates an alert.

- Service Request

You can create new service requests, review your requests, and review the requests of other users.

You can file requests manually, or you can provide credentials and contact information for assets and configure Oracle Enterprise Manager Ops Center to generate service requests. When an incident occurs with an asset, a service request is automatically created using the credentials and contact information that you provided.

Note: You cannot display service requests created outside of Oracle Enterprise Manager Ops Center. To see the status of 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 Manager Ops Center in Connected Mode.

To access the My Oracle Support database, your user must be registered as a My Oracle Support user. This is the same account that is used to access My Oracle Support at

To determine if you are running in Connected Mode and have access to My Oracle Support, check the icons in the upper right corner of the UI, as shown in [Figure 9–14](#). If an icon does not contain color, you are not connected.

- The World icon indicates the status of the Internet connection.
- The Shield icon indicates the status of the connection to the Oracle Knowledge Base.
- The Phone icon indicates the status of the connection to My Oracle Support Services.

Figure 9–14 *Connection Icons*



Viewing Contract and Warranty Information

You can display contract information by asset, or you can obtain entitlements associated with all contracts that are associated with a user. Contract and warranty information is available for managed servers that have a serial number associated with a contract in the My Oracle Support database.

The contract and warranty information in Enterprise Manager Ops Center is updated each week so contract changes or new contracts might take up to seven days to appear in the user interface. When a contract or warranty is about to expire, an alert is displayed as an Incident in the Message Center and the contract details appear in orange text in the asset's Summary tab.

Note: Updating contract and warranty information requires running the product software in Connected mode. If you change to Disconnected Mode, the contract information becomes outdated.

To View Contract and Warranty Information For an Asset

1. Select a hardware asset in the Navigation pane, from either the **All Asset** list or from a group.
2. Click the **Summary** tab.

When Enterprise Manager Ops Center is in Connected Mode and the serial number of the selected asset is associated with a contract, a Support row is added to the summary. The Support field contains the contract ID and an expiration date.

- If the contract is within 90 days of expiration, the information is displayed in an orange font.

- If the contract has expired, the information is displayed in a red font.

To View All Contracts Associated with a My Oracle Support Account

1. Click the **Enterprise Controller** in the Administration section of the Navigation pane.
2. Click **Edit Authentications** in the Actions pane. The Edit Authentications window is displayed with online user names and associated contracts.

Viewing Service Requests

You can see all current and completed service requests that are filed through Oracle Enterprise Manager Ops Center. The service request

To View Service Requests

1. Click **Message Center** in the Navigation pane.
2. Click **Open Service Requests, My Service Requests, or Service Requests Opened by Others** to display a list of requests.
3. To view details of a particular service request, highlight a row, then click the **View Service Request** icon.

Figure 9–15 View Service Request

Oracle Enterprise Manager Ops Center - View Service Request

Information				
Request Number	3-1863062401			
Severity	1-Critical			
Summary	Problem detected on: hs-x4100-2 - 172.20.28.190			
Last Updated	Tue Oct 19 2010 15:13:34 GMT-0600 (MST)			
Contact	MOSPatchOCMCollector Test			
Status	Open			
Sub Status	New			
SR Email	mospatchtest14@sleepycat.com			
SR Telephone	415-999-0000			
Support ID	17251035			
Address	Oracle UK Headquarters Oracle Parkway CA Reading RG6 1RA United Kingdom			

Description				
Problem detected by Ops Center instance: https: * Ops Center Problem ID: 432 Problem Severity: CRITICAL * Problem Description: hs-x4100-2 - 55.775578% of space is used on / filesystem.				
Fri Oct 01 2010 22:37:10 GMT-0600 (MST)				
Problems reported by Ops Center: Current Problem: Severity: CRITICAL ID: 432 State: UNASSIGNED Description: hs-x4100-2 - 55.775578% of space is used on / filesystem. Creation Date: Fri Oct 01 16:36:02 MDT 2010				
Associated Alerts:				
Alert Type	Alert Source	Attribute	Current Status	Highest
Threshold	hs-x4100-2	FileSystemUsages.name=/.usedSpacePercentage	CRITICAL	CRITICAL

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 incident or from an asset. See [Requirements for Oracle Services](#) for requirements that must be met before successfully filing a service request ticket. For example, if the asset is not registered in My Oracle Support, the service request job fails. If the Open Service request action is disabled, there is no connection to My Oracle Support.

To File a Service Request From a Incident

1. Click Message Center in the Navigation pane.
2. Click **My Incidents** or **Unassigned Incidents**.
3. Select the incident, then click the **Open Service Request** icon in the center pane.

To File a Service Request From an Asset

1. Select the hardware in the Assets section of the Navigation pane.
2. Click **Open Service Request** in the Actions pane.

Automated Service Requests

Instead of manually filing a service request, you can configure Oracle Enterprise Manager Ops Center to automatically create service requests for known issues. When Automated Service Requests (ASRs) are enabled, Oracle Enterprise Manager Ops Center automatically generates service requests based on critical incidents. Contact information for the ASR is taken either from Oracle Enterprise Manager Ops Center or from the Customer Service Identifier (CSI) associated with the asset. Annotations are added to the incident to indicate the status of the ASR creation. Once they are created, ASRs are identical to other service requests and can be viewed and managed using the same processes and tools.

A user with the Ops Center Admin role must enable the ASR feature. See the *Oracle Enterprise Manager Ops Center Administration Guide* for how to configure and use automated service requests.

Related Resources for Incident Management

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

- See [Chapter 4, "Monitoring Rules and Policies"](#) for information about the monitoring rules that trigger incidents.
- See [Appendix E, "Icons"](#) for a description of the icons.
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm. The following How To documents relate to monitoring:
 - Managing Incidents
 - Using Service Requests

This chapter includes the following information:

- [Introduction to Reports](#)
- [Roles for Reports](#)
- [Actions for Reports](#)
- [Location of Report Information in the User Interface](#)
- [Creating a Firmware Report](#)
- [Creating an Operating System Report](#)
- [Creating Additional Operating System Reports](#)
- [Creating Incident Reports](#)
- [Creating System Information Reports](#)
- [Creating Exalogic Reports](#)
- [Related Resources for Reports](#)

Introduction to Reports

The Reports of Oracle Enterprise Manager Ops Center provides information about assets, such as job history, firmware, operating system updates, and incidents. Reports are exported or used to start jobs on targeted assets. You can schedule to run the report using any of the following parameters:

- **Now:** Select the current date and time to generate the report.
- **At a later date/time:** Select a date and time to generate the report.
- **On a Recurring Schedule:** Select the month and day when you want to generate the report. Select the Start Time, End Time, and Number of Hours between runs. This is to set the number of times the report is generated between the specified start and end time. For example, when you set the start time at 6.00 a.m, end time at 12.00 a.m, and the number of hours between runs as 2, then the report is run at 6.00 a.m, 8.00 a.m, 10.00 a.m, and 12.00 a.m.

After the report is generated, the Report Results pane lets you to export the report result in CSV and PDF formats, view the report interactively, and delete the report. View Interactive option helps you to view the generated report in detail. Report result displays the Report name, Report type, Run Date, Targets of the Report, OS updates applicable to selected targets. Report parameters show the target name, product name, description. You can save the report as a template and also rerun the report.

Figure 10–1 displays an Interactive view of Report Result.

Figure 10–1 Interactive View of Report Result

Oracle Enterprise Manager Ops Center - Report Details

Report Name: **test2** Report Type: **Change History Report**
 Run Date: **03/20/2012 09:35:06 am UTC** Status: **Success**
 Description: **hjk**

Report Result **Report Parameters**

The following targets were run with the change history report

▼ **Targets of the Report (0)**

Target Name	Number of Changes
No data	

The following detailed change history information applies to the specified target(s).

▼ **OS Updates Applicable to Selected targets (0)**

Host Name	User Name	Job Name	Action	Component	Distribution	Date	Time	Target Boot Env.
No data								

Roles for Reports

The following table lists the tasks and the role required to complete the task. Contact your administrator when you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant. Table 10–1 illustrates the roles and permissions for reports.

Table 10–1 Reports Tasks and Roles

Tasks	Roles
View Reports	All
Create Reports	Asset Administrator

Actions for Reports

You can perform the following actions in the reports section.

- [Managing Report Templates](#)
- [Viewing and Exporting Report Results](#)

Managing Report Templates

A report template is a pre-formatted file that serves as a starting point to create a new report. When you save a file created from a template, you are prompted to save a copy of the file so that you do not over-write the template. Templates are provided within a software or a program or it is created by the user. Most major software support templates. If you want to create a similar document or report over and over again, it is a good idea to save one of them as a template. You can open the report template and start creating reports from there. Parameters in the report template are specified when the report is created or run. You can save the criteria as a template, after creating the report criteria. The saved report template enables us in:

- [Generating a Report from a Report Template](#)
- [Editing a Report Template](#)
- [Deleting a Report Template](#)

Generating a Report from a Report Template

To generate a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report template from the center pane.
3. Click the **Generate Report** icon to run the report. The corresponding report is generated and the results are displayed under Report Results.

Editing a Report Template

To edit a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report template to edit from the center pane.
3. Click the **Edit View** icon to edit the selected report template. The corresponding report wizard is displayed.
4. Edit the report parameters as required in the wizard.
5. Click **Run and Close** to run the report or click **Save Template and Close** to save the report template. When you click **Run and Close**, the report is generated but the edits for the report are not saved.

Deleting a Report Template

To delete a report from the report template, do the following:

1. Select **Reports** from the Navigation pane.
2. Select a saved report template to delete from the center pane.
3. Click the **Delete** icon to delete the report template.
4. Click **Ok** to confirm the delete action. The selected report template is deleted.

Viewing and Exporting Report Results

Oracle Enterprise Manager Ops Center provides an interactive result viewer to view the results. The generated report results are displayed under the Report Results in the All Reports page. You can select a report result from the Report Results pane to view the report interactively. You can rerun the report, delete the report, view, export, and save the output of report in the format of CSV and PDF.

You can do the following operations:

- [Viewing the Report Result](#)
- [Saving the Report Result](#)

Viewing the Report Result

To view the report result, do the following:

1. Select **Reports** in the Navigation pane. The All Reports page in the center pane displays all the report templates and report results.

2. Select a report result under the Report Result section in the center pane.
3. Click the **Interactive** icon to view the report result. The **Interactive Result viewer** opens the report result and displays the following information.
 - Report detail: This displays the name, type, run time, and the status of the report.
 - Report Result: This displays the targets on which the report is run and the corresponding operating system updates that are applicable.
 - Report Parameters: This displays the parameters that are used to generate the report.

Saving the Report Result

To save the report result, do the following:

1. Select **Reports** in the Navigation pane. The All Reports page in the center pane displays all the report templates and report results.
2. Select a report result under the Report Result section in the center pane.
3. Click the **View CSV** or **View PDF** icon. You can save or open the report in CSV or PDF formats.

Figure 10–2 illustrates the report output in PDF format.

Figure 10–2 Exporting a Report Result in PDF

Report Name:	test1		
Description:	test1		
Run Date:	Mon Mar 19 20:40:17 MDT 2012		
Report Type:	Change History Report		
<table border="1"> <thead> <tr> <th>Target Name</th><th>Number of Changes</th></tr> </thead> </table>		Target Name	Number of Changes
Target Name	Number of Changes		

Location of Report Information in the User Interface

To see report information, expand Reports in the Navigation pane. This displays OS reports, System Information reports, Incident reports, and Firmware reports. Click OS reports to view the various types of OS reports that run on Oracle Solaris/Linux OS updates and Windows OS updates. Click Incident reports to view the different types of incident reports. Table 10–2 illustrates the location of reports in the user interface.

Table 10–2 Location of Report Information in the UI

Object	Location
OS Reports	Expand the Reports in the Navigation pane. Click OS Reports to view the various reports that run on Oracle Solaris/Linux OS updates and Windows OS updates.
System Information Reports	Expand the Reports in the Navigation pane. Click System Information Reports. This displays the Create System Information Report wizard in the Actions pane.

Table 10–2 (Cont.) Location of Report Information in the UI

Object	Location
Incident Reports	Expand the Reports in the Navigation pane. Click Incident Reports. This displays the types of incident reports.
Firmware Reports	Expand the Reports in the Navigation pane. Click Firmware Reports. This displays the Create Firmware Report wizard in the Actions pane.
Exalogic Reports	Expand the Reports in the Navigation pane. Click Exalogic Reports. This displays the Create Exalogic Report wizard in the Actions pane.

Creating a Firmware Report

Firmware Compliance Reports enables us to maintain consistent firmware versions across the datacenter. The Firmware Report feature compares the firmware images specified in a firmware profile to the firmware images installed on hardware assets. The report indicates whether the firmware on the asset complies with the profile's specifications. You can update the firmware on any non-compliant asset by clicking the Make Targets Compliant button in the Interactive report.

Note: The Firmware Reports action is disabled for Exalogic systems.

To Create a Firmware Report

1. Select **Reports** in the Navigation pane.
2. Select **Firmware Reports**.
3. Select **Create Firmware Report** from the Actions pane. The Create Firmware Report Wizard is displayed.
4. Define the report parameters, including:
 - **Report Name:** The name of the report.
 - **Description:** A description of the report.
 - **Schedule:** If you do not plan to create this report routinely, deselect the **Create Schedule** option.
 - **Output Format:** Select the output format of the report result. CSV and PDF formats are available.
 - **Profile:** Select the firmware profile for a service processor or for a storage component, such as a RAID controller, expander, or disk.
5. Click **Next** to display the Select Targets.
6. Select the targets you want to test against the profile. Select the asset from the Available Items hierarchy and click **Add to Target List**. When you have selected all the targets, click **Next** to display the Schedule page.
7. Select a desired schedule to run and generate the report.
8. Click **Next** to display the Summary.
9. Review the summary and click **Run Report** to create the report job.

The report job starts at the time you specified and compares the values in the profile to the existing values on the targets you selected. The report displays whether a target asset is compliant, not compliant, or not applicable:

- A compliant asset has the firmware images specified in the profile.

- A non-compliant asset does not have the same firmware images as specified in the profile. You can update the firmware on the asset by either clicking the Make Targets Compliant button in the Interactive report or using the procedure in Updating Firmware.
- A non-applicable asset indicates that a firmware image in the profile does not match the model of service processor in the asset. This condition can occur when either the profile does not recognize the model that the service processor is reporting or the profile includes firmware images that are not designed for the service processor.
 - a. Compare the model of the service processor displayed in the asset's Summary tab with the model of the service processor included in the profile. If they are different, add the name in the profile to the asset's data.
 See *Oracle Enterprise Manager Ops Center Administration Guide* for information about adding a product alias.
 - b. When the firmware profile was created, only images that matched the service processor could be included. However, if the service processor did not report all the firmware types it supported, an image that did not match the service processor could have been included in the profile. To update the software with all the service processor's supported firmware types, use the **Refresh** action to update the information about the service processor. When the job is completed, view the service processor's **Summary** tab to see all firmware types.
 - c. Repeat the procedure to create a firmware report.

Creating an Operating System Report

Use Operating System update reports to check for new software updates and security advisories. For auditing purposes, create a change history report. Various types of update reports are available for Linux, Oracle Solaris, and Windows operating systems. You can export report results to CSV or PDF format.

Use operating system reports to obtain information about managed Oracle Solaris, Linux, and Windows operating systems.

Table 10–3 illustrates the various reports that are run in Oracle Enterprise Manager Ops Center and displays the type of report that are supported and run on an operating system.

Table 10–3 Compatibility of Reports on Operating System

Report Name	Linux OS	Oracle Solaris 8, 9, and 10	Oracle Solaris 11	Microsoft Windows
Change History Report	Yes	Yes	Yes	No
CVE Compliance Report	Yes	Yes	No	No
RSC Report	Yes	Yes	No	No
System Catalog Report	Yes	Yes	Yes	No
Oracle Solaris Update Compliance Report	No	Yes	No	No
Baseline Analysis Report	No	Yes	No	No
Update Compliance Report	Yes	Yes	No	No

Table 10–3 (Cont.) Compatibility of Reports on Operating System

Report Name	Linux OS	Oracle Solaris 8, 9, and 10	Oracle Solaris 11	Microsoft Windows
Incident Compliance Report	Yes	Yes	No	Yes
Host Compliance Report	Yes	Yes	No	Yes
Distribution Update Report	Yes	Yes	No	No
Service Pack Compliance Report	Yes	No	No	No
Package Compliance Report	No	Yes	No	No

Update Compliance Reports

The following reports are available for Linux, Oracle Solaris, and Windows operating systems:

- **Host Compliance:** Provides information on whether your system is compliant with security and bug fixes incidents.
- **Incidence Compliance:** Provides information about the number of systems to which the selected operating system updates apply.

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 Operating System update, install, and uninstall jobs completed on managed systems.
- **CVE Compliance:** Provides information on incidents that are related to specific Common Vulnerability and Exposure Identifiers (CVE IDs) and the systems that have these incidents installed. CVE IDs are unique, common identifiers for publicly known security vulnerabilities.
- **Distribution Update:** Provides a mapping between selected updates, CVE's, and selected distributions to find out whether the updates are installed.
- **Package Compliance:** Provides the details of the selected packages on managed system that are compliant or not compliant with the latest recommended version available.
- **RSC: Recommended Software Configuration –** Provides information about the system compliance for installing a specific application, such as the Oracle 11g Database, on an Oracle Solaris, or Linux Operating System.
- **Service Pack Compliance (Linux only):** Provides information on incidents created by the publication and release of a service pack by a vendor. This helps in determining whether the system has the latest service packs released by the vendor.
- **Oracle Solaris Update Compliance (Oracle Solaris Operating System only):** Provides information on whether an Oracle Solaris system is compliant with a specific update.
- **Baseline Analysis (Oracle Solaris Operating System only):** This helps to check the compliance of systems against newly released Oracle Solaris baselines.

Creating a Change History Report

The Change History Report provides a history of operating system update install and uninstall jobs completed on managed Oracle Solaris or Linux systems. The report also displays the deployments made by the specific user, enabling you to track a team of operators.

To Create a Change History Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Change History Report** from the Actions pane.

The Create Change History Report Wizard is displayed.

Figure 10–3 Defining Report Parameters for Change History Report

Oracle Enterprise Manager Ops Center - Create Change History Report

Create Change History Report

Steps: 1. Define Report Parameters, 2. Report Schedule, 3. Summary

Define Report Parameters

* Report Name: Sample

Description:

Date Range: ☒ Specify Date Range

Start Date: 11/11/2010

End Date: 11/11/2010

Actions: ☒ Install, ☒ Uninstall

Schedule: ☒ Create Schedule

Output Format: ☒ CSV, ☒ PDF

Select Targets

Assets	Product Name	Description
All Assets		
Servers		
cosoqa09	Sun Fire B100x B1a Server	System
cosoqa09	Red Hat Enterprise Operati	ServerP
cosoqa14	Sun Fire B100x B1a Server	
cosoqa14	Solaris 10 Operatin GlobalZ	
cosoqa15	Sun Fire B100x B1a Server	

Target List(1): cosoqa14

Buttons: Add to Target List, Remove from Target List, Previous, Next, Cancel

4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Date Range: Specify the start date and end date between which the report will cover.
 - Actions: Select the actions to be reported. You can select Install, Uninstall or both.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to schedule the report.
6. Select a desired schedule to run and generate the report.

7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - Run and Close: Runs the report and closes the wizard window.

The report results are displayed under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result.

Creating a Baseline Analysis Report

An Oracle Solaris baseline is a dated collection of Oracle Solaris updates, update metadata, and tools. Oracle releases Oracle Solaris baselines on a monthly basis. A Baseline Analysis Report checks the compliance of Oracle Solaris systems against newly released Oracle Solaris baseline. When you install the updates of a baseline on a host, that system is considered to be compliant with that baseline.

Each dated baseline contains these update sets:

- Full: Includes all Oracle Solaris updates
- Recommended: Includes Oracle Solaris recommended updates and security updates
- Security: Includes only Oracle Solaris security updates

All baselines include updates for a specific time frame. However, the Full baseline often contains Oracle Solaris operating system updates that are not included in the Recommended baseline. The Full baseline includes additional updates based on feedback from various customer support groups within Oracle. Recommended baseline does not includes these updates.

To install the Recommended and Security baselines, you must either deploy two jobs or have a job that includes multiple tasks. This might result in multiple reboots, for example, if both tasks (baselines) include updates that have Single User mode requirements.

Oracle Enterprise Manager Ops Center's Knowledge Base (KB) is updated with the information about the baselines. This is done a few days after the official release of baselines by Oracle.

Note: The Oracle Solaris 8 Operating System was placed into End of Service Live (EOSL) on March 31, 2009. Oracle Solaris 8 Operating System baselines are available through March 2009. The KB might contain artificial baselines after that date. Do not use baselines dated after March 2009.

Oracle Solaris baselines enables you to easily identify the update level of your hosts. For example, install some test hosts with a particular baseline. Test these hosts for a period to see whether the updates in this baseline are stable enough to be used on production hosts. When the testing reveals that this baseline is stable, install the same baseline on production hosts.

Oracle Solaris baselines are available as a component in the recommended component list. This contains a list of dated baselines.

The Baseline Analysis report helps to verify the compliance of your system against the newly-released baselines (as and when they are available in Knowledge Base).

The Baseline Analysis Report (BAR) enables you to determine whether the managed system is compliant with recently released Oracle Solaris baselines. Baselines pertain only to Oracle Solaris systems. This section describes Oracle Solaris baselines, white list, black list, and how to run a Baseline Analysis report in connected and disconnected mode of the Enterprise Controller.

The Baseline Analysis Report (BAR) describes how to generate a BAR. The report gives the compliance status of the managed system with the selected Oracle Solaris baseline that was released.

You can generate two types of BARs:

- Agent-based BAR
- Database-based BAR

In an agent-based BAR, a simulated job is run against the managed hosts. This type of report takes time to complete because it checks for dependent components and missing dependencies, and then downloads the updates that must be installed. When you run a compliance job from this report result, the job is completed quickly because the updates are downloaded. However, to improve the report performance of a BAR, skip the downloads in a simulated job by deselecting this option.

In a database-based BAR, the report is run against the database of the management server, the selected baselines are broken down into individual update IDs, and then formed into an incidents list. The report is generated based on the information that are available on the database. Based on the report result, run a compliance job.

White List

A white list is the list of updates that is required to install in addition to the updates in the baseline. To establish a white list, create a profile using the Required setting. You can also specify a white list when generating a Baseline Analysis Report. Select the white list either from the created profile or enter the update IDs separated by new lines.

For example, baseline B includes updates X, Y, and Z, and the white list has updates U, V, and W. When the Baseline Analysis report is created, the host is marked compliant only when all six updates X, Y, Z, U, V, and W are present.

Black List

A black list is a list of updates that you do not want to install. Create a black list by creating a policy with the specified action for the updates. Select a black list option while creating a Baseline Analysis Report. Select the black list either from the created policy or enter the update IDs separated by new lines.

If a particular update in the profile is set with the policy component setting as Never for an install action, then the update is not installed. If the update is installed, it will not be uninstalled or removed.

For example, if baseline A has updates X, Y, and Z, and the black list specifies only Y and Z, the system is compliant if X is installed. If the updates Y and Z are installed, they will not get uninstalled if you run a compliance job from the report results. If Y and Z are not installed, they are not listed in the non compliant result and are not added in the compliance job.

To Create a Baseline Analysis Report

This report provides information about the hosts that are compliant with a baseline operating system.

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Baseline Report** from the Actions pane. The Create Baseline Analysis Report Wizard is displayed.
4. Define the report parameters:
 - Report Name: Name of the report.
 - Description: The description of the report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the Oracle Solaris baselines.
6. In Select Baseline(s), select the following options:

Figure 10–4 Selecting Baselines for Baseline Analysis Report

Oracle Enterprise Manager Ops Center - Report Details

Report Name: **BAR1** Report Type: **Baseline Analysis Report**
 Run Date: **01/10/2011 01:19:18 pm IST** Status: **Successful But Not Compliant**

Description: **Report Result** Report Parameters

The following targets were run with the Baseline Analysis Report

Targets of the Report (1)

Baseline	Target	OS	No of Changes	Status
Baseline 2010/12/08 [Full]	ipgs14	SOLARIS_10_0_X86	161	NONCOMPLIANT

The following detailed Baseline Analysis information applies to the specified target(s).

OS Updates Applicable to Selected targets (161)

Target	Baseline	Action	Incident	Synopsis	Link
ipgs14	Baseline 2010/12/08 [Full]	Install	125389-03	SunOS 5.10_x86: SNIA	http://sunsolve.sun.com/
ipgs14	Baseline 2010/12/08 [Full]	Install	125216-03	SunOS 5.10_x86: wget	http://sunsolve.sun.com/
ipgs14	Baseline 2010/12/08 [Full]	Install	122260-03	SunOS 5.10_x86: SunFr	http://sunsolve.sun.com/
ipgs14	Baseline 2010/12/08 [Full]	Install	142235-01	SunOS 5.10_x86: ntp.xr	http://sunsolve.sun.com/
ipgs14	Baseline 2010/12/08 [Full]	Install	119118-52	Evolution 1.4.6_x86 patc	http://sunsolve.sun.com/

Make Targets Compliant Rerun Report Close

- Run Against Database or Run Report Against Agent.

When you select **Run Report Against Agent**, check the **Download** check box to download the updates that are installed on the target.

- Select the distribution type and select the baselines from the list. You can select targets of multiple distribution. For each distribution, select the corresponding baselines. A warning message is displayed when the baselines are not selected for a distribution.

Note: If you have multiple distributions, then you must select baselines for at least one distribution to continue further in the wizard. If you have not selected baselines for a distribution, then the targets of that distribution are not in the report result.

- Click **Add** or **Add All** to select all the baselines.
7. Click **Next** to modify the update lists that are applied to the report.
 8. Select any of the following White List options:
 - None: No white list.
 - Manual Input: Enter a list of updates.
 - Specify with Profile: Select a profile to import as a white list.
 9. Select any of the following Black List options:
 - None: No black list.
 - Manual Input: Enter a list of updates.
 - Specify with Policy: Select a policy to import as a black list.
 10. Click **Next** to schedule the report.
 11. Select a desired schedule to run and generate the report.
 12. Click **Next** to display the Summary.
 13. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate reports later.
 - Run and Close: Runs the report and closes the wizard window.

The report result displays under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating a Profile Analysis Report

A Profile Analysis Report provides information about Oracle Solaris or Linux systems' compliance with the Operating System Update Profiles that you define in Oracle Enterprise Manager Ops Center. The update profiles include both the system-defined and user-defined profiles in Oracle Enterprise Manager Ops Center.

Note: Avoid running reports for system-defined profiles like Perform Reboot+Reconfigure and Perform Reboot as these profiles do not contain any updates.

Before You Begin

You can modify the update list that is applied to generate the report by selecting a white list and a black list.

A white list is the list of updates to install. To establish a white list, create a profile using the required setting. Select the white list either from the created profile or enter the update IDs separated by new lines.

For example, baseline B includes updates X, Y, and Z, and the white list has updates U, V, and W. When the Baseline Analysis Report is created, the host is marked compliant only when all six updates (X, Y, Z, U, V, and W) are present.

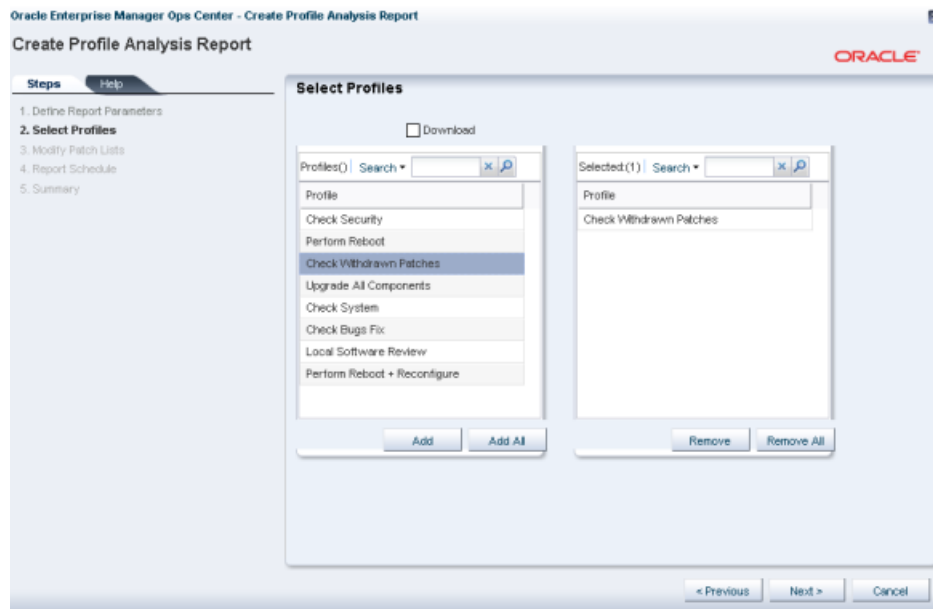
A black list is a list of updates that you do not want them to be installed. You create a black list by creating a policy with the specified action for the updates. Select the black list either from the created policy or enter the update IDs separated by new lines.

When a particular update in the profile is set with the policy component setting as Never for the install action, then the update is not installed. When the update is installed, it is not uninstalled or removed.

For example, when baseline A has updates X, Y, and Z, and the black list specifies only Y and Z, the system is compliant when X is installed. Even if the updates Y and Z are installed, they are not uninstalled when you run a compliance job from the report results.

To Create a Profile Analysis Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Profile Report** from the Actions pane. The Create Profile Report Wizard is displayed.
4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the profiles.
6. Select the profiles from the list and click **Add** or **Add All** to select all the available profiles.

Figure 10–5 Selecting Profiles for Profile Analysis Report

7. Check the Download check box to download the updates that must be installed for the system compliance.
8. Click **Next** to modify the update lists that are applied to the report.
9. Select any of the following White List options:
 - None: No white list.
 - Manual Input: Enter a list of updates.
 - Specify with Profile: Select a profile to import as a white list.
10. Select any of the following black list options:
 - None: No black list.
 - Manual Input: Enter a list of updates.
 - Specify with Policy: Select a policy to import as a black list.
11. Click **Next** to schedule the report.
12. Select a desired schedule to run and generate the report.
13. Click **Next** to display the Summary.
14. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - Run and Close: Runs the report and closes the wizard window.

The report result displays under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating a Recommended Software Configuration Report

A Recommended Software Configuration provides information about the system compliance for installing a specific application, such as the Oracle 11g Database, on an Oracle Solaris or Linux operating system.

The Knowledge Base provides a list of application configuration requirements with which you can check your system compliance status.

For example, you can check the system compliance status of Oracle Solaris operating system for installing Oracle 11g Database. The report provides information about the updates that must be installed, uninstalled, or upgraded for installing the Oracle database.

For an Oracle Solaris operating system, you cannot upgrade a update component from the existing lower version to the recommended higher version. Such instances will be marked as Error in the RSC report result. In such scenarios, you cannot make the target system fully compliant with the recommended software components by the report.

Before You Begin

You can generate different types of RSCs:

- Agent-based RSC
- Database-based RSC

In an agent-based RSC, the report is generated based on the information from the target system. The dependencies for the updates are checked and downloaded when required. This report takes time to generate because it checks dependencies and downloads updates that must be installed.

In a database-based RSC, the report is generated based on the target system information that is available on the database of the Enterprise Controller. The dependencies are not checked and required updates are not downloaded. This type of report is generated quickly.

To Create a Recommended Software Configuration Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Recommended Software Configuration Report** from the Actions pane. The Create Recommended Software Configuration Report Wizard is displayed.
4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the recommended software configurations.

6. In **Select Recommended Software Configurations**, select any of the following options:
 - **Run Against Database or Run Report Against Agent.**
When you select **Run Report Against Agent**, then click the **Download check box** to download the updates that must be installed on the target.
 - **Select the Distribution type.**
 - Select the recommended software component from the list and select the required configuration. The recommended configuration describes the prerequisite list of updates for the selected application. You can select targets of multiple distribution. For each distribution, select the corresponding RSCs. A warning message is displayed when you have not selected RSCs for a distribution.

Note: When you have multiple distributions, then you must select RSCs for at least one distribution to continue further in the wizard. When you have not selected RSCs for a distribution, then the targets of that distribution are not in the report result.

7. Click **Next** to schedule the report.
8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.
10. Review the report parameters and select one of the option as required:
 - **Save Template and Close:** Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - **Run and Close:** Runs the report and closes the wizard window.

The report result displays under the **Report Results** in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating an Oracle Solaris Update Compliance Report

The Oracle Solaris Update Compliance report determines whether a specific Oracle Solaris system is compliant with a particular released Update.

To Create an Oracle Solaris Update Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select **Additional Reports** from the Reports section.
3. Select **Solaris Update Compliance** from the Actions pane. The Solaris Update Compliance Report Wizard is displayed.
4. Specify the report parameters:
 - **Name:** The name of the report.
 - **Description:** A description of the report.
 - **Schedule:** Select **Create Schedule** to schedule the report.
 - **Output Format:** Select the output format of the report result. CSV and PDF formats are available.

- Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
- 5. Click **Next** to select the target asset.
The Select Targets page is displayed.
- 6. Add the targets by selecting them from the list on the left by clicking **Add to Target List**. Click **Next** to display the Summary page.
- 7. Click **Save Report** to save the report for future use. This returns you to the Reports tab, where you can run the report by selecting it from the Saved Reports section and clicking **Re-run Report**.
- 8. Click **Run Report** to run and display the report.
- 9. Click **Export to CSV** to export the report result.
- 10. Click **Done** to close the report.

Creating an Incident Compliance Report

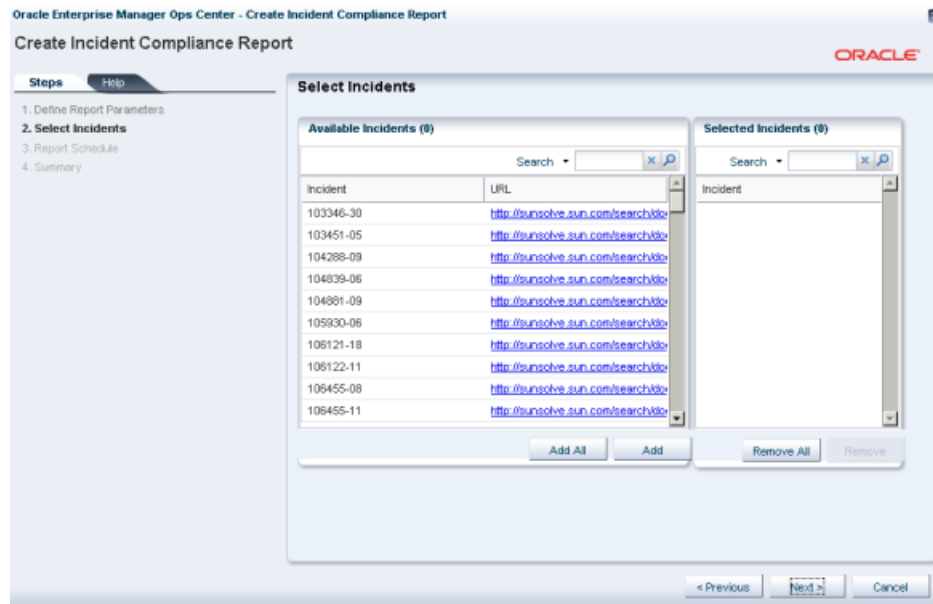
Incidents are the updates that are available for an application or feature. Incidents apply to one or more packages or RPMs. You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

Incident Compliance Report for Oracle Solaris or Linux

You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

To Create an Incident Compliance Report for Oracle Solaris or Linux

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Incident Report** from the Actions pane. The Create Incident Compliance Report Wizard is displayed.
4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Criteria: You can select either Select Updates or Filter Updates, for selecting the updates that are used as a comparison. Depending on the selection of criteria the wizard steps vary.
 - Compliant: Select either **Compliant** or **Non-compliant** for compliance status.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the updates.
6. When you have selected Select Updates in the previous step, the list of available incidents is displayed.

Figure 10–6 Selecting Incidents for Incident Compliance Report

7. Select the incidents and click **Add** or **Add All** to select all the listed incidents.
8. If you have selected Filter Updates in the first step, then select the following:
 - Select Packages: You can select the updates based on the category, update type and releases date. Select the packages and click **Add** or **Add All** to select all the packages in the Available Packages list. Click **Next** to select the CAN IDs.
 - Select CAN IDs: Select from the list of Available CAN IDs. Click **Add** or **Add All** as required.
9. Click **Next** to schedule the report.
10. Select a desired schedule to run and generate the report.
11. Click **Next** to display the Summary.
12. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - Run and Close: Runs the report and closes the wizard window.

The report result displays under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating an Incident Compliance Report for Microsoft Windows

You can run an incident compliance report to determine whether the incidents on the managed hosts are compliant with the latest released version.

The Incident Compliance Report provides information about whether your systems are compliant with the Windows updates incidents. This report displays the number of systems to which the selected Windows updates apply, how many systems have the updates installed, and how many systems require the updates to be installed to make the systems compliant. You can create a Windows update job based on the results of an Incident Compliance Report.

To Create an Incident Compliance Report for Windows

1. Select **Reports** from the Navigation pane.
 2. Select Windows Incident Compliance Report from the Actions pane. The Windows Incident Compliance Report Wizard is displayed.
 3. Specify the report parameters. They include:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Specify the Windows OS updates on which to run the report – You can specify filter criteria such as Category, Severity, Superseded, and Release Date for Windows OS updates, or you can select specific Windows OS updates to run the report.

Click **Next**.
 4. Based on your selection in Step 3, either the Define Updates Filter window is displayed or the Select Updates window is displayed. When the Define Updates Filter window is displayed, go to Step 6. When the Select Updates window is displayed, go to Step 7.
 5. Make your selections in the Define Updates Filter screen. They include:
 - Category: Includes Application, Critical Updates, Definition Updates, Drivers, Service Packs, Security Updates, Tools, Update Rollups, and WSUS Infrastructure Updates. You can select either All available updates under all category or Selected categories only. Use the Control key on the keyboard to select multiple items in the list under Selected category only.
 - Severity: Includes Critical, Important, Moderate, Low, and Default. You can select either All updates with any severity or Selected severities only. Use the Ctrl key on the keyboard to select multiple items in the list under Severity.
 - Superseded: Enables you to select all or just the most recent updates.
 - Release Date: Refers to the date that the update updates were released. You can select the range of release dates to include in your report by filling in the From and To fields. Click **Next**. Go to Step 8.
 6. Make your selections in the **Select Updates** window. Under Search, Select All enables you to include a bulletin ID, article ID, and title in your search, or you can select specific fields to narrow your search. Use the Control key on the keyboard to make multiple selections in the list under Available Windows Software Updates.
- Click **Add to Updates List**, and then click **Next** to select the targets.
7. Add the targets by selecting them in the list of Available Items, by clicking **Add to Target List**.
- Click **Next** to display the Summary page.
8. Click **Finish** to run the report.
- The results of the report are displayed under the Report Results list.

Creating a Host Compliance Report

You can run a host compliance report to determine whether the hosts are complaint with security and bug fix incidents. This report displays the number of updates that are applicable to each system, and whether the updates are installed or must be

installed to make the system compliant. You can also create an update job based on the results of a Host Compliance Report.

Host Compliance Report for Oracle Solaris or Linux

The Host Compliance Report provides information if your systems are compliant with update incidents.

To Create a Host Compliance Report for Oracle Solaris or Linux

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create Host Compliance Report** from the Actions pane.
The Create Host Compliance Report Wizard is displayed.
4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Update Level: Select whether you want the compliant status for Security and Bug Fixes or for only Security Updates.
 - Compliance: Select either **Compliant** or **Non-Compliant**.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them in the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to schedule the report.
6. Select a desired schedule to run and generate the report.
7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - Run and Close: Runs the report and closes the wizard window.

The report result displays under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating Host Compliance Report for Microsoft Windows

The Host Compliance Report for Windows provides information if your systems are compliant with the Windows updates incidents. This report displays the number of Windows updates that are applicable to each system, and whether the updates are installed or must be installed to make the system compliant. You can also create a Windows update job based on the results of a Host Compliance Report.

To Create a Host Compliance Report for Windows

1. Select **Reports** from the Navigation pane.

2. Select **Windows Host Compliance Report** from the Actions pane. The Windows Host Compliance Report Wizard is displayed.
3. Specify the report parameters. They include:
 - Report Name: A name for the report.
 - Description: A description of the report.
 - Specify the Windows OS updates on which to run the report. You can specify filter criteria such as Category, Severity, Superseded, and Release Date for Windows OS updates, or you can select specific Windows OS updates to run the report.
4. Click **Next**. Based on your selection in Step 3, either the Define Updates Filter window is displayed or the Select Updates window is displayed. When the Define Updates Filter window is displayed, go to Step 5. When the Select Updates window is displayed, proceed to Step 6.
5. Make your selections in the Define Updates Filter screen. They include:
 - Category: Includes Application, Critical Updates, Definition Updates, Drivers, Service Packs, Security Updates, Tools, Update Rollups, and WSUS Infrastructure Updates. You can select either All available updates under all category or Selected categories only. Use the Control key on the keyboard to select multiple items in the list under Selected category only.
 - Severity: Includes Critical, Important, Moderate, Low, and Default. You can select either All updates with any severity or Selected severities only. Use the Ctrl key on the keyboard to select multiple items in the list under Severity.
 - Superseded: Enables you to select all or just the most recent updates.
 - Release Date: Refers to the date that the updates were released. You can select the range of release dates to include in your report by filling in the From and To fields. Click **Next**. Go to Step 8.
6. Make your selections in the Select Updates window. Under Search, Select All enables you to include a bulletin ID, article ID, and title in your search, or you can select specific fields to narrow your search. Use the Control key on the keyboard to make multiple selections in the list under Available Windows Software Updates. Click **Add to Updates List**. Click **Next**.
7. Add the targets by selecting them from the list of Available Items. Click **Add to Target List**. Click **Next** to display the Summary page.
8. Click **Finish** to run the report.

The results of the report are displayed under Report Results list.

Creating a Common Vulnerability and Exposure (CVE) Report

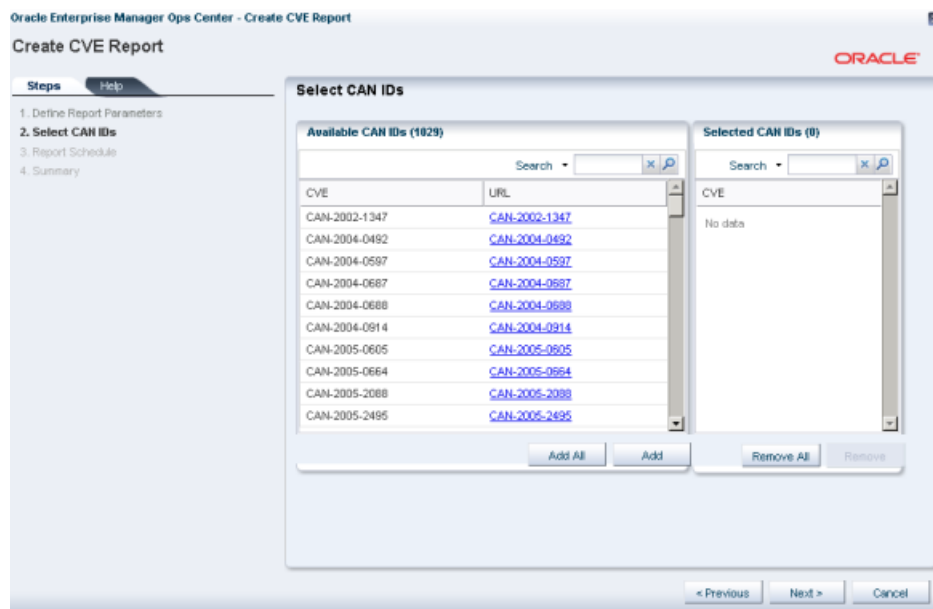
The CVE report provides information about incidents that are related to specific Common Vulnerability and Exposure Identifiers (CVE IDs) and the systems that must have these incidents installed. CVE IDs are unique, common identifiers for publicly known security vulnerabilities. The updates from a list of vendors are published as common vulnerabilities and security exposure incidents. CVEs are identified by a candidate ID (CAN ID).

To Create a CVE Report

1. Select **Reports** from the Navigation pane.

2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Create CVE Report** from the Actions pane. The Create CVE Report Wizard is displayed.
4. Define the report parameters. They include:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Compliance: Select either **compliant** or **non-compliant** report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
5. Click **Next** to select the CAN IDs.
6. Select one or more CAN IDs and click **Add** or **Add All** to select all the available CAN IDs.

Figure 10–7 Selecting CAN IDs in CVE Report



7. Click **Next** to schedule the report.
8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.
10. Review the report parameters and select one of the options as required:
 - Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
 - Run and Close: Runs the report and closes the wizard window.

The report result is displayed under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating a System Catalog Report

A System Catalog Report lists the current catalog of one or more systems. A system catalog contains a list of operating system software components that are installed on a managed system. Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems.

After an operating system is available and selected, you can view and modify the catalogs, and create historical catalogs. Historical catalogs are snapshots of the system. The software automatically takes a snapshot of the operating system after running a job on the operating system, including when you discover and manage the operating system. A snapshot is stored as a catalog with the time stamp and job details after every update job that you run on a system.

You can create a new catalog at any time and use it to record the state of a system. Catalogs enables us to rollback our system to any previous configuration or to create a profile that is used to apply a consistent configuration throughout our datacenter.

See [Chapter 13, "Operating System Updates"](#) for more information on System Catalogs.

To Create a System Catalog Report

1. Select **Reports** from the Navigation pane.
2. Select **System Catalog Report** from the Actions pane.
The System Catalog Report Wizard is displayed.
3. Define the report parameters, including:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Targets: Add one or more targets by selecting them in the list of Available Items by clicking **Add to Target List**.
4. Click **Next** to display the Schedule.
5. Select a desired schedule to run and generate the report.
6. Click **Next** to display the Summary.
7. Review the Summary, then click **Run and Close**.

The report result is displayed under the Report Results in the center pane.

See [Viewing and Exporting Report Results](#) for more information about viewing a report result and generating a compliance job from the result.

Creating Additional Operating System Reports

Use Additional Operating System Reports to obtain information from Service Pack Compliance Report, Distribution Update Report, and Package Compliance Report. You can export report results to CSV format.

Creating a Distribution Update Report

A Distribution Update Report provides a mapping between selected updates, and CVEs and selected distributions to find out whether the updates are installed. This report determines whether a specific distribution (E.G. SOLARIS10_SPARC) has been updated with specific updates, or CVEs.

To Create a Distribution Update Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Distribution Update Report** from the Actions pane. The Distribution Update Report Wizard is displayed.
4. Define the report parameters. They include:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Category: Select the required categories.
 - Type: Select required types. Package and Update types are available.
 - Released: Mention the start date and the end date.
5. Click **Next** to display the Distributions.
6. Select the required Distribution and click **Add** or **Add All** to add distributions. Click **Remove** or **Remove All** when you do not require any distribution.
7. Click **Next** to select the updates.
8. Click **Next** to select the packages.
9. Click **Next** to select the CVEs. Click **Add** or **Add All** to add CVEs and **Remove** or **Remove All** when you do not require any CVE.
10. Click **Next** to display the Summary.
11. Review the Summary and select one of the options as required:
 - Save Report: Saves the report as a template and closes the wizard.
 - Run Report: Runs the report and closes the wizard window.

The report result is displayed under the Report Results in the center pane.

Creating a Service Pack Compliance Report

A Service Pack Compliance Report provides information on updates created by the publication and release of a service pack by a vendor. This report enables you to determine whether the target system has the latest service package installed that is provided by the vendors.

To Create a Service Pack Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Service Pack Compliance Report Creation** Wizard from the Actions pane. The Service Pack Compliance Report Creation wizard is displayed.
4. Define the report parameters. They include:

- Report Name: The name of the report.
 - Description: A description of the report.
 - Status: Select the compliant or non compliant status.
 - Services: Select the required services.
5. Click **Next** to select the Targets. Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
 6. Click **Next** to display the summary.
 7. Review the Summary and select one of the options as required:
 - Save Report: Saves the report as a template and closes the wizard.
 - Run Report: Runs the report and closes the wizard window.
- The report result is displayed under the Report Results in the center pane.

Creating a Package Compliance Report

A Package Compliance Report provides a mapping between the selected packages and the selected target systems to find out the installed packages.

To Create a Package Compliance Report

1. Select **Reports** from the Navigation pane.
2. Select Oracle Solaris/Linux OS Updates from the Reports section.
3. Select **Package Compliance Report Creation** Wizard from the Actions pane. The Package Compliance Report Creation Wizard is displayed.
4. Define the report parameters. They include:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Status: Select the compliant or non compliant status.
 - Level: Select security updates or security and bug updates.
5. Click **Next** to select the Targets. Add the targets by selecting them from the list of Available Items by clicking **Add to Target List**.
6. Click **Next** to select the Packages.
7. Click **Next** to display the summary.
8. Review the Summary and select one of the options as required:
 - Save Report: Saves the report as a template and closes the wizard.
 - Run Report: Runs the report and closes the wizard window.

The report result is displayed under the Report Results in the center pane.

Creating Incident Reports

You can create incident reports to obtain information about incidents.

The following types of incident reports are available:

- The Incident Summary Report: This is a historical report that summarizes information about all alerts and incidents for a specified category, such as alarm

state, alarm owner, asset type, date range, severity levels, and affected asset groups.

- The Incident Detail Report: This contains detailed information about one or more incidents. In addition to a summary, the report includes an audit trail consisting of state-change annotations, alert annotations, suggested-fix annotations, comment annotations, operation annotations.

Each incident has four pages, after the summary page. They are:

- Details of the incident
- Suggested actions, if any
- Alerts History
- Any annotations that are associated with the incident

When you create a report, you can save the report as a template, or you can generate the report. After a report is created, you can view the report, re-run the report to get updated information, or save it as a template.

To Create an Incident Summary Report

1. Select **Reports** in the Navigation pane.
2. Select Incident Reports, and then click Incident Summary Report.
3. Select **Create Incident Summary Report** from the Actions pane.
4. Define the report parameters:
 - Report Name: The name of the report.
 - Description: A description of the report.
 - Schedule: Select **Create Schedule** to schedule the report.
 - Output Format: Select the output format of the report result. CSV and PDF formats are available.
 - Select Assets: Select either all assets or specific assets.
5. Click **Next** to specify the Incident Parameters.
6. Define the incident parameters by providing the following information:
 - To create a historical report of all incidents and the date that each incident was detected, select **All Creation Dates**.
 - To create a summary report for incidents detected during a specific date range, select **Range of Creation Dates**, then enter the beginning date in the From field and the ending date in the To field.
 - To filter by severity level, owner, or state, highlight the fields to include in the report. Use Ctrl+Enter to select multiple options.
 - To filter by one or more criteria, add the criterion in the Description Contains field. Use a comma-delimited list for multiple criteria. For example, FileSystemUsage or FileSystemUsage,SwapUsage.
7. Click **Next** to display the Schedule.
8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.
10. Review the summary and select one of the options as required:

- **Save Template and Close:** Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
- **Run and Close:** Runs the report and closes the wizard window.

The report results are displayed under the Report Results in the center pane.

To Create an Incident Detail Report

1. Select **Reports** in the Navigation pane.
2. Select Incident Reports, and then click Incident Detail Report.
3. Select **Create Incident Detail Report** from the Actions pane.
4. Define the report parameters:
 - **Report Name:** The name of the report.
 - **Description:** A description of the report.
 - **Output Format:** Select the output format of the report result. CSV and PDF formats are available.
5. Click **Next** to display the Summary.
6. Review the summary and select **Save Template** and **Close**. This saves the report as a template and closes the wizard. You can use the report template to generate the report later.

The report results are displayed under the Report Results in the center pane.

Note: The Incident Compliance Report refers to the Microsoft Windows operating system incidents, while Incident Summary Report and Incident Detail Report refers to alerts and alarms that are raised by Oracle Enterprise Manager Ops Center incidents.

Creating System Information Reports

Create a system information report to obtain the information on assets such as operating systems, servers, chassis, logical domains, global zones, non-global zones, and M-Series servers.

To Create a System Information Report

1. Select **Reports** from the Navigation pane.
2. Select System Information Reports from the Reports section.
3. Select **Create System Information Report** from the Actions pane. The Create System Information Report Wizard is displayed.
4. Define the report parameters, including:
 - **Report Name:** The name of the report.
 - **Description:** A description of the report.
 - **Schedule:** Select Create Schedule to schedule the report.
 - **Output Format:** Select the output format of the report result. CSV and PDF formats are available.
 - **Select Targets:** Add one or more targets by selecting them in the list of Available Items by clicking **Add to Target List**.

5. Click **Next** to display the Attribute Selection.
6. Select one or more attributes from the list and click **Add** or **Add All** to choose all the attributes.

Click **Next** to display the Attribute Filters.

7. To set a filter for an attribute, select the attribute and specify its condition. Click the **Add icon** to set filters for other attributes.

Click **Next** to display the Schedule, when you are finished setting the filters.

8. Select a desired schedule to run and generate the report.
9. Click **Next** to display the Summary.

10. Review the report parameters and select one of the options as required:

- **Save Template and Close:** Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
- **Run and Close:** Runs the report and closes the wizard window.

The report results are displayed under the Report Results in the center pane.

To view the report, select it from the Reports Results section and then click one of the icons to choose the format: View interactively, View CSV, or View PDF.

See [Viewing and Exporting Report Results](#) for more information about generating a compliance job from the result.

Creating Exalogic Reports

The Exalogic report is all about viewing the rack setup for each of the rack within the system including the asset details related to the firmware. Exalogic Report parameters display the name and model of the Exalogic system. This also displays the number of assets in the rack by their types, such as compute nodes, switches, and storage nodes in the system.

To Create an Exalogic Report

1. Select **Reports** from the Navigation pane.
2. Select Exalogic Reports from the Reports section.
3. Select **Create Exalogic System Report** from the Actions pane. The Create Exalogic System Report Wizard is displayed.
4. Define the report parameters, including:
 - **Report Name:** The name of the report.
 - **Description:** A description of the report.
 - **Schedule:** Select **Create Schedule** to schedule the report.
 - **Output Format:** Select the output format of the report result. CSV and PDF formats are available.
5. Click **Next** to display the Schedule.
6. Select a desired schedule to run and generate the report.
7. Click **Next** to display the Summary.
8. Review the report parameters and select one of the options as required:

- Save Template and Close: Saves the report as a template and closes the wizard. You can use the report template to generate the report later.
- Run and Close: Runs the report and closes the wizard window.

The report results are displayed under the Report Results in the center pane.

9. Once the report is run successfully, select the report in the Reports Results section and click the PDF or CSV icon to open the report.

Related Resources for Reports

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

- Oracle Enterprise Manager Ops Center Documentation Library at http://docs.oracle.com/cd/E27363_01/nav/index.htm
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm

This chapter describes hardware management features that are available in Oracle Enterprise Manager Ops Center:

- [Introduction to Managing Hardware Assets](#)
- [Roles for Hardware Management](#)
- [Actions for Hardware Management](#)
- [Location of Hardware Information in the User Interface](#)
- [Profiles for Hardware Management](#)
- [Configuring the Service Processor](#)
- [Configuring a RAID Controller](#)
- [Configuring a Dynamic System Domain](#)
- [Configuring a Rack and Placing Components](#)
- [Hardware Monitoring](#)
- [Monitoring Power Utilization](#)
- [Maintaining Hardware Assets](#)
- [Firmware Provisioning](#)
- [Related Resources for Hardware Management](#)

Introduction to Managing Hardware Assets

Oracle Enterprise Manager Ops Center provides comprehensive lifecycle management for the hardware assets in your data center. The hardware assets can be handled individually or as a group. After the discovery and management of the hardware assets, as described in [Asset Management](#), you can configure and then monitor them to gather the information to maintain them.

Configuring Hardware Assets

Use these deployment plans:

- [Install Server](#)
- [Update BIOS Configuration](#)
- [Configure Service Processor](#)
- [Configure Server Hardware and Install OS](#)

- Configure RAID
- Configure M-Series Hardware, Create and Install Domain
- Configure and Install Dynamic System Domain

All the plans are based on hardware resource profiles. See [Hardware Resource Profiles](#)

Monitoring Hardware Assets

As soon as a hardware asset is managed, Oracle Enterprise Manager Ops Center starts to monitor it, according to the asset type's monitoring profile. The center pane displays information for a selected asset in a series of tabbed windows. The tabs and the type of information is specific for the asset type but, in general, Oracle Enterprise Manager Ops Center reports the following:

- Health status
- Power state
- Power usage
- Hardware variables and connectivity

You can change the monitoring thresholds in the standard profile to specify the conditions that generate an alert. You can also create custom profiles with different rule sets and alert parameters and apply the profile to a specific system, a group of homogeneous systems, or a group that you define.

For more information about monitoring policies and rules, see [Chapter 4](#).

For information about monitoring power consumption, see [Energy Tab](#).

Maintaining Hardware Assets

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

- Update management credentials. See [Chapter 2, "Asset Management"](#)
- [Setting and Changing the Power Policy](#)
- Power systems on and off
- Place in maintenance mode
- Reset a server
- Get access to the serial console
- Use locator lights to identify a specific asset
- Check firmware compliance and update firmware

Use these deployment plans or customize them:

- Update Firmware
- Update Storage Appliances
- Update BIOS Configuration

Roles for Hardware Management

[Table 11-1](#) lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See

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

Table 11–1 Hardware Tasks and Roles

Task	Role
Configure and Deploy Server	Server Deploy Admin
Install Server	Server Deploy Admin
Configure RAID	Server Deploy Admin
Update Management Credentials	Security Admin
Importing and uploading firmware images	Storage Admin
Edit Attributes	Asset Admin
Power On, Power Off, Power on with Net Boot	Asset Admin
Set Power Policy	Asset Admin
Reset Servers, Reset Service Processors, Refresh	Asset Admin
Locator Light On/Off,	Asset Admin
Snapshot BIOS Configuration, Update BIOS Configuration	Asset Admin
Update Firmware	Update Admin
Simulate a firmware update	Update Admin
	Update Sim Admin
Launch LOM Controller	Asset Admin
Edit Tags	Asset Admin

Actions for Hardware Management

After you manage your assets, you can perform the following actions:

- Use a resource profile to configure a hardware asset
- View utilization of systems
- Modify energy consumption
- Reset a server
- Power systems on and off, including a forced power off
- Place in maintenance mode
- Get access to the serial console
- Use locator lights to identify a specific asset
- Check firmware compliance
- Update firmware
- Use a provisioning profile to update firmware

Location of Hardware Information in the User Interface


Hardware assets are visible in the All Assets section of the user interface. For assets with a service processor, both the service processor and the system are included, as shown in [Figure 11–1](#).

Figure 11–1 Managed System



Each type of hardware is also added to the appropriate group in the Resource Management view, as shown in [Figure 11–2](#).

Figure 11–2 Example of Resource Management View



[Table 11–2](#) shows where to find information.

Table 11–2 Location of Hardware Information in the BUI

To See	Location
All servers	Expand Servers in the Assets pane. The center pane lists up to 50 servers.
Details for a specific server	Expand Servers in the Assets pane. Then select one of the servers. The center pane includes a series of tabbed displays.
Monitoring Rules for a specific server	Expand Servers in the Assets pane. Then select one of the servers. Then select the Monitoring tab.
Monitoring rules for a hardware type	Expand Plans and then Operational Plans and then Monitoring Policies. Select a type. The center pane shows all of the rules and thresholds.
Current firmware version	Expand Servers in the Assets pane. Then select one of the servers. Then select either the Summary tab or the Hardware tab. Each one includes the Firmware table.
A report of all attributes	Expand Reports and then select System Information Report.

Profiles for Hardware Management

The work flow of an asset deployment can be captured and enacted in a repeatable fashion using plans and the profiles included in the plans. Oracle Enterprise Manager Ops Center provides default profiles for configuring hardware asset types consistently. You can use the default profiles, make copies of the profile to edit, or create new ones.

Deployment proceeds from configuring the hardware, installing the correct firmware, provisioning the OS, and applying the required updates. The deployment is not only for servers with an OS installed but also for chassis, racks, power distribution units, and M-Series servers.

Hardware Resource Profiles

Use hardware resource profiles in a deployment plan to configure, install, or update systems. The Oracle Enterprise Manager Ops Center software provides the following hardware resource profiles:

- [Configuring the Service Processor](#)
- [Configuring a RAID Controller](#)
- [Configuring a Dynamic System Domain](#)
- [Configuring a Rack and Placing Components](#)

You can use the default profiles, make copies of the profile to edit, or create new ones. Use the profiles in a deployment plan to configure your hardware assets.

Firmware Provisioning Profiles

Hardware depends on firmware to perform operations. Part of monitoring and managing a hardware asset is to make sure that it has the appropriate version of firmware. Oracle Enterprise Manager Ops Center uses firmware profiles to provision, or update, firmware on each type of asset.

A firmware profile is a set of actions and values that define how to provision one or more assets and specifies one or more firmware images. A firmware profile updates existing firmware assets completely and consistently. When you apply a deployment plan that contains a firmware profile, Oracle Enterprise Manager Ops Center compares the versions of each firmware image specified in the profile with the versions of the existing firmware on the asset and then takes the action you specify in the profile.

See [Firmware Compliance Reports](#) and [Firmware Provisioning](#).

Configuring the Service Processor

The **Declare Unconfigured Asset** action includes the service processor in the Oracle Enterprise Manager Ops Center environment. You can configure only unconfigured service processors, that is, a processor in its factory default state. Use a deployment plan to configure the service processor:

- Configure Service Processor
- Update BIOS Configuration
- Configure and Deploy Server
- Install Server

An alternative to specifying the configuration in the profile, is to duplicate an existing configuration. You can create a snapshot of the BIOS of a working service processor, which creates also creates profile. You then apply the profile to an unconfigured service processor.

Creating a Service Processor Configuration Profile and Plan

1. Expand Plans in the Navigation pane and the select Profiles and Policies.
2. Expand Service Processors.

3. Click Create Profile. The first step of the wizard is displayed, as shown in [Figure 11-3](#). Depending on the subtype and target you select, more steps are added.

Figure 11-3 Create Profile - Service Processor

Oracle Enterprise Manager Ops Center - Create Profile - Service Processor

Create Profile - Service Processor

ORACLE

Steps Help

1. Identify Profile
2. Summary

Identify Profile * Indicates Required Field

* Name:

Description:

☒ Create a deployment plan for this profile.

* Subtype:

- Subtype
- CMM SP with ILOM 3.0
- Server SP with ILOM 3.0
- M-Series SP

Target Type:

- Target Type
- Blade Chassis

4. Complete the specification of the service processor and click Finish.
The new profile and plan are available from the Assets pane.

Creating a BIOS Configuration Profile and Plan

1. Expand Plans in the Navigation pane and the select Plans.
2. Expand Update BIOS Configurations.
3. Click **Create Plan from Template**. [Figure 11-4](#) shows the first step of the wizard, which includes the Update BIOS profile.

Figure 11–4 Create Plan - BIOS

Oracle Enterprise Manager Ops Center - Create a Deployment Plan

Create a Deployment Plan ?

* Indicates Required Field

* Plan Name:

Description:

Failure Policy: ☒ Stop at failure ☐ Complete as much as possible

Target Type: Hardware

Template Name: Update BIOS

Deployment Plan Steps

Step	Profile/Plan Type	Associated Profile/Deployment Plan
Update BIOS	BIOS Profile	

4. Click the Update BIOS profile.
5. Click the **Create Profile** icon. Figure 11–5 shows the window for creating the profile.

Figure 11–5 Create Profile-BIOS Configuration

Oracle Enterprise Manager Ops Center - Create Profile - BIOS Configuration

Create Profile - BIOS Configuration

Steps Help

1. Identify Profile

2. Summary

Identify Profile

* Indicates Required Field

* Name:

Description:

* Subtype: Subtype

BIOS

UEFI

Target Type: Target Type

Servers

6. Enter a name and description and then select the type of BIOS and the type of target.
7. Click **Next** to review and then click **Finish** to submit the job.
8. When the job is completed, return to the Create a Deployment Plan window.
9. Specify a name for the plan and select the new profile.
10. Click **Save** to submit the job.

The new profile and plan are available from the Assets pane.

Creating a Snapshot of a Service Processor Configuration

1. Expand Assets in the Navigation pane and select Servers.
2. Select the configured server.
3. Click the **Hardware** tab in the center pane.
4. In the Component Navigation section, select **Service Processor**. The Service Processor Configuration section displays the configuration attributes and values.
5. In the Service Processor Snapshots section, click the **Create Snapshot** icon to display the window shown in [Figure 11–6](#).

Figure 11–6 Create Service Processor Configuration Snapshot

Note: An alternative method of displaying this window if you are familiar with the current configuration is to select the server and then choose **Snapshot BIOS Configuration** in the Actions pane.

6. Specify a name for this snapshot and accept the default of creating a profile from the snapshot.
7. Click **Create Snapshot** to submit the job.
8. When the job completes, the snapshot is located in the EC local library.

Applying a Snapshot of a Server Processor Configuration

This procedure applies a snapshot of one server's service processor to another, unconfigured service processor. Use the **Update BIOS Configuration** action or use the snapshot profile as a step in the Install Server deployment plan.

View a Service Processor Snapshot

To verify the snapshot configuration before you attempt to configure a new server, view the snapshot.

1. Expand Libraries in the Navigation pane and then select the current EC local library.
2. In the center pane, scroll down to the BIOS Configuration section of the library.

3. Select the snapshot and then select the **View Snapshot** icon.
4. You can apply the snapshot directly to a server or you can use a deployment plan.

Apply a Service Processor Snapshot to a Server

1. Expand Assets in the Navigation pane and then select Servers.
2. Select the unconfigured server.
3. Click **Update Service Processor** in the Action pane. A one-step wizard is displayed, including the list of snapshots in the EC local library that are appropriate for the selected server.
4. Select the snapshot from the list.
5. Click **Next** to review the Summary and then click **Finish** to submit the job.

When the service processor is configured, the server reboots.

Apply a Service Processor Snapshot Profile to a Server

1. Expand Plans in the Navigation pane and then select Profiles.
2. Expand Update BIOS Configuration and select the profile for the snapshot.
3. Review the profile in the center pane. To modify the profile, click the **Edit Profile** icon.
4. Create a plan that includes this profile as its only step or edit an existing deployment plan to include this profile.

Configuring a RAID Controller

Oracle Enterprise Manager Ops Center provides a default profile for configuring RAID controllers. You can configure and update the hardware devices that can be done only through the host OS. This profile provisions a reduced OS image along with the management pack on the target to configure RAID.

Caution: When you reconfigure an existing RAID controller, all the data on the disk is lost.

Configuring a Dynamic System Domain

You can create domains on a M-series server.

- Verify that a user account with `platadm` privilege exists on the XSCF processor and is included in the profile. Disable the audit policy for this user account, using the following command:

```
setaudit -a opsadm=disable
```

- Verify that you can use `ssh` to log in to the XSCF processor.

To configure a Dynamic System Domain, you apply a deployment plan to an M-Series server, as described in [Chapter 8](#). The plan includes the profile for the Dynamic System Domain.

Configuring a Rack and Placing Components

The rack asset is a group that includes the other managed assets installed in the physical rack, such as servers or compute nodes, storage arrays or appliances, switches, and power distribution units (PDU). After you create the rack, you put assets in the rack so that Oracle Enterprise Manager Ops Center can present them and manage them as if it were a physical asset.

Creating a Rack

1. Expand Assets in the Navigation pane and select Racks.
2. Click **Create Rack** in the Action pane.
3. Enter the name for the rack.
4. For a rack containing an Oracle Engineered System, specify the serial number of the rack.
5. Enter the total number of slots in the rack. A full rack has 42 slots. Default to 42.
6. You have the option to place assets into the rack immediately after this procedure ends or to perform this task at a later time.
7. Enter a description for this rack.
8. Add semantic tags that are appropriate for this rack.
9. Click **Create Rack** to submit the job.

Use the **Place/Remove Assets in Rack** action and the **Place/Remove PDU in Rack** action to specify each asset in the physical rack and its location in the rack so that Oracle Enterprise Manager Ops Center can represent the physical rack accurately to remote users.

Placing Assets in a Rack

The rack asset is a group that includes the managed assets installed in the physical rack, such as servers or compute nodes, storage arrays or appliances, switches, and power distribution units (PDU).

Use the **Place/Remove Assets in Rack** action and the **Place/Remove PDU in Rack** action to specify each asset in the physical rack and its location in the rack so that Oracle Enterprise Manager Ops Center can represent the physical rack accurately to users who cannot examine the rack. Update the type and location of the assets in the rack asset when the configuration of the physical rack changes.

Placing a Power Distribution Unit in a Rack

1. Expand Assets in the Navigation pane and select Racks from the Resource Management Views.
2. Select the rack.
3. Click **Place/Remove Assets in Rack** in the Actions pane or click the **Details** tab to navigate to the icon. The Place Assets in *name* Rack window opens, which displays a list of all managed assets in the physical rack. You can filter the list by type or attribute.
4. Click an asset and then enter its slot number.
5. Click **Place Asset in Rack**. The Assets in the Rack pane is updated to show the location.
6. Continue to select assets and place them in the rack.

7. To change an asset's slot, select it and click **Edit Placement**. Enter a new slot number.
8. To remove an asset, select it and click **Remove Asset**. The asset is deleted from the Assets in the Rack pane and is available to be placed.
9. When you are satisfied with the configuration displayed in the Assets in the Rack pane, click **Submit**.

To add PDUs to the rack asset, use the **Place/Remove PDUs** action in the same way. Select the rack and the action. The assets you have placed in the rack are shown in the Assets in Rack pane. Add each PDU.

Hardware Monitoring

The Oracle Hardware Management Agents use Simple Network Management Protocol (SNMP) to monitor your Oracle hardware and storage devices. The software uses the Intelligent Platform Management Interface (IPMI) protocol to access the Oracle ILOM service processors.

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

Hardware Status

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

- Good – The hardware asset is working properly.
- Unknown – Oracle Enterprise Manager Ops Center is unable to retrieve information from the sensor. The hardware asset is connected but is not reporting information.
- Unreachable – The hardware asset cannot be contacted. This state indicates a network problem.
- Warning Failure – Oracle Enterprise Manager Ops Center has detected a potential or impending fault condition. Take action to prevent the problem.
- Critical Failure – A fault condition has occurred. Take corrective action.
- Nonrecoverable Failure – The hardware asset has failed. Recovery is not possible.
- Faulted – The hardware asset reports a fault. Contact service personnel to repair.

Groups of Hardware Assets

Oracle Enterprise Manager Ops Center monitors hardware assets according to the monitoring profile for that type of asset. To see the default profile for monitoring a hardware type, see [Hardware Monitoring](#).

To change a threshold value, see [Editing Monitoring Rules](#) in [Chapter 4](#).

Connectivity Status

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

Service Processor Details

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

- Name and SNMP Community
- Whether Auto DNS through DHCP is in use, and any DNS Servers
- Search Path, if any
- Time Zone
- Whether an NTP server is in use and its identifier

RAID Controller Details

- RAID volume name
- RAID level
- Number of disks
- Stripe zone
- RAID Controller ID

Sun ZFS Storage Appliance Details

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

Note: The Create Filesystem and Delete Filesystem actions are disabled for the appliance installed in Exalogic systems.

See [Sun ZFS Storage Appliance](#) in [Chapter 2](#) for discovery requirements.

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 Details

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

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 adds details to the information in the Dashboard tab. For the Power status, the reported status is for the server's domains. When any domain is powered on, the status is reported as powered on. When all domains are powered off, the Summary tab shows a status of Powered Off; the M-Series server itself remains powered on.

- 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 tab include the following information.

- The Unallocated Resources table lists all the physical system boards and their status: PSD ID, Assignment Status, Power Status, Connection Status, Diagnostics Status, and Operational Status
- The Allocated Resources table lists all domains that are using the physical system boards and their status: Domain ID, PSB ID, XSB ID, LSB ID, Assignment Status, Power Status, Connection Status, Diagnostics Status, and Operational Status
- The Dynamic System Domain table lists all the domains and their details: Domain ID, MAC Address, Autoboot Policy, Secure Mode Policy, CPU Mode, Diagnostics Level, Domain Degradation Policy, and Operational Status

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

Oracle Enterprise Manager Ops Center monitors the voltage for the Board and IO Unit components and the speed for the Fan components. The Monitoring tab shows the actual value and the threshold values.

See [Discovering a SPARC Enterprise M-Series Server](#) in [Chapter 2](#) for discovery requirements.

Switch Details

Oracle Enterprise Manager Ops Center can manage Sun Ethernet 10GbE Fabric switches, Sun Datacenter InfiniBand switches, and Cisco Catalyst 4948 Ethernet switches. These switches reside in the system or blade system and provide the switch fabric.

- For the Sun Ethernet 10GbE Fabric switch, see <http://docs.oracle.com/cd/E19934-01/index.html>
- For the Sun Datacenter InfiniBand switch, see <http://docs.oracle.com/cd/E19654-01/index.html>
- For the Cisco Catalyst 4948 Ethernet switch, see the product information at www.cisco.com

Oracle Enterprise Manager Ops Center reports hardware information on the Summary tab:

- Name
- Description
- 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)

Rack Details

The rack's Dashboard displays the following information:

- Name
- Rack ID
- Description
- Number of Slots
- Tags
- Support

The Details tab shows the configuration of each slot in the rack. For each component in the rack, this tab displays the position, name, description, type, model and health status. The Power Distribution Units are included and their current status is displayed.

The Firmware tab displays the name, description, and version of the firmware for each component and the slot for each component:

- Compute Nodes
- Switches
- Storage Appliances
- Power Distribution Units

The rack's Charts tab displays the following plots, as described in [Charts Tab](#):

- Aggregate Power Usage
- Power Usage
- Average Fan Speed

The rack's Energy tab displays the information described in [Energy Tab](#)

Note: The Power Off and Reset actions are disabled for Exalogic systems and for single-rack configurations.

PDU Details

The Dashboard tab shows:

- Name and description
- Model and serial number

- Management IP address and MAC address, if any

The Details tab reports the same information and adds whether SNMP and HTTP is enabled and the version of the firmware.

Oracle Solaris Cluster Details

The Dashboard tab shows:

- Name, description, and ID
- Number of possible quorum votes and the current quorum votes

The Network tab shows the public and private interconnects used by the cluster.

The Quorum tab shows the status of each member of the quorum and the number of quorum votes for each member.

Oracle Exalogic Details

Oracle Enterprise Manager Ops Center supports bare metal provisioning and discovery of hardware assets of an Exalogic machine. The Oracle Exalogic architecture includes InfiniBand and Ethernet switches and the Sun ZFS Storage 7320 Appliance. Using Oracle Enterprise Manager Ops Center, you can manage these components individually and as a group, as well as get access to the browser user interface to the storage appliance and serial console access to the compute nodes and switches.

For detailed information about how to use Oracle Enterprise Manager Ops Center with Oracle Exalogic, see the Oracle Exalogic documentation:

http://docs.oracle.com/cd/E18476_01/doc.220/e18478/toc.htm

Monitoring Power Utilization

Input power is the power pulled into a power supply from an external resource. The power consumption of a hardware asset is the sum of the input power consumed by each power supply of the asset. Output power is the amount of power provided from the power supply to the system components, measured at the power supply output. Input power is calculated from output power by applying an efficiency function to the output power from each power supply.

Calculating power compensation for the blades is difficult because the power supplies are shared. Each blade gives a report based on the power consumption of the local components, but this is not an accurate power consumption value for an individual blade.

To measure the input power, the interfaces must be exposed and the service processors must be able to retrieve and report data with one-minute accuracy. Servers that can report power usage have a Charts tab.

You can see current power usage and change the display of power graphs using the controls on the Energy tab and the Charts tab.

Energy Tab

The asset's Energy tab reports power consumption as the current value and for a period of time, as well as attributes of the fan and power supplies. 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 temperature
- Power Policy
- Cost Per KiloWatt-hour
- Currency units used to compute cost. The price per currency unit is set by the **Edit Energy Cost** action in the Administration section of the Navigation pane. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information.
- 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 Celsius, and the average fan speed in RPM. Click 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 the **Live** button in the graph's toolbar to make the information static so that you can set 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 heterogeneous group of servers.
 - The Power option displays power utilization for the five highest or lowest power consumers in the group or virtual pool.
 - The Aggregate Power option displays the power utilization, using the sum of all members that report power consumption. The number of systems in the aggregate is included. For heterogeneous group, the Chart tab includes a table of all systems in the group and their various power attributes for the selected time period. From this table, you can power off and power on selected servers to conserve power.

Maintaining Hardware Assets

- Update management credentials. See [Updating Management Credentials](#) in [Chapter 2](#)
- Set power policy
- Power systems on and off
- Place in maintenance mode
- Reset a server
- Get access to the serial console
- Enable and disable ports
- Use locator lights to identify a specific asset
- Check firmware compliance and update firmware. See [Firmware Provisioning](#)

Setting and Changing the Power Policy

An asset's use of power can be optimized for performance or for conservation:

- Performance – Disable CPU power management to use of more power.
- Elastic – Enable CPU power management to conserve power.

View an Asset's Power Policy

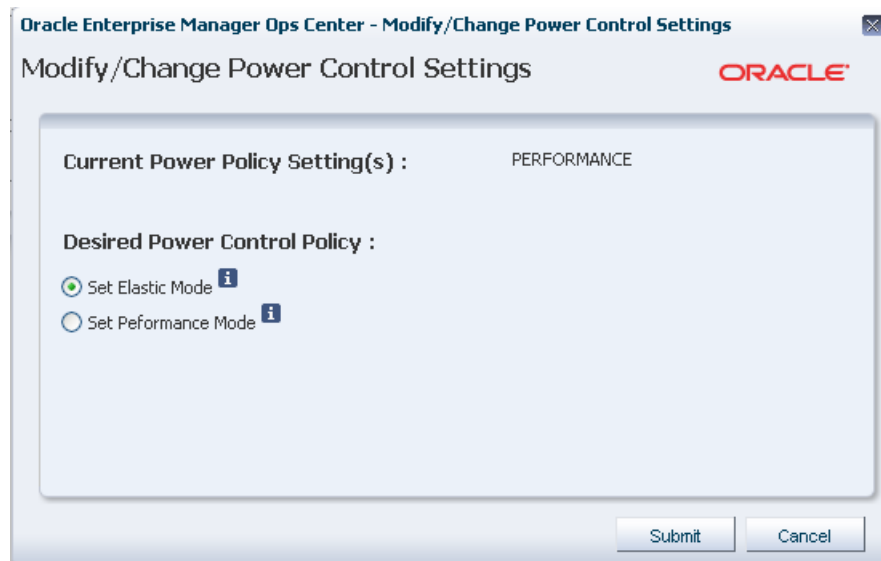
1. Expand the Assets in the Navigation pane.
2. Select a hardware asset.
3. Click the **Energy** tab in the center pane.

Changing an Asset's Power Policy

1. Expand the Assets in the Navigation pane.
2. Select a hardware asset.

3. Click **Set Power Policy** in the center pane to display the window shown in [Figure 11-7](#).

Figure 11-7 Set Power Policy



4. The current power policy is displayed. Choose the alternative policy.
5. Click Submit to create a job that sets the power policy.

Replacing a Failed Power Distribution Unit in a Rack

1. Expand Assets in the Navigation pane and then expand Racks.
2. Navigate to the rack type and expand to show its components. Expand Power Distribution Units.
3. Select the failed power distribution unit.
4. Click **Place in Maintenance Mode** in the Action pane.
5. Go to the rack's location and remove the failed PDU.
Install the new PDU and connect it to the network, according to the procedures in the [Power Distribution Units User's Guide](#).
6. Verify that the new PDU has the same IP address as the failed PDU.
7. Use the PDU's web interface to configure the administrator user account, according to the procedure in the [Power Distribution Units User's Guide](#).
8. Return to the Oracle Enterprise Manager Ops Center user interface and navigate to the same rack and the new PDU. Select the PDU.
9. Click **Remove From Maintenance Mode** in the Action pane.
10. To update the credentials for the administrator user, click **Update Management Credentials** in the Action pane. The wizard opens.
11. At the Management Type step, select HTTP credentials and then select the Create a new set of credentials option. Click Next.
12. Enter new credentials for the administrator user.

13. Click Apply to update credentials.
14. To create new SNMP credentials, click **Update Management Credentials** in the Actions pane again.
15. Select SNMP credentials and Create a new set of credentials.
Specify a community string that is different than the previous community string.
16. Press Apply to update credentials.

Installing and Upgrading Oracle Solaris Cluster

You use cluster profiles in a deployment plan to perform the following operations:

- Install Oracle Solaris Cluster software
- Upgrade Oracle Solaris Cluster software

See [Appendix B](#) for instructions on obtaining the current profiles and supporting scripts.

Firmware Provisioning

The Oracle Enterprise Manager Ops Center provisioning feature installs firmware on the managed hardware assets, using the protocols and ports listed in [Table 7-6](#). You initiate the installations from the UI, rather than from the asset itself.

Oracle Enterprise Manager Ops Center provides default profiles for configuring firmware for servers and for disk storage. An alternate procedure is to create a Firmware Report. You then use the report results to update the firmware.

Note: The Update Firmware and Firmware Reports actions are disabled for Exalogic systems, except for its storage resource, the Sun ZFS Storage Appliance.

The benefit of using a profile to install firmware is that the firmware is installed consistently, no matter how many assets you provision. The benefit of using the Firmware Compliance Report is to identify the firmware on a specific asset or set of assets.

Firmware Profiles

The general procedure for provisioning firmware has the following steps:

1. Import a file with the firmware and the associated metadata into a software library, according to the procedure in *Keeping Your Firmware Up-to-Date* in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.
2. Create a firmware profile, based on one or more firmware images, according to the procedures in *Keeping Your Firmware Up-to-Date* in the How To Library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.
3. Shut down the server gracefully. Most firmware requires that the server is not running when the firmware is updated. Most firmware images include a power-off command for a running server, which causes a hard shutdown of the server.
4. Apply the firmware profile.

Firmware Compliance Reports

The Firmware Report feature compares the firmware images specified in a firmware profile to the firmware images installed on one or more hardware assets. Use the **Make Targets Compliant** button in the Interactive type of report to update the firmware on any non-compliant assets.

The report shows whether a target asset is compliant, not compliant, or not applicable:

- A compliant asset has the firmware images specified in the profile.
- A non-compliant asset does not have the same firmware images as specified in the profile. Update the firmware by either clicking the Make Targets Compliant button in the Interactive report or using the procedure in [Updating Firmware](#).
- A non-applicable asset indicates that a firmware image in the profile does not match the model of service processor in the asset. This condition can occur when either the profile does not recognize the model of the service processor or the profile includes firmware images that are not designed for the service processor.
 1. Compare the model of the service processor displayed in the asset's Summary tab with the model of the service processor included in the profile. If they are different, add the name in the profile to the asset's data.
 2. When the firmware profile was created, only images that matched the service processor could be included. However, if the service processor did not report all the firmware types it supported, an image that did not match the service processor could have been included in the profile. To update the Oracle Enterprise Manager Ops Center software with all the service processor's supported firmware types, use the **Refresh** action. When the job is completed, view the service processor's Summary tab to see all firmware types.
 3. Create a new firmware compliance report.

See [Creating a Firmware Report](#) for the procedure to create the report.

Updating Firmware

To update the firmware on one or many assets, you use a deployment plan to apply a firmware profile for the type of asset. For a server, the profile updates the firmware on a service processor, and restarts the service processor and operating system. For storage components, profiles update firmware on a RAID controller, an expander, or disk.

To see the deployment plans that update firmware, expand the Deployment Plans section of the Navigation pane and then click Firmware. A list of existing plans and profiles is displayed.

To update the firmware of one asset, an alternative to a deployment plan is to use the **Update Firmware** action. Select the asset from the Asset section of the Navigation pane and then click **Update Firmware** in the Actions pane.

Before You Begin

The software library must contain the images that provision the firmware. Perform the Uploading a Firmware Image procedure in the *Keeping Your Firmware Up-to-Date* document in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

If you are updating the firmware on a server, shut down the server before you update the firmware. A firmware update to a server's service processor usually requires that

the server is not running. If you start to update the firmware on a running server's service processor, the procedure performs a hard shutdown of the server.

Requirements for ALOM Service Processors and M-Series Servers

The firmware provisioning process for M-Series servers and servers that have ALOM service processors relies on a temporary account that performs an FTP operation.

Note: Advanced Lights Out Management (ALOM) is a Sun Microsystems standard for these systems: SunFire V125/V210/V215/V240/V245/V250/V440/T1000 and Sun Netra 210/240/440.

If your site does not allow a temporary account, use the following procedure to prepare for the provisioning operation:

1. On the Enterprise Controller, open the `/var/opt/sun/xvm/hal.properties` in an editor.
2. Add the following properties to the file:

```
ftp.user.name=username  
ftp.user.password=password
```
3. Restrict access to the file to root user:

```
chmod 600 /var/opt/sun/xvm/hal.properties
```
4. On the Proxy Controller that provisions the firmware, enable the `ftp` service on Oracle Solaris or the `vsftpd` service for Oracle Linux systems.

You can now apply the firmware provisioning deployment plan. The FTP operation retrieves the credentials from the file.

If a network failure occurs while updating the firmware, repeat the firmware update procedure. If you do not repeat the procedure, the firmware inventory list might be incomplete.

Option for Deferring the Stop and Restart of the Operating System and Server

The procedure is available for Oracle Solaris 10 Update 10 operating systems running on servers with the ILOM x86 (3.0 and higher) service processor.

When you update the firmware of a service processor, the procedure stops the operating system and the server before the update and restarts them after the update so that the new BIOS takes effect. If you prefer to stop and restart at a convenient time, keeping the current BIOS in effect, use the following procedure to change the action of both Oracle Enterprise Manager Ops Center and the firmware's metadata:

1. On the Enterprise Controller, open the `/var/opt/sun/xvm/hal.properties` in an editor.
2. Add the following property to the file:

```
ilom.fwp.skipAutoReboot=true
```
3. On the Proxy Controller that provisions the firmware, enable the `ftp` service on Oracle Solaris or the `vsftpd` service for Oracle Linux systems.
4. At a later time, reboot the servers.

Related Resources for Hardware Management

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

- For the Sun Ethernet 10GbE Fabric switch, see <http://docs.oracle.com/cd/E19934-01/index.html>
- For Sun Datacenter InfiniBand switch, see <http://docs.oracle.com/cd/E19654-01/index.html>
- [Power Distribution Units User's Guide](#)

Operating Systems

This chapter provides an overview of the operating system (OS) management features that are available in Oracle Enterprise Manager Ops Center.

The following information is included:

- [Introduction to Operating Systems](#)
- [Roles for Operating Systems](#)
- [Actions for Operating Systems](#)
- [Location of Operating System Information in the User Interface](#)
- [Operating System Profiles](#)
- [Using Agent Management for Operating Systems](#)
- [Monitoring Operating Systems](#)
- [Using Analytics](#)
- [Provisioning Operating Systems](#)
- [Overview of Oracle Solaris Boot Environments](#)
- [Overview of Oracle Solaris 11 Boot Environments](#)
- [Overview of Oracle Solaris 10 Boot Environments](#)
- [Related Resources](#)

See [Chapter 13, "Operating System Updates"](#) for information about patching and updating your operating systems.

Introduction to Operating Systems


Oracle Enterprise Manager Ops Center provides comprehensive lifecycle management for Oracle Solaris and Linux operating systems in your datacenter. The software also enables you to patch, or update, Microsoft Windows operating systems.

The discovery feature makes adding operating systems and other assets quick and easy. After the operating systems are added, they are considered managed and you can begin using the monitoring, analytics, OS provisioning, and update features to gather information and perform tasks.

Your managed operating systems are visible in the All Assets section of the user interface. The operating system appears under the service processor and hardware, as shown in [Figure 12-1](#).

Figure 12–1 Managed System and Operating System

The operating system also appears in the appropriate platform-specific group in the Operating Systems view, as shown in [Figure 12–2](#).

Figure 12–2 Operating Systems View

You can create user-defined groups and subgroups to refine your administration tasks. For example, you might want to create groups for Critical Systems, Training, Region 1 and Region 2. Groups are useful when you want to organize the systems to apply different monitoring standards, implement update requirements, or job scheduling times. You can create rules for your groups to automatically add existing and newly managed operating systems to the correct group or subgroup. See [Chapter 2, "Asset Management"](#) for details on using the discovery feature to add assets and for more information on creating user-defined groups.

Two types of OS management are available: agent managed and agentlessly managed. In some cases, the features and actions that you can perform on an operating system are determined by the management type. See [Using Agent Management for Operating Systems](#) for more information about the features available for agent and agentlessly managed operating systems.

The following features are available for operating systems:

- **Monitoring:** A series of monitoring rules and parameters monitor your managed assets. Alerts and incidents are raised for components that are not performing as expected.
- **Performance:** Analytics provides you with a detailed view into OS performance.
 - System resource graphs, processes information, and a view of the top consumers
 - Resource usage of virtualized OS instances
- **Provisioning:** Install Oracle Solaris or Linux operating systems onto your systems, making it easy to install one system or many servers simultaneously.

Note: The Enterprise Controller must be installed on an Oracle Solaris operating system in order to perform the following tasks:

- Provision Oracle Solaris 10 using JET customization.
 - Provision Oracle Solaris 11 (requires that both the Enterprise Controller and Proxy Controller are running on an Oracle Solaris 11 operating system).
 - Provision Oracle VM Server for SPARC (control domain).
-
- Manage Oracle Solaris Boot Environments: Create and manage Oracle Solaris boot environments in a repeatable and consistent manner from a single user console.
 - OS Updates: Apply update packages to keep your operating systems up-to-date. See [Chapter 13, "Operating System Updates"](#) for information.
 - Reports and Snapshots: A variety of reports are available for your operating systems. See [Chapter 10, "Reports"](#) for OS reports. See [Chapter 13, "Operating System Updates"](#) for information on how to use the System Catalogs to maintain snapshots of your operating system.

The monitoring feature provides extensive monitoring capabilities that are enabled as soon as you begin managing an operating system. A series of three escalating status levels notifies you when something is not operating as expected. The first level is informational, then warning, and finally a critical status. A set of default monitoring attributes and alert triggers are included with the software. You can tune the monitoring thresholds and triggers to define what you want to generate an alert and when. You can create custom monitoring rule sets and alert parameters and apply the customized monitoring rules to a specific operating system, a group of homogeneous operating systems, or a group that you define, such as critical systems or regional systems.

The Analytics feature provides extensive information about a specific operating system in one location so that you can maximize performance and utilization. The Analytics information includes process details, defined monitoring thresholds for operating systems, metrics and historical information on the top consumers, and a extensive list of metrics data. The Summary contains details on the top five consumers for CPU, memory, network, and I/O utilization. Use the graphical representation to quickly view utilization trends and high resource consumers. You can drill down to get detailed utilization and process information, and kill a process that is consuming too many resources.

Roles for Operating Systems

[Table 12–1](#) lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 12–1 OS Management Roles and Permissions

Task	Role
Reboot an OS	Asset Admin

Table 12–1 (Cont.) OS Management Roles and Permissions

Task	Role
Charts and Utilization	Asset Admin Cloud Admin
Analytics	Read Asset Admin
Kill action in Analytics	Operating System Management
Update Management Credentials	Security Admin
Any Actions related to changing credentials	Security Admin
Import image	Storage Admin
Upload image	Storage Admin
Unconfigure, SCCM Configuration	Oracle Enterprise Manager Ops Center Admin
Reboot, upgrade Agent Controller	Asset Admin
Edit Tags	Asset Admin
Edit Attributes	Asset Admin
View Boot Environment	Read Asset Admin Update Admin
Create Boot Environment	Asset Admin Update Admin
Update an Alternate Boot Environment	Update Admin
Activate and Reboot a Boot Environment	Asset Admin Update Admin
Synchronize Boot Environments	Asset Admin Update Admin
Delete an Alternate Boot Environment	Asset Admin Update Admin
Monitor Boot Environment Attributes	Asset Admin

Actions for Operating Systems

You can manage an operating system in one of two modes: agent managed or agentless managed. The management mode determines the features that are enabled for your operating system.

Agent managed is the more robust management mode because the Agent Controller enables a greater level of communication with the Proxy Controller and Enterprise Controller than the agentless managed operating systems. You can use the features and perform the actions described in this chapter with an agentless managed operating system, but OS update functionality requires an agent managed operating system. You can manage your operating systems by installing an Agent Controller on the OS or by using SSH to perform tasks. See [Using Agent Management for Operating Systems](#) for more details on managing operating systems with or without an Agent Controller.

See [Chapter 2, "Asset Management"](#) for details on using the discovery feature to add assets and for more information on creating user-defined groups.

After you manage your assets, you can perform the following actions:


- Monitor your physical and virtual operating systems
- View OS utilization for Oracle Solaris and Linux operating systems
- Provision Oracle Solaris and Linux operating systems
- Manage Oracle Solaris boot environments

Note: To perform actions, such as OS provisioning, updating, and managing boot environments, on an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must be installed on an Oracle Solaris 11 operating system.

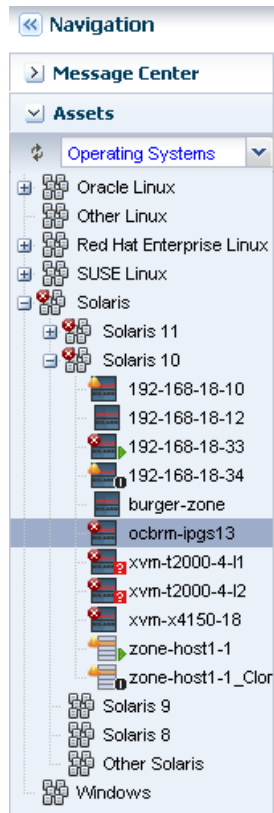
Location of Operating System Information in the User Interface

When you manage a physical or virtual operating system, it appears under the associated server in the All Assets view, as shown in [Figure 12-3](#) and it appears in a special pre-defined operating system group, as shown in [Figure 12-4](#).

Figure 12-3 All Assets View



Operating systems are automatically added to a homogenous group of operating systems. The group contains directories for each release. In [Figure 12-4](#), the list of Oracle Solaris 10 operating systems includes physical operating systems, virtual hosts, and zones. You can add user-defined groups and create rules to automatically add newly discovered assets to that group, or you can manually add them at any time. See [Using Groups](#) for details.

Figure 12–4 Operating Systems View

To view information about a specific OS, select the OS from the Assets pane. OS details appear in the tabs across the center pane.

Table 12–2 Location of Operating System Information in the UI

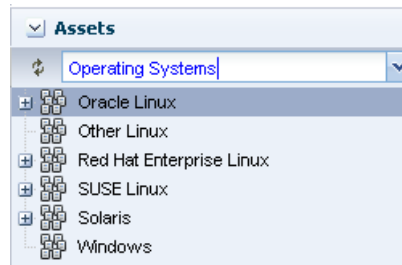
To Display	Select
Managed operating systems	Expand the Assets pane. Each operating system appears in the Assets Navigation tree under the system on which it is installed. To only view operating systems, click the drop-down next to All Assets and select Operating Systems. The systems are grouped by platform.
Operating system details for a specific operating system	Select an operating system in the Assets pane, then click the Summary tab.
Operating system details for a group of operating systems	Select an operating system in the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Summary tab.
Unresolved incidents and alerts for a specific operating system	Select an operating system in the Assets pane, then click the Incidents tab. The details are in the Unresolved Incidents and Alert subtabs.
Unresolved incidents and alerts for a group of operating systems	Expand the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Incidents tab. The details are in the Unresolved Incidents and Alert subtabs.
Monitoring Rules for a specific operating system	Select an operating system in the Assets pane, then click the Monitoring tab.

Table 12–2 (Cont.) Location of Operating System Information in the UI

To Display	Select
Monitoring Rules for a group of operating systems	Select an operating system in the Assets pane, then click the drop-down next to All Assets and select Operating Systems. Select a group, then click the Monitoring tab.
Analytics	Select an operating system in the Assets pane, then click the Analytics tab.
OS Provisioning	Select a Provision OS plan in Plan Management.
Boot environments	Select an operating system in the Assets pane, then click the Boot Environments tab.

Viewing Operating Systems

You can view all managed assets in the Assets section of the Navigation pane, or you can filter to display a specific group of assets. [Figure 12–5](#) shows the asset filter menu where you can choose the Operating System group or user-defined groups.

Figure 12–5 Asset Filter Menu

See [Using Groups](#) for information about creating user defined groups.

Displaying Operating System Details

Details about a specific operating system appear in a number of tabs in the center pane of the UI. The tabs and contents might vary based on the type of operating system selected and whether the operating system is agent or agentlessly managed. When a tab is not applicable, it will not appear in the UI. For example, the Libraries tab does not appear when there are no libraries associated with the operating system.

To view the details in the UI, perform the following steps:

- Expand **Assets** in the Navigation pane.
 - To display All Assets, expand **Servers**.
 - To display the Operating System groups, select **Operating Systems** from the menu.
 - To display user-defined groups, select **All User Defined Groups** from the menu, then select the group.
- Select an operating system in the Navigation pane.

The following tabs appear in the center pane:

- Dashboard:** Displays the summary of the selected operating system, including the name, server name, whether the operating system is agent managed, the number of unassigned incidents, the current alert status, and the OS release.

The dashboard also includes the operating system's membership graph, the status of incidents, and compliance reports, if any.

- **Summary:** Displays the name of the selected Oracle Solaris operating system, description, server name, operating system version, number of CPU threads, active boot environment, the state, the length of running time, and zone patching information. Tables display total CPU utilization and total ZPool Utilization for Oracle Solaris operating systems.
- **Libraries:** Displays library details when a library is associated with an operating system. See [Chapter 5, "Software Libraries"](#) for more information.
- **Storage:** Displays logical unit (LUN) details and iSCSI targets, when applicable. LUN details include the initiator, MPxIO details, the GUID, type, size, status, vendor, and product. The iSCSI details include the initiator, iSCSI target address, address type, port, and IP address. You can discover iSCSI addresses from this tab. See [Block Storage](#) for more information.
- **Analytics:** Displays utilization and metrics for the selected operating system and zones, when applicable. See [Using Analytics](#) for more information.
- **Connectivity:** Displays Linux OS network interface of the system, including network connectivity and aggregated links. See [Network Connectivity](#) for more information.
- **Networks:** Displays Oracle Solaris network connectivity, IPMP groups, and aggregated links. For an Oracle Solaris 11 OS, the tab also provides bandwidth management. See [Network Connectivity](#) and [Bandwidth Management](#).
- **Incidents:** Displays all incidents reported from the selected operating system. See [Chapter 9, "Incidents"](#) for more information.
- **Monitoring:** Displays the alert monitoring rules and service dependencies of the selected operating system. [Chapter 4, "Monitoring Rules and Policies"](#) for more information.
- **Terminal:** Enables you to establish an SSH connection to the terminal window.
- **Boot Environments:** Displays Oracle Solaris boot environment details, including all available boot environments, the active and enabled boot environment, the size, and the creation or synchronization date. For a selected boot environment, you can view snapshot details, file system details, and any associated zone boot environments. This tab is only available for Oracle Solaris operating systems. See [Overview of Oracle Solaris Boot Environments](#) for more information.
- **Jobs:** Displays current and historical job information for that operating system. See [Chapter 3, "Jobs"](#) for more information.
- **Service Request:** Displays the service request for the selected operating system. See [Oracle Services and Service Requests](#) for more information.
- **Configuration:** Displays access points, or resources, that are associated with the operating system. See [Using Access Points](#) for more information.

Operating System Profiles

The following categories of OS Management profiles are available:

- OS Provisioning
- Monitoring Profile


- Oracle Solaris Zone
- Boot Environments

Using Agent Management for Operating Systems

Two management modes are available, agent-managed and agentlessly-managed. An agent managed operating system has an Agent Controller installed to gather information for the Enterprise Controller. Some OS-specific features require an Agent Controller. When you do not want an Agent Controller installed, you can agentlessly manage your operating systems. To gather information on an agentlessly managed operating system, the Proxy Controller uses SSH to perform certain tasks and periodically check on the operating system.

The current management mode of an operating system appears on the Dashboard for the operating system, as shown in [Figure 12–6](#) and [Figure 12–7](#).

Figure 12–6 Agent Managed



When you install the Agent Controller on the operating system, the following actions occur:

- The software registers the Agent Controller with the Enterprise Controller. It takes at least five (5) minutes for the software to register the Agent Controller. After registered, you can update the operating system.
- The software sends you a notification when it has enabled the update function for the operating system.
- The Agent Controller checks the inventory of patches and packages and creates the System Catalog. The catalog lists the patches and packages, and the versions that are currently installed on the operating system.


Some features are not available when the operating system is managed agentlessly. [Table 12–3](#) shows the information that are available for each management type.

Table 12–3 Information Available for Agent Managed and Agentlessly Managed Assets

Tab or Feature	Agent Managed	Agentlessly Managed
Dashboard	Yes	Yes
Summary	Yes	Yes
Libraries	Yes	No
Storage	Yes	No
Utilization	Yes	Yes
Analytics	Yes	Limited
Virtualization Analytics for Oracle VM Server	Yes, if the guest is agent-managed	No

Table 12–3 (Cont.) Information Available for Agent Managed and Agentlessly Managed

Tab or Feature	Agent Managed	Agentlessly Managed
Virtualization Analytics for Oracle Solaris 10 Zones	Yes, if the global zone is agent-managed or if the non-global zone is agent-managed.	No
Networks	Yes	No
Incidents	Yes	Yes
Monitoring	Yes	Yes
Charts	Yes	Yes
Reports	Yes	No See Chapter 10, "Reports"
System Catalogs	Yes	Oracle Solaris 11: Yes Oracle Solaris 8-10, Linux, Windows: No
Terminal	Yes	No
Jobs	Yes	Yes
Configuration	Yes	Yes
OS update	Yes	Oracle Solaris 11, Windows: Yes Oracle Solaris 8-10, Linux: No
OS provisioning	Yes	No
Zone management, including create, read, update, and delete zones	Yes	No

Figure 12–7 Agentlessly Managed

You can use the following methods to change the agent management mode:

- Unmanage the asset, then rediscover the operating system using a profile with the alternative mode.
- Use the Switch Management Access feature.

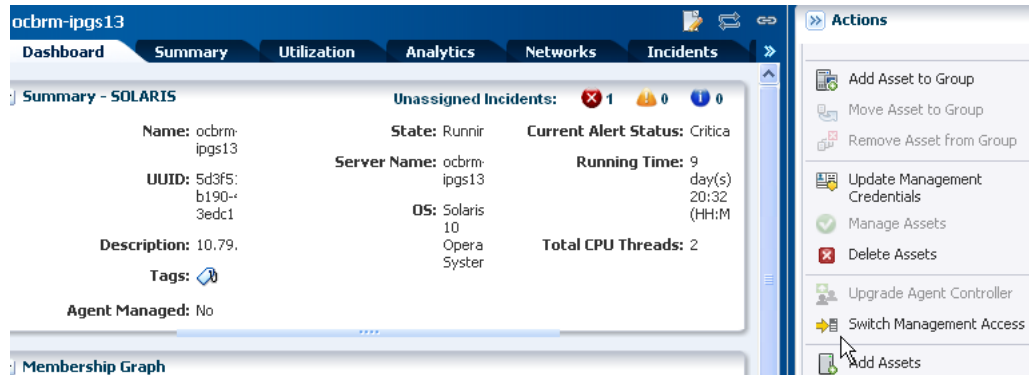
The Switch Management feature enables you to move back and forth from agentlessly managed to agent managed. When the operating system is agent managed and you use this action, the software removes the Agent Controller and changes the operating system to agentless. Select or create new credentials for the Proxy Controller to use to obtain information from the asset.

To Use Switch Management Access

1. Expand Assets in the Navigation pane.

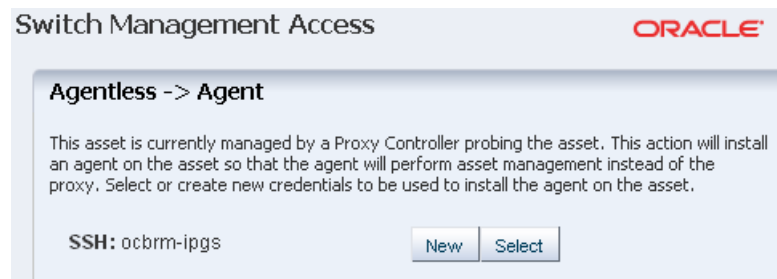
2. Expand the Assets tree, then select the operating system.
The current management status appears in the Dashboard tab.
3. Click **Switch Management Access** in the Actions pane.

Figure 12–8 Switch Management Access



4. Add or select the credentials for the system, then click **Finish**.
 - (Optional) To create a new set of credentials, click **New** and complete the Create Credentials Wizard, then click **OK**.
 - (Optional) To select from a list of existing credentials, click **Select**, highlight the credentials from the list of available credentials, then click **OK**.

Figure 12–9 Switch Management Access Credentials



Monitoring Operating Systems

The software monitors the status of your operating systems right after asset discovery. Oracle Enterprise Manager Ops Center uses the default rules and thresholds. Monitoring rules state the values and boundaries for an asset's activity. A monitoring policy is a set of rules. You can change the rules or thresholds to adjust the type and level of monitoring you want. Analytics provides details on operating system activity and utilization.

A monitoring policy defines alert configurations to be performed on one or more managed resources. A policy is for a specific type of resource, such as operating systems. A more specific policy might apply to all Oracle Solaris operating systems. Each monitoring policy contains several alert monitors for a specific type of resource. Alert monitors watch the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

Applying a monitoring policy to all the assets enforces consistency. Each monitoring policy contains rules for threshold levels. Default policies for monitoring hardware,

operating systems, and Oracle Solaris Clusters are included in the software. You can use the default policies, but you cannot edit them. To edit or add monitoring rules to a monitoring policy, you must make a copy. Modifying a monitoring rule for a specific asset creates a custom set of monitoring rules for the asset.

See [Chapter 4, "Monitoring Rules and Policies"](#) for information about how to change the threshold limits and how to change the deactivate or activate the auto delete policy. See [Chapter 9, "Incidents"](#) for details about how incidents are generated, severity badges, how to assign and close an incident.

Using Analytics

The Analytics feature provides a view into operating system and zone performance and status. The charts, reports, and utilization data provide details of an individual eligible operating system or zone. You can use the information to analyze the behavior of an operating system to aid in peak performance and to diagnose and correct incidents.

Monitoring uses the Agent Controller to gather information. When an operating system is not an agent-managed system, or it is Microsoft Windows, the software uses an SSH connection from a Proxy Controller to perform remote monitoring. Remote monitoring over an SSH connection limits the available metric information. The Summary view, Process view, Historical view and Virtualization Analytics are not available for these operating systems.

The Analytics view provides information for the following agent-managed platforms:

- Linux
- Oracle Solaris 11 and 10 OS
- Oracle Solaris 10 non-global zone, when the global zone or non-global zone is agent-managed
- Oracle Solaris 11 non-global zone, when the global zone is agent-managed

When your operating system is agentlessly-managed, information is not available for zones and less information is available for Linux and Oracle Solaris 10 and 11 operating systems. [Table 12–4](#) shows a list of features and whether the feature is supported on an agent-managed or agentlessly-managed asset.

Table 12–4 Supported Analytics Information

Feature	Supported on Agent Managed OS	Supported on an Agentlessly Managed OS
Summary	Yes	Yes
Virtualization Analytics for Oracle VM Server	Yes, for agent-managed guests	No
Virtualization Analytics for Oracle Solaris 10 Zones	Yes, for an agent-managed global zone or non-global zone	No
Services	Yes	Yes
Processes	Yes, for an agent-managed guest, global zone, or non-global zone	No
Threshold	Yes	Yes
History	Yes, for an agent-managed guest, global zone, or non-global zone	No

Table 12–4 (Cont.) Supported Analytics Information

Feature	Supported on Agent Managed OS	Supported on an Agentlessly Managed OS
Metrics	Yes	Yes

The current management mode appears in the Dashboard Summary page for the operating system. See [Using Agent Management for Operating Systems](#) for details and for the steps to switch the management mode.

Displaying Analytics Information

For each operating system, the information appears in the **Analytics** tab in the center pane.

The following diagnostic pages are available for an agent-managed operating system:

- OS Analytics view: Displays analytics details for an operating system.
 - System resource graphs
 - Top-consumers views
 - Processes information
 - View historical data and set thresholds
- Virtualization Analytics view: Displays analytics details for a zone.
 - Virtualized OS instance
 - Breakdown of resource usage of the physical server
 - Running OS instance
 - View historical data and set thresholds

To Display the Analytics View

1. Select an operating system, zone, or guest in the Navigation pane.
2. Click the **Analytics** tab.

To Display the Both the OS Analytics and Zone Analytics Views

1. Select a global zone that has zones in the Navigation pane.
2. Click the **Analytics** tab to see the OS Analytics.
3. Click the **Zones Analytics** tab to view the Summary, Zones, History, and Resource Pool Utilization for all zones of the global zone.
4. See [Displaying Virtualization Analytics](#) for information about analytics for virtualized platforms.

The Analytics view contains current and historical information about an operating system's use of resources, including CPU, Network, disk I/O, memory, alert history, and total thread and process counts. Information presented in this section assumes that the operating system is agent managed. See [Table 12–4](#) for the list of supported features for an agentlessly managed OS.

The information appears in a series of charts in the following subtabs:

- **Summary:** View a high-level overview of the top consumers, by process, for the CPU, memory, network, and disk I/O resources. See [Displaying the Analytics Summary](#) for details.
- **Processes:** View process-specific details. See [Displaying the Processes View](#) for details.
- **Services:** View and monitor Oracle Solaris 10 and 11 SMF services. See [Displaying the Services View](#) for details.
- **Thresholds:** View and edit the threshold limits for a selected monitored attribute. See [Displaying Thresholds](#) for details.
- **History:** View a history for the top consumers. See [Displaying Historical Data](#) for details.
- **Metrics:** View specific details an operating system element, such as the percentage of memory used. See [Displaying Metrics](#) for details.
- **Charts:** You can create a variety of utilization charts, define the coordinates, and export the chart data to CSV or XML output. See [Displaying and Creating Charts](#) for details.

You can view analytics information for all supported and managed operating systems, both physical and virtual. The information presented varies, depending on the data available for the OS. For example, when the System Tap is not installed on a Linux operating system, information about the top network and I/O consumers might not be available. Also, the top consumers for I/O data might not be available with some versions of the Linux kernel.

When an agent-managed operating system is on a virtualization platform, you can view the information in the Analytics view of the virtualization platform. See [Displaying Virtualization Analytics](#) for information about analytics for virtualized platforms.

Displaying the Analytics Summary

The Summary view provides details for an individual operating system. The Summary provides a graphical representation and a list of the top five (5) resource usage consumers, by process, for the CPU, memory, network, and disk I/O resources. Click the **View** icon for a row to display process details.

The following details appear in the Summary view:

- **CPU Utilization:** A graph of the percentage of CPU Utilization over time and a detailed list of the top five (5) CPU processes. The list includes the process identification number (PID), name, and CPU Usage %.
- **Memory Utilization:** A graph of the top 5 memory processes, including the PID, the process name, and the percentage of memory used.
- **Network Utilization:** A graph of the top 5 network processes, including the PID, the process name, and the network usage in Kilobytes (KB) per second.
- **I/O Utilization:** A graph of the top 5 I/O processes, including the PID, the process name, and the I/O usage per second.

Oracle Enterprise Manager Ops Center collects information every five minutes on every managed asset and displays the last hour of data on the Summary tab. Each list has an icon before the PID column for each process. Click an icon to view details of the process. The type of resource determines the available details.

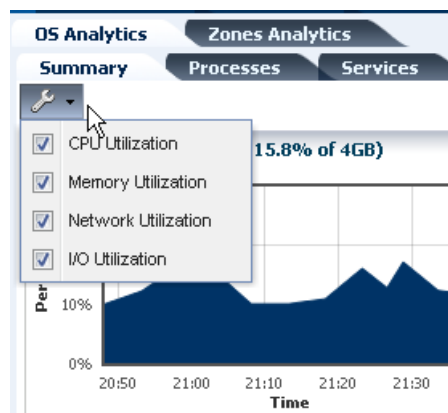
- **Process Details:** Includes the contract ID, Service FMRI, creator, elapsed time, project name, project ID, and a process tree.
- **Thread Information:** Includes the thread's light-weight process (LWP) ID, the number of threads (NLWP), the user, the priority, the state (such as sleeping), the percentage of CPU used, the CPU time, the percentage of memory used, how much memory the process has marked for allocation (VSZ memory), the processors, and the command.
- **Handles:** Includes the port details, such as the family, local address, local port, the protocol, remote address, remote port, device, and node.
- **Process Environment:** Includes details about the environment, such as the arguments, data model, and flags.
- **Memory Information:** Includes physical, virtual, and anonymous memory and dirty page details. Details include the virtual address and the number of KB in the virtual mapping size, the resident physical memory, the anonymous memory, and the dirty pages. The lock status, permissions, and mapping name details are displayed, when available.

See [Displaying and Creating Charts](#) to create a chart with different values for the x and y axis of the graph and to export the chart data.

Displaying and Configuring Graphs

When the icon that looks like a wrench appears on a page with graphs, click the icon to configure the graphs or change the graphs that appear on the page. [Figure 12–10](#) shows an example where you can choose the system resource graphs that display on the Summary page.

Figure 12–10 *Configure Graphs Menu*



Displaying the Processes View

You can drill down to display process-specific data, based on current data from the operating system. Some data, such as CPU usage, might also be available in the History view.

The following details are available in the Process view:

- Process ID (PID)
- Process Name
- User

- State
- CPU Usage %
- Memory Usage %
- Physical Memory size in MB
- Virtual memory size
- Target

When you click a process in the Processes table, two icons are enabled in the center pane, one to view more details and the other to kill the process.

[Figure 12-10](#) shows an example where you can click check boxes to select or deselect system resource graphs from displaying on the Summary page.

When the software is configured to work with Oracle Enterprise Manager Cloud Control, information about available Oracle Enterprise Manager Cloud Control targets appears in the Process view.

Displaying the Services View

The Services tab provides a view of Oracle Solaris 10 and 11 Service Management Facility (SMF) services. The Oracle Solaris SMF feature defines the relationships between applications, or services, and is a method of managing them by providing a framework for startup scripts, `init` run levels, and configuration files. Each service is identified by a Fault Managed Resource Identifier (FMRI).

The Services tab contains SMF service instances, state, dependencies, and severity information. You can drill down to see specific service details, including the configuration, dependencies, and the processes that are in the service contract.

The following actions are available from the Services tab:

- View a list of services currently installed and their states
- View a list of dependencies and dependents for FMRI
- View the relationship between services and processes
- View details about why a service is not available
- Obtain logs for debugging
- Clear faults for FMRI
- Invoke the disable, enable, and restart actions on FMRI
- Read configuration files

Services information is available for agent managed and agentlessly managed Oracle Solaris 10 and 11 operating systems, including global and non-global zones.

Displaying Thresholds

The Thresholds tab contains information about all monitored attributes for the selected operating system, including per-instance attributes such as File System Usage for each file-system on the asset.

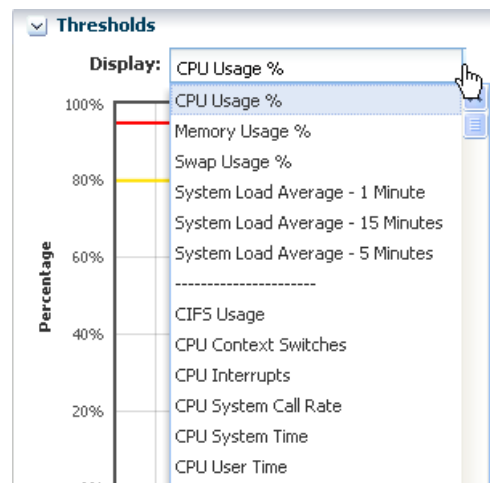
To Display Thresholds

1. Display the Analytics view.
2. Click the **Thresholds** tab.

3. Click the **Display** menu to show the available monitored attributes.

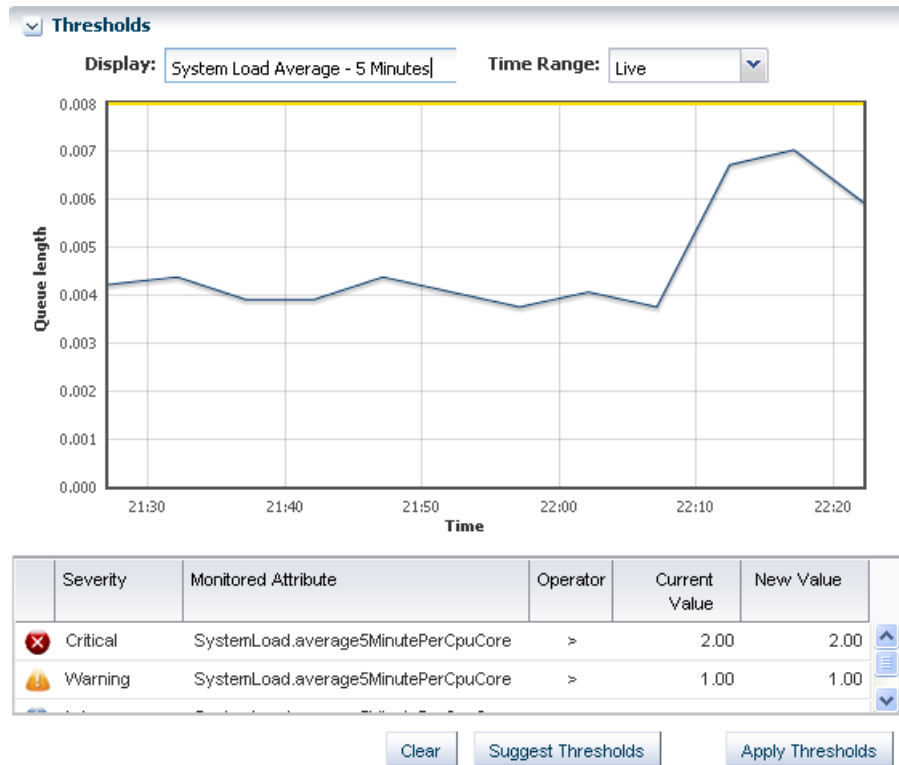
Figure 12–11 shows a partial list of the monitored attributes available for display.

Figure 12–11 *Thresholds Display Options*



Charts for the historical data of those attributes show the alert monitor threshold levels, if configured, for that attribute. Figure 12–12 shows the chart for the System Load Average - 5 minutes. The time frame selected is Live. Under the chart are the severity levels for this monitored attribute (`SystemLoad.average15MinutePerCpuCore`), the operator, and the threshold values. Instead of editing the thresholds in the Monitoring tab, you change the values on this page. You can either add values in the New Value column or click **Suggest Thresholds** to populate the fields with suggested values. Click **Apply Thresholds** to change the existing thresholds.

Note: When the threshold monitor is modified, the asset is removed from the default monitoring profile and a custom profile created.

Figure 12–12 Thresholds Chart


Displaying Historical Data

The History tab contains the history for the top consumers. By default, the current date displays. The time selection box is used to select a time of day for which the top consumer processes are listed. It does not affect the chart, which can only show discrete intervals of 1 hour, 1 day, 5 days, 3 weeks, 6 weeks, and 6 months. The history chart displays the most applicable chart interval for the selected date and time.

[Figure 12–13](#) is an example of a Top Consumer History chart.

To Display Historical Data

1. Display the Analytics view.
2. Click the **History** tab.
3. Click the wrench icon to select from a list of options to chart.

Figure 12-13 Top Consumer History

Displaying Metrics

The Metrics tab provides you with specific details about various operating system statistics, such as the percentage of memory used, and view graphs. When monitoring thresholds are enabled for an element, you can reset the thresholds. When an element does not have a monitoring threshold, you can configure a new monitoring threshold.

To Display Metrics

1. Display the Analytics view.
2. Click the **Metrics** tab to see OS-specific components.


The following are the categories of component details available for an agent-managed Oracle Solaris OS:

- Load
- File Systems
- Networks
- Users
- Buffer Activity
- Disk Usage
- Paging Activity
- Message Activity
- Tables Status
- TTY Activity
- Kernel Memory
- DNLC
- IPC Message Queue
- IPC Shared Memory
- IPC Semaphore Usage
- CPU Detail
- File Access
- Disk Errors

- Memory Utilization
3. Expand a component to view the available elements.


Click an element to view the details. In [Figure 12–14](#), the Load component is expanded and the elements, such as percentage of memory used, CPU utilization percentage, CPU core count, and system update options are available in the Load list. Memory Used % is selected and the details and a graph are visible.

Figure 12–14 Component Navigation



The type of resource determines the available details. [Figure 12–15](#) is an example of the details and graph for the Memory Used % element. This example shows a system in trouble. The Details shows the last known value memory usage at 95.59111 percent and the established monitoring parameters. A warning incident is generated at 80 percent and a critical incident is generated when the used memory reaches 95 percent.

View the bar graph to see how the memory usage has trended over the selected time period. The yellow and red horizontal lines indicate the warning and critical thresholds.

Figure 12-15 Memory Used % Details and Graph

The wrench icon appears above the graph. For this graph, you can click the icon to change the graph style to line, bar, or area. You can also change the definition of the X axis from 1 day to 1 hour, 5 days, 3 weeks, 6 weeks, or 6 months. The export icon enables you to export the graph to file in CSV or XML format.

You have the opportunity to reset the threshold values for the monitoring rule from this page, or to view recommended threshold values based on past performance. When you click the Set Thresholds button, you are taken to the Thresholds tab, where you can enter new threshold values. When the Thresholds tab appears, the drop-down menu contains the name of the monitored attribute from the metrics view, and the associated chart appears. Editing a threshold value creates a custom set of monitoring rules for the asset. See [Chapter 4, "Monitoring Rules and Policies"](#) for more information about editing individual monitoring rule parameters and creating a custom monitoring policy.

Displaying and Creating Charts

The Charts tab enables you to modify the following charts to change the type of chart and to see utilization data over longer periods of time, up to six months:

- Power Utilization: Servers, chassis

- CPU Utilization: Operating system, operating system for a virtualization host, virtual host, server pool
- Memory Utilization: Operating system, virtual host, server pool
- Network Utilization: Operating system, operating system for a virtual machine, virtualization host, server pool
- File System Utilization: Oracle Solaris OS and Linux OS
- System Load: Oracle Solaris OS and Linux OS

For the first five days of operation, Oracle Enterprise Manager Ops Center collects data every five minutes. After the fifth day, the reported data is an average, according to the following:

- Five days to three weeks: Average for each hour
- Three weeks to six weeks: Average for each 12-hour period
- Six weeks to six months: Average for each 24-hour period

Oracle Enterprise Manager Ops Center deletes the data after six months.

To Display Charts

1. Display the Analytics view.
2. Click the **Charts** tab.

You can create charts 24 hours after you first manage the asset. The asset must be operating and the Enterprise Controller must be able to get access to the asset. You can use the default format or change the charts to use a line, bar, or area format or to use different time intervals.

To Create Charts

1. Display the Analytics view.
2. Click the **Charts** tab.
3. Select an option from the Chart type menu to display the output as a line, bar, or area chart.
4. Select **Live**, **1 Hour**, **1 Day**, **5 Days**, **3 Weeks**, **6 Weeks**, or **6 Months** from the X Axis (Time) menu to change the timeframe for the data.
5. Select **Percentage** or **Unit of Measure** to configure the chart's Y Axis.
6. Select the component and options to plot from the Plot selection.

By default, the graphs show one day of data. Select **Live** in the X Axis menu to change to Live mode, which reports new information every 5 seconds. You can also change the graph to one of the following periods:

- One hour (1H): One point for every 5 minutes
- One day (1D): One point for every 5 minutes
- Five days (5D): One point for every 5 minutes
- Three weeks (3W): One point every hour
- Six weeks (6W): One point every 12 hours
- Six months (6M): One point for every day. To make a graph with the minimum of two points, a system must have been managed for at least 10 minutes to view a one-hour graph and for at least two days to view the six-months graph.

The software stores the data for these time periods separately. For example, when a server is managed for two hours and you select the 6W view, the graph does not display because only one point of data of that type is available; the second point has not yet occurred. When you select the 1D view, the graph displays 24 points of data (120 minutes in 5-minute intervals). However, the graph displays these points over a 24-hour period and not over the actual two-hour period. For the most accurate representation of the data, choose a time period that is less than or equal to the time that the selected server has been managed.

You can export the data for either the current view or all available data to a file in either CSV or XML format. Click **Export Chart Data** to choose options for exporting the data.

When the graph is blank for a server, one of the following conditions has occurred:

- Server does not have the appropriate ILOM version.
- Server has not been discovered through the ILOM driver.
- Server is unreachable.

Displaying Virtualization Analytics

The Virtualization Analytics view displays resource usage of the physical server for each running operating system instance, showing the physical resources consumed by the control domain (global zone or Oracle VM Server), and each non-global zone or guest. Metrics for Oracle VM Server for x86 are available through the Oracle VM Manager.

- Virtualization Analytics Summary
- Virtualization Analytics Zones or Virtualization Analytics Guests
- Virtualization Analytics History

The information is refreshed every 20 seconds for a guest running on a virtualization server container, including global zones, Oracle VM Server for SPARC, and Oracle VM Server for x86.

Customizing the Analytics View

The default display is the resource usage data. You can change the display settings for each view and choose which system resource graphs and details to display. The changes made on a given analytics page affects the views of all analytics screens for the same OS or virtualized OS.

To Customize the Analytics View

1. Select an operating system (global zone) or zone in the Navigation pane.
2. Click the **Analytics** tab.
3. Click the **Settings Menu** icon in the center pane.
4. If you select a global zone that has zone, click the **Zones Analytics** subtab, then click the **Settings Menu** icon in the center pane to change the view for the Zones Analytics page.

Provisioning Operating Systems

The provisioning feature provides a method of automatically and consistently installing operating systems on managed systems from the Oracle Enterprise Manager Ops Center user interface.

You can provision the following:

- Linux
- Oracle Solaris (SPARC or x86)
- Oracle VM Server for SPARC
- Oracle VM Server for x86
- Oracle Solaris Clusters

Note: To provision Oracle Solaris 11, both the Enterprise Controller and Proxy Controller must be installed on a system that is running Oracle Solaris 11.

To provision Oracle Solaris 9 and 10, you can use a Proxy Controller that is installed on a system that is running either Oracle Solaris or Linux.

To provision Oracle Solaris 10 using JET customization, the Enterprise Controller must be installed on a system that is running an Oracle Solaris operating system.

Table 12–5 shows the actions that you can perform based upon which operating system that the Enterprise Controller and Proxy Controller are installed.

Table 12–5 Provisioning Actions Determined by the Operating System on Which the Enterprise Controller and Proxy Controller are Installed

Action	Enterprise and Proxy Controllers on Oracle Solaris 11	Enterprise and Proxy Controllers on Oracle Solaris 10	Enterprise and Proxy Controllers on Linux
OS Provisioning Oracle Solaris 11	Supported	Not Supported	Not Supported
OS Provisioning Oracle Solaris 10	Supported	Supported	Supported
	JET Supported	JET Supported	JET not Supported
OS Provisioning Linux	Supported	Supported	Supported
Provisioning Oracle VM Server for SPARC	Supported	Supported	Enterprise Controller – Supported
			Proxy Controller – Not Supported
Provisioning Oracle VM Server for x86	Supported	Supported	Supported

For information about provisioning virtual machines, see [Chapter 15, "Oracle VM Server for SPARC"](#) and [Chapter 16, "Oracle VM Server for x86"](#).

Automated OS provisioning is performed through a managed service processor. You can provision the OS on a single system, one or more groups of systems, or a combination of systems that are attached to your network.

OS provisioning is handled by a profile and an associated deployment plan. The profile defines the image and the plan defines the tasks to perform and the targets to provision. Some plans are simple single step plans, others are multi-step plans where OS provisioning is a step in a larger plan.

When you run an OS provisioning job on a SPARC machine, Oracle Enterprise Manager Ops Center resets the configuration to the factory default configuration and removes the user-defined commands that are executed during start-up and that are stored in the NVRAMRC file in the non-volatile RAM (NVRAM). The Control Domain OS provisioning profile does give you the option to preserve the information in the NVRAMRC file. See [Creating an OS Provisioning Profile](#) for more information about provisioning a Control Domain.

You have many options available when you provision an operating system. The following steps are needed to complete a basic provisioning job:

1. Determine the method of importing the OS image files, with either Dynamic Host Configuration Protocol (DHCP) or a WAN boot.
2. Import a file with the OS image into a Software Library. A provisioning profile is automatically created when you import the OS image. You can use this profile, edit the profile, or create a new profile. See [Images](#) and [Images for OS Provisioning](#) for details on importing images. See [Overview of OS Provisioning Profiles](#) for the types of profiles you can create.
3. Create or configure a deployment plan that includes your profile.
4. Manage the service processor for one or more systems to provision. See [Discovering and Managing Assets](#) for details.
5. Apply the deployment plan on the targets. When you choose a group as a target, it must be an homogeneous group where all members are the same.

After you import the OS images and create the OS profiles, you can reuse the profiles in a variety of plans that have OS provisioning as a step.

About Dynamic Host Configuration Protocol (DHCP)

DHCP dynamically assigns IP addresses to devices on a network. A typical OS provisioning job requires an installation server and a DHCP server on the same subnet as that of the client systems. A JumpStart boot server must be on the same subnet as that of the client systems.

Before you can provision, you must configure DHCP services on the Proxy Controllers. You can use basic DHCP services, with or without defined subnets, or an external DHCP server. See the *Administration Guide* for information about how to configure DHCP and subnets for OS provisioning.

About WAN Boot

The WAN boot installation method enables you to boot and install software over a wide area network (WAN) by using HTTP. By using WAN boot, you can install the Solaris OS on SPARC based systems over a large public network where the network infrastructure might be untrustworthy. You can use WAN boot with security features to protect data confidentiality and installation image integrity.

WAN boot is the default connection for Oracle Solaris 11 provisioning. Oracle Solaris 11 provisioning does not use a Flash Archive (FLAR) image.

Oracle Solaris 10 provisioning can use a WAN boot or DHCP connection. With a WAN boot connection, Oracle Solaris 10 provisioning enables you to provision a FLAR

image on a SPARC system using an HTTP web server. WAN boot installation is useful when DHCP does not meet your organization's security policies or you have SPARC-based systems that are located in geographically remote areas and you need to install servers or clients that are accessible only over a public network. Because WAN boot uses an HTTP server, it works across your corporate firewall and does not require DHCP or a JumpStart boot server to be on the same network as the client systems.

The WAN boot installation method uses port 5555 and HTTP to boot and install software on SPARC-based ILOM, ALOM, or M-series systems over a wide area network (WAN). The WAN boot security features protect data confidentiality and installation image integrity over a large public network where the network infrastructure might be untrustworthy. You can use private keys to authenticate and encrypt data. You can also transmit your installation data and files over a secure HTTP connection by configuring your systems to use digital certificates. See [Security Configurations Supported by WAN Boot](#) in the *Oracle Solaris 10 10/09 Installation Guide: Network-Based Installations* document for more information about secure WAN boot installation configuration.

Requirements for a WAN Boot Connection

Oracle Solaris 11 and 10 WAN Boot require that *devalias net* is aliased to the proper device. Oracle Enterprise Manager Ops Center does this for you when you use the *GB_X/netY* syntax in your OS Profile when specifying the NIC. See [Adding or Re-Aliasing a Device for WAN Boot](#) for more information about how Oracle Enterprise Manager Ops Center handles device aliases.

The following are required to use WAN boot with Oracle Enterprise Manager Ops Center:

- Oracle Solaris 11
 - The target is a SPARC ALOM-CMT, ILOM-SPARC or M-Series platform that has a supported OBP or XCP. See [Checking Your Client OBP for WAN Boot Support](#) for how to check your client OBP. For M-Series, the XSCF Control Package (XCP) file should be at least version 1082. The XCP file contains the hardware's control programs and includes the XSCF firmware and the OpenBoot PROM firmware.
 - The Enterprise Controller is installed on an Oracle Solaris operating system. You can use a SPARC or x86 platform for the Enterprise Controller.
 - WAN boot is enabled for Oracle Solaris 11 in Administration. See [Disabling and Enabling WAN Boot](#).

Note: When the target does not have the required OBP firmware version, the Oracle Solaris 11 provisioning profiles do not appear in the target list for the server.

- Oracle Solaris 10
 - The target is a SPARC (ALOM-CMT, ILOM-SPARC or M-Series) and has the minimum OBP or XCP version
 - A FLAR image is used
 - The Enterprise Controller is installed on an Oracle Solaris operating system. You can use a SPARC or x86 platform for the Enterprise Controller.

- WAN boot is enabled for Oracle Solaris 10 in Administration. See [Disabling and Enabling WAN Boot](#).

Note: When the target does not have the required OBP firmware version, Oracle Solaris 10 provisioning reverts to a DHCP connection, or OBP/XCP versions + if an ISO is used.

To determine if your client system has a WAN boot-enabled PROM, see [Checking Your Client OBP for WAN Boot Support](#).

Checking Your Client OBP for WAN Boot Support

To determine if your client system has a WAN boot-enabled PROM, perform the following to check the client Open Boot PROM (OBP) for WAN boot support.

1. Log in to a terminal window as root.
2. Enter the following to check the OBP configuration variables for WAN boot support:


```
# eeprom | grep network-boot-arguments
```
3. The OBP supports WAN boot installations when the variable *network-boot-arguments* appears, or when the command returns the output *network-boot-arguments: data not available*. For example:


```
# eeprom | grep network-boot-arguments
network-boot-arguments: data not available
```
4. If the command in Step 2 does not return any output, the OBP does not support WAN boot installations. Use [Firmware Provisioning](#) to update the OBP to the required level.

Setting Up a WAN Boot Connection

When Oracle Enterprise Manager Ops Center is installed on an Oracle Solaris operating system, the Enterprise Controller is automatically configured to be a WAN boot server.

Oracle Solaris 11 uses WAN boot. For earlier versions of Oracle Solaris, WAN boot is the default connection for provisioning when the requirements are met and you choose to use a FLAR image. When you launch an OS provisioning on an eligible SPARC-based system and you choose a FLAR image, the software automatically uses WAN boot. If you have a group of systems to provision, the software determines whether to use WAN boot or DHCP for each system.

Disabling and Enabling WAN Boot

WAN boot is automatically installed and enabled when the Enterprise Controller is running on an Oracle Solaris operating system. You can check the status, disable, or enable WAN boot in the Enterprise Controller configuration file. WAN boot is enabled when the value for the property is set to true in the Enterprise Controller configuration file. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about how to check the status, disable or enable the WAN boot connection.

1. Expand the **Administration** section in the Navigation pane, then click **Enterprise Controller**.
2. Click the **Configuration** tab.

3. Select **OS Provisioning** from the Subsystem menu.
 - For Oracle Solaris 11, see the following property: `usesS11WANBoot`.
 - For Oracle Solaris 10, see the following property: `usesS10WANBoot`.
4. (Optional) To disable WAN boot, change the value for the property to **false**.
5. (Optional) To enable WAN boot, change the value for the property to **true**.

Adding or Re-Aliasing a Device for WAN Boot

Oracle Solaris 11 and 10 WAN Boot require that *devalias net* is aliased to the proper device. The default uses the first NIC. Some option-card NIC devices, such as 10 GbE, are not automatically aliased. To add a net entry, such as `net4`, you must re-alias net to the right side of the net device that you want.

Oracle Enterprise Manager Ops Center performs a re-alias if you use the *GB_X/netY* syntax when you specify the NIC in the OS Provisioning profile. For example, in [Figure 12-16](#), the default *GB_1* value is changed to *GB_1/net4*.

Figure 12-16 Network Interface Page

Network Interfaces (2)				
<div> + × </div>				
Network	VLAN ID	NIC	Boot	Address Allocation Method
10.79.204.0/24.1	-	GB_0	<input checked="" type="radio"/>	Use Static IP
10.79.204.0/24.1	-	GB_1/net4	<input type="radio"/>	Use Static IP

Note: For the software to re-alias the device, you must use the syntax *GB_X/netY*.

When you have more *GB_X* than there are *devalias netX* devices showing, you must manually re-alias the net. For example, the factory default for a T3 might only show *devalias net0* through *net3*. You can define additional *devalias*es for higher numbered *netX* devices.

You can directly alias net to point to one of those higher devices, if you use *nvalias*; *nvstore*. A more useful option is to use *nvalias*; *nvstore* for the higher numbered *netX* aliases, which enables Oracle Enterprise Manager Ops Center to perform the *devalias* procedure. For example, use *nvalias net10* as the *netdevice* path for first option-card NIC and *nvalias net11* as the *netdevice* path for second option-card NIC. When you use *nvalias net10*, you can use *GB_10/net0* in the UI and the software performs the *devalias* procedure for the higher numbered devices. The software performs the *devalias* procedure based on what you entered for *net10* and sets that on net for you before issuing the `boot net - install` command.

Importing Images

Import your OS images into the Software Library before you can start a provisioning job. You can save the images in a Local Software Library on the Enterprise Controller or in a Network Attached Storage (NAS) Software Library that is created on an NFS

server that the Enterprise Controller can access. See [Creating a Software Library](#) for details.

Overview of OS Provisioning Profiles

The provisioning profiles collect all the information such as type of target, OS image, time zone and language setup, required JET modules, disk partitions, naming services and network details. Depending on the type of operating system you are provisioning, include its parameters in the profile that installs the operating system.

When you specify the OS Parameters, you have the option to include custom scripts. This feature is disabled for Oracle Solaris 11 profiles.

Specifying Common Linux Parameters

Specify the following parameters:

- Installation number: The number that enables you to install all of the Linux software that is included in your subscription.
- Partition action: Use this parameter when you want to change the disk partition of the system.
 - You can opt to remove all the existing Linux partitions and retain the non-Linux partitions. You can provide specification for the new partitions.
 - You can opt to preserve all the existing partitions. You must define new partitions, outside of the partitions that exist, in which to install the OS.
 - You can opt to remove all the existing partitions. Define specification for the new partitions.
- Install protocol: Specify HTTP or NFS as the install protocol.
- Kernel parameters: Enter kernel parameters for the GRUB menu of the target system, when needed.
- MD5 Checksum: Select this option to use MD5 encryption for user passwords.
- Reboot action: Select whether you want to reboot the target system after OS installation.
- Disk label initialization: Select this option to initialize labels on new disks. This option creates labels that are appropriate for the target system architecture.
- Shadow passwords: Select this option to use an `/etc/shadow` file to store passwords on the target system.
- Clear master boot record: Select this option to clear all invalid partition tables.
- Linux packages: You can specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude any package, enter the package name preceded by a dash (-).

Specifying SuSE Parameters

To create a profile that installs the SuSE Linux OS, specify the following parameters:

- FTP proxy server: Enter the name of the FTP proxy server to support FTP services.
- HTTP proxy server: Enter the name of the HTTP proxy server to support HTTP services.
- Install protocol: Specify HTTP or NFS as the install protocol.

- Enable proxy servers: Select this option to enable the FTP and HTTP proxy servers that you specified in the FTP Proxy Server and HTTP Proxy Server fields.
- Kernel parameters: Enter kernel parameters for the GRUB menu of the target system, when necessary.
- Reboot action: Select whether you want to reboot the target system after OS installation.
- Linux packages: You can specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude any package, enter the package name preceded by a dash (-).

About Oracle Solaris or Linux OS Provisioning Profiles

Each profile is defined by the OS image that is either in the Software Library, or is imported into the Software Library. When you import an OS image into the software library, a default provisioning profile is automatically created. You can edit the default profile or you can create a new profile.

For each profile, you specify the setup parameters, file system parameters, the management mode, and the naming service for the operating system. For Oracle Solaris only, you can optionally use JET Modules to specify additional Installation Parameters. To learn more about JET and how to use JET in a profile, see [About JumpStart Enterprise Toolkit \(JET\) for Oracle Solaris](#).

When you choose the agent managed mode, you can perform software updates and create operating system reports. When you choose agentlessly managed, SSH credentials are required to monitor the operating system. You can change the management mode after the OS is provisioned.

A provisioning profile requires a DHCP-enabled network interface for the boot interface. You can add multiple networks, as long as the networks are available and defined in the Enterprise Controller. You can select a NIC from the list of available logical interfaces for each network or, as shown in [Figure 12–17](#), you can use the Address Allocation Method for the selected networks. You cannot use the Address Allocation method for the boot interface. When you use a static IP address, you must provide the IP address when you apply a deployment plan that uses the profile. The IP address is assigned to the target system after provisioning.

Figure 12–17 Select Networks

Select Networks

Select all networks that you want to use on the provisioning process.
VLAN IDs are used in logical domain provisioning, but are disregarded in bare metal OS provisioning.

Network Interfaces (1)

Network	VLAN ID	NIC	Boot	Address Allocation Method
192.168.18.0/24@xvm-x4150 -		GB_0	<input checked="" type="radio"/>	Use Static IP

By default, the profile reboots the OS. You can choose the Manual Net Boot option to enable manual control of network boot operations for the target system.

Note: The Enterprise Controller cannot remotely control the network boot process on systems that do not have a service processor. When your target system does not have a service processor, you must select the Manual Net Boot option.

NVRAMRC File

Automatic booting on a SPARC system uses the default boot device that is defined in the non-volatile RAM (NVRAM). User-defined commands that are executed during start-up are stored in the NVRAMRC file in the NVRAM.

When you run an OS provisioning job on a SPARC machine, Oracle Enterprise Manager Ops Center resets the configuration to the factory default configuration. This process removes the devalias and other information this is stored in the NVRAMRC file on the server.

The Control Domain OS provisioning profile gives you the option to preserve the NVRAMRC file. When selected, Oracle Enterprise Manager Ops Center preserves the information before resetting the server to the factory defaults, and then restores the information after the reset.

About Oracle VM Server for SPARC OS Provisioning Profile

To create a profile that installs an Oracle VM Server Control Domain, define the following parameters:

- CPU Threads
- Memory
- Crypto Units
- Virtual Console Port Range
- Enable JASS

Enabling JASS results in installation of `SUNWJass` package to harden the system when installing Oracle VM Server 1.2 or lower versions. JASS is not supported for Oracle VM Server for SPARC 1.3 or higher versions. When you select Enable JASS in the profile that is installing Oracle VM Server for SPARC version 1.3 or 2.0, the software disregards this option and JASS is not enabled.

- Enable Multiplexed I/O

Enabling multiplexed I/O on a system used for storage enables the Fibre Channel ports on the system.

About JumpStart Enterprise Toolkit (JET) for Oracle Solaris

JumpStart Enterprise Toolkit (JET) provides a framework to simplify and extend the JumpStart functionality provided within the Oracle Solaris 9 and 10 operating system. The `SUNWjet` and `JetFLASH` packages are installed on the Proxy Controller during installation when the Proxy Controller is installed on an Oracle Solaris 10 operating system.

Using JET provides more options for defining the Jumpstart parameters. When you install JET on a JumpStart server, you have the following advantages:

- Install multiple versions of Oracle Solaris
- Deploy flash archives

- Utilize multiple boot methods
- Install recommended patches
- Configure all your network interfaces

Note: You cannot define IPMP groups or link aggregation for a JET template profile.

Table 12–6 describes the JET modules that are installed on the Proxy Controller.

Table 12–6 JET Modules and Associated Packages

JET Module Name	JET Package	Description
base_config	SUNWjet	Provides the standard installation configuration for the client, including the information required to set up the JumpStart server to allow the client to boot and build.
custom	SUNWjet	Adds functionality to the JumpStart framework to handle packages, patches, scripts, and files.
flash	JetFLASH	Adds the ability for the JumpStart server to deliver Solaris images in Solaris Flash format.

See [Appendix D, "JumpStart Enterprise Toolkit"](#) for a list of SUNWjet parameters.

To Create a JET Template

To create a profile that uses the JumpStart Enterprise Toolkit (JET), select a JET template that defines all the parameters for OS provisioning.

Place the JET template on a directory that the Enterprise Controller can access. You can also create a JET template on the Enterprise Controller in the directory `/opt/SUNWjet/Templates`, using the following command:

```
./make_template template_name
```

A sample template is provided. You can make a copy and change the values in the JET template as required. During provisioning, the OS provisioning parameters are read from the template. After you create the JET templates you want and save them on the Enterprise Controller, you can use them in your Oracle Solaris provisioning profiles.

About Provisioning Plans

When you create a profile, the default action creates a deployment plan and associates the plan with the profile. To provision an OS, you must have an OS image. Importing OS images and creating OS profiles are one-time tasks for each OS-target combination. After an OS image and an associated OS profile are created, you can use deployment plan to provision the operating system onto one or more managed systems at the same time.

Note: It is a good practice to place the systems that you are going to provision in Maintenance Mode so that you can take the system offline without generating alerts and incidents.

The following is the general workflow for provisioning an OS:

1. Import a file with the OS image into a Software Library. The import process creates a profile that specifies the name and location of the OS image.
2. Use a deployment plan that includes the profile, either the default one, or one that you create.
3. Apply the deployment plan on the targets. Oracle Enterprise Manager Ops Center enables OS provisioning on single systems, groups of systems, a combination of the single systems and groups, or Oracle Solaris Clusters. OS provisioning for groups of systems requires using homogeneous groups.

To provision an operating system on an Oracle Solaris Cluster, you provision the same operating system on all nodes of a cluster. The Cluster OS Profile profile handles the pre-action and post-action operations. See the Oracle Solaris Cluster documentation and [Appendix B, "Oracle Solaris Cluster"](#) for how to install a new Oracle Solaris Cluster and to maintain existing Oracle Solaris Clusters.

About Deployment Plans That Provision an Operating System

Using the provisioning profiles created for different targets, create a deployment plan with the associated profile. Several multi-step and complex plans use an OS provisioning profile and plan as a step in the larger plan. Provisioning an Oracle Solaris 11 operating system is a different procedure from provisioning earlier versions.

The following deployment plans include a step to provision an operating system:

- Provision an Oracle Solaris or Linux OS
- Provision an Oracle Solaris 11 OS
- Provision and Update Logical Domains
- Provision an OS on an Oracle Solaris Cluster

Provisioning Linux and Oracle Solaris 8, 9, and 10

You can create OS profiles for provisioning Linux OS on x86 systems or Oracle Solaris on x86 or SPARC systems. See [Provisioning Oracle Solaris 11](#) for details on provisioning Oracle Solaris 11.

Planning

Perform the following before you provision the operating system:

- Import the OS image. A default OS provisioning profile for the image is automatically created. Uploading the packages from Oracle to the library can take several hours.
- (Optional) Edit an existing OS provisioning profile or create a new profile.
- Discover the service processors of the target systems.
- Verify that the Dynamic Host Configuration Protocol (DHCP) services are enabled on Proxy Controllers. You cannot create a profile or assign any network if the DHCP services are not enabled. The Install Server option to provision OS on a server is not enabled if the DHCP is not enabled on any of the interfaces.
- Verify that any scripts the profile uses are in a directory that the Enterprise Controller can access. You can save scripts in a local directory of the Enterprise Controller, or in a directory that the Enterprise Controller mounts using NFS.
- When you are provisioning a dynamic system domain of an M-Series server, the domain must have an IP address.

- Disable monitoring for the target systems to prevent events related to a system going offline. See [Maintenance Mode](#) for more information.

Complete the following steps to provision an operating system:

1. Verify that the OS image and profile are available and configured with the parameters you want to use.
2. Discover the service processors of the target systems.
3. Place the asset in [Maintenance Mode](#) to prevent events related to a system going offline
4. Select the OS Provisioning plan, or a plan that uses OS Provisioning as a step.
5. Define the targets for the plan, make any last minute changes in the plan, then submit the job.

Provisioning Oracle Solaris 11

Oracle Solaris 11 uses a new OS provisioning technology, called the Automated Installer. This feature replaces the older JumpStart Enterprise Toolkit technology that Oracle Enterprise Manager Ops Center uses to provision earlier versions of Oracle Solaris.

Oracle Enterprise Manager Ops Center reduces the complexity by using a local copy of the Oracle Solaris 11 Software Library on the Enterprise Controller and creating an Automated Installer server on the Proxy Controller for your use.

Note: To provision an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must both be running on an Oracle Solaris 11 operating system. The repository resides on the Enterprise Controller and the Automated Installer server resides on the Proxy Controller. When the Enterprise Controller and Proxy Controller are not running on Oracle Solaris 11, the Oracle Solaris 11 library and OS provisioning actions are not available.

Oracle Solaris 11 uses the following components for provisioning an operating system:

- Oracle Solaris 11 Software Library: A local version of the software package repository. You can update the Oracle Solaris 11 Software Library, as needed, and then provision multiple systems without using a network connection to Oracle for each provisioning job. See [Chapter 5, "Software Libraries"](#) for more information about adding content to libraries.
- Installation manifest: Defines the system configuration, including what software to install and details on the virtualized environments to provision. A default manifest is included with each Image Packaging System (IPS) software repository.
- DHCP or WAN boot connection
 - x86 client: Requires a DHCP connection.
 - SPARC client: Requires a DHCP or WAN boot connection. Oracle Enterprise Manager Ops Center automatically sets up WAN boot connection.

Note: Oracle Solaris 11 only supports a ISC DHCP Server. Oracle Solaris DHCP Server is not supported.

You can provision Oracle Solaris 11 zones as part of the OS installation. After a system is bootstrapped with a minimized operating system, the operating system is installed from the Oracle Solaris 11 Software Update Library in Oracle Enterprise Manager Ops Center. The zones are provisioned during the initial system reboot after the base operating system is installed.

Oracle Solaris 11 Provisioning Profiles

You can create multiple profiles to respond to subtle variations in hardware attributes, software profiles, or your organization's requirements.

Before you can create provisioning profiles, verify that the Oracle Solaris 11 Software Library is configured on the Enterprise Controller and the package you want is available in the library.

When you create an Oracle Solaris 11 OS provisioning profile, you select the architecture, either SPARC or x86, to display the boot image and distribution for your architecture. You must provide non-root user credentials and root user credentials. You can only use non-root user credentials to login or SSH to the client after install.

The boot image has several associated software groups. You must select one of the following software groups:

- large-server
- small-server
- desktop

The boot image might have additional Software groups that you can select, such as storage-server and NAS attached. You can select one or more of these groups. Each group has a tool tip that describes the group.

To install the OS to an iSCSI disk, select the Use iSCSI Disk option when you create the profile, then specify the iSCSI disk settings. When you use this option, you must provide the following parameters when you deploy the OS Provisioning plan:

- Storage server IP
- SCSI disk LUN

Steps for Oracle Solaris 11 Provisioning Plan

The following are the main steps to provision an Oracle Solaris 11 operating system:

1. Create an Oracle Solaris 11 OS provisioning profile.
2. From the Library section, initialize the Oracle Solaris 11 Software Update Library.
3. From the Plan Management section, create a Solaris 11 x86 or SPARC profile (the Solaris 11 images are pre-imported as part of MSR creation).
4. Target the plan to an x86 server or WAN boot-capable SPARC server.

Defining Network Interface Names for Oracle Solaris 11

The hardware configuration determines whether you need the Oracle Solaris 11 vanity NIC handling, or not. Other operating systems do not require vanity NIC handling.

It is important to use *netY* in the *GB_X/netY* syntax. Including *netY* in the syntax enables you to create names instead of using the Oracle Solaris 11 vanity NIC names, which is useful when the names are not what you might expect. That could happen in the example above where you have a pair of option-card 10 GbE NIC's installed. In

that case, Oracle Solaris 11 might display the two option card NIC's as *net0* and *net1* and the on board NIC's as *net2* through *net5*.

Following that same example, when you want to boot from the *GB_0*, but you want the Oracle Solaris11 manifest to reflect the correct Oracle Solaris 11 vanity NIC name, use *GB_0/net2*.

Selecting an Alternate Network Interface for Provisioning Logical Domains

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 that runs independently on an Oracle VM Server for SPARC Control Domain.

When you provision OS on logical domains, *GB_0* is the default network interface. When you want to select an alternate boot interface other than the primary *GB_0* to provision logical domain OS, you must specify *GB_x* as the corresponding alternate network interface in the **Select Networks** page of the OS provisioning deployment plan. For example, select *GB_3* instead of the default *GB_0*.

Perform the following steps to select an alternate network interface for booting for provisioning OS on logical domains:

1. Expand the **Plan Management** section in the Navigation pane, then expand **Deployment Plans**.
2. Expand **Provision OS**.
3. Right-click the plan to modify.
4. Click **Apply Deployment Plan**.
5. Select the targets and click **Add to Target List**.
6. Proceed through the wizard pages until you reach the **Select Networks** page.
7. Specify *GB_x* in the NIC column to boot over an alternate network interface.
8. Proceed through the wizard pages.
9. Review the selections on the **Summary** page, then click **Apply** to deploy the plan and provision OS on the logical domains.

Overview of Oracle Solaris Boot Environments

Boot environments are a feature of Oracle Solaris. A boot environment is an instance of a bootable Oracle Solaris image plus additional software packages that are installed onto the image, and the set of all file systems and devices (disk slices and mount points) that are required to operate an Oracle Solaris OS instance. You can have disk slices, also known as partitions, on the same disk or distributed across several disks.

A dual boot environment consists of a live, or active, boot environment (BE) and one or more inactive alternate boot environments (ABE). A system can have only one active boot environment, which is the booted environment. An alternate boot environment is an inactive environment that is not currently booted. A system can have many inactive boot environments. You can activate an alternate boot environment at any time.

You can use a dual boot environment within Oracle Enterprise Manager Ops Center to manage your Oracle Solaris software. A dual boot environment is often used to manage updates because it can significantly reduce the service outage time that is usually associated with patching. Maintaining multiple boot environments also enables quick and easy rollback to a version before the patches were applied, if

needed. The boot environment technology enables you to duplicate a boot environment and perform the following tasks without affecting the currently running system:

- Run an Oracle Solaris software update simulation on the inactive boot environment. You can run the simulation with or without downloading the patches.
- Update your Oracle Solaris OS on the inactive boot environment and test the update before deploying it as your active environment.
- Maintain multiple boot environments with different images. For example, you can create one boot environment that contains all current patches and another that contains only security patches.

Understanding the Differences Between Oracle Solaris 11 and Oracle Solaris 10 Boot Environments

The boot environment feature changed beginning with Oracle Solaris 11, including the supported file systems, file system requirements, zone support, and how boot environments are created.

Oracle Solaris 11 Boot Environments use the `beadm` utility and ZFS file systems to create and manage boot environments. You do not need to create boot environments ahead of time. The software creates them automatically. You can use an agent-managed or agentless-managed operating system with the Oracle Solaris 11 Boot Environments feature.

Oracle Solaris 10 and earlier use the Live Upgrade feature with `lucreate` scripts and ZFS or UDFS file systems to create alternate boot environments. You must use an agent-managed operating system with the Oracle Solaris Live Upgrade feature.

The Boot Environments profile defines the boot environment for the OS update deployment plans. The Oracle Solaris release determines how the profile is used:

- Oracle Solaris 11: The profile indicates the policy that is used when an OS update plan is executed. Use this profile to define the creation policies and as a step in an OS update deployment plan.
- Oracle Solaris 10: The profile defines the `lucreate` script that is used to create your alternate boot environments. Use this profile to specify how to create alternate boot environments for the eligible operating systems in your data center, create alternate boot environments, synchronize and activate boot environments, and as a step in an OS update deployment plan.

The user interface assists you in easily navigating the differences in the boot environments features and provides a unified view whenever possible. See [Overview of Oracle Solaris 11 Boot Environments](#) and [Overview of Oracle Solaris 10 Boot Environments](#) for detailed version-specific information.

Monitoring Boot Environments

Monitoring rules and thresholds are defined in the Monitoring tab. The following rules apply to boot environments:

- Number of Boot Environments: Defines the number of boot environments in the selected zone.
- Boot Env Usage Percent in a ZPool: Defines the percentage disk utilization on the boot environments in any of the zpools.

Each rule has two defined thresholds:

- **Warning:** Generates a warning alert and displays the yellow warning badge in the Asset navigation tree. A warning alert contains a suggested action.
- **Critical:** Generates a critical alert and displays the red critical badge in the Assets tree. By default, a critical alert contains an automated action.

See [Monitoring Operating Systems](#) for more information.

To Clear a Boot Environment Incident

Perform the following steps to manually clear the disk space and close the incident:

1. Click an OS in the Assets tree to view the Boot Environment in the center pane, then click the **Incidents** tab.
2. Delete one or more boot environments to clear the disk space.
3. Select one or more incidents, then click the **Close Incidents** icon in the center pane.

See [Chapter 4, "Monitoring Rules and Policies"](#) for details on how monitoring rules and policies work in the software and see [Chapter 9, "Incidents"](#) for information on how to manage and close an incident.

Viewing Boot Environments

The active boot environment is identified in the Boot Environments tab and in the Summary tab. Inactive boot environments appear in the Boot Environments tab for each operating system. For global zones, the Boot Environments tab displays the relationship between the associated zones and existing boot environments.

Note: If there are no alternate boot environments, the Boot Environments tab does not appear in the center pane.







Boot environment support for Oracle Zones is available beginning with Oracle Solaris 11, enabling boot environments from the non-global zone to appear in the UI. Boot environments for non-global zones appear in the Boot Environment tab, not the Assets tree. Select a non-global zone in the Assets tree, then click the Boot Environments tab in the center pane.

The Summary tab provides the name of the active boot environment, the zpool utilization of all zpools for the selected OS and the amount of zpool utilization attributed to the boot environments. Details are available in the Boot Environments tab.

1. Click an OS or zone in the Assets tree. Details about the active boot environment appear in the Summary tab.

Non-bootable snapshots are available beginning with Oracle Solaris 11. When a boot environment has one or more associated non-bootable snapshots, the snapshots appear in the Snapshots subtab, as shown in [Figure 12-18](#).

Figure 12–18 Boot Environment Tab and Snapshots Subtab

Boot Environments				
  				
Search <input type="text"/>				
Active	BE Name	Size	Description	Created/Last Synchronized on
	solaris	19.14 GB		11/15/2011 6:58:06
	B1	165.00 KB		11/16/2011 9:46:29
	solaris-backup-1	242.00 KB		02/27/2012 4:36:59
	solaris-backup-2	576.00 KB		02/28/2012 3:57:07
Snapshots File Systems Associated Zone BEs				
Snapshots of selected BE: solaris				
 				
Search <input type="text"/>				
Name	Size	Created on		
solaris@2012-02-28-10:57:06	357.00 KB	02/28/2012 3:57:06 am MST		
solaris@2012-02-27-11:36:58	7.91 MB	02/27/2012 4:36:58 am MST		

- Click the **Boot Environments** tab in the center pane to see details about the associated alternate boot environments.
If the tab is disabled, no alternate boot environments exist for that OS in the zone or non-global zone.
- To display file system details, click the boot environment in the table, then click the **File Systems** subtab:
 - Oracle Solaris 11:** Mount point information, boot environment file systems, zone boot environment information, and non-global zones in the selected boot environment. When a global zone boot environment is selected, the zone boot environment details are listed separately, a zone per row. You can expand all rows to display a complete view of all zone boot environments under the selected global zone boot environment.
 - Oracle Solaris 10:** Mount point information and boot environment file systems.
- To display associated zone boot environment details, click the boot environment in the table, then click the **Associated Zone BEs** subtab:
 - Oracle Solaris 11:** Mount point information, boot environment file systems, zone boot environment information, and non-global zones in the selected boot environment. When a global zone boot environment is selected, the zone boot environment details are listed separately, a zone per row. You can expand all rows to display a complete view of all zone boot environments under the selected global zone boot environment.
 - Oracle Solaris 10:** Mount point information and boot environment file systems.

Monitoring rules and thresholds are defined in the Monitoring tab.

Managing Boot Environments

You can view all available boot environments for a system, and choose to activate an alternate boot environment or delete inactive environments.

Activating and Reboot a Boot Environment

Activating a boot environment makes an inactive or alternate boot environment the active boot environment. You can activate a single boot environment, or you can select an OS group and activate an alternate boot environment for each operating system.

To Activate a Single Boot Environment

Information about boot environments appears in the Boot Environments tab for each operating system.

Perform the following steps to activate a single boot environment:

1. Click an Oracle Solaris 11 global or non-global zone or an Oracle Solaris 10 global zone in the Assets tree.
2. Click the **Boot Environments** tab.

Existing boot environments for the OS appear in the center pane.

To Activate Boot Environments for All Members of a Group

You can activate an alternate boot environment for all members of a group. If each OS in the group has a single alternate boot environment, all operating systems are booted into the alternate boot environment. When some systems have more than one alternate boot environment, you are prompted to select the alternate boot environment.

Perform the following steps to activate boot environments for all members of an operating system group:

1. Select an Oracle Solaris 11 or an Oracle Solaris 10 Operating system group from the Assets section in the Navigation pane.
2. Click **Activate Boot Environment and Reboot** in the Actions pane.
3. For systems with multiple alternate boot environments, select the one that you want to boot into. A Filter ABE by name option is available to identify similarly named alternate boot environments across multiple targets.

Deleting a Boot Environment

When you delete a boot environment, you delete all associated snapshots and unshared file systems. A snapshot is a non-bootable copy of a boot environment. If there are non-global zone boot environments associated with the global zone boot environment, they are deleted too. Shared file systems are not deleted.

You cannot delete the active boot environment.

1. Click an Oracle Solaris 11 global or non-global zone, or an Oracle Solaris 10 global zone in the Assets tree.
2. Click the **Boot Environments** tab.
3. Select one or more boot environments or snapshots that you want to remove, then click the **Delete** icon. [Figure 12-19](#) shows a snapshot selected for deletion.

Figure 12–19 Delete an Oracle Solaris 11 Snapshot

Overview of Oracle Solaris 11 Boot Environments

Oracle Solaris 11 uses ZFS file systems, where the swap and dump volumes are shared within the pool. For unshared file systems, ZFS is the only supported file system for boot environments on Oracle Solaris 11. With the ZFS-based boot environment, the boot environments are clones of the existing ZFS partitions. This saves disk space and you do not need to reserve disk partitions for additional boot environments. The create and activate boot environment functionality is much faster than in previous versions.

Oracle Solaris 11 boot environments are managed through the `beadm` utility. A new boot environment is created whenever the kernel or system packages are installed or updated. This can result in high disk space utilization levels. By default, Oracle Enterprise Manager Ops Center monitors the disk (zpool) utilization of boot environments. If the utilization levels exceed defined thresholds, you can delete unwanted boot environments.

When a boot environment is created in the global zone, the following occurs:

- A Boot Environment of the source boot environment is created (the boot environment from which it is cloned).
- The currently active boot environment in all of the non-global zones is cloned and it is associated with the global zone boot environment that was just created.

When a global zone boot environment is activated, the active boot environment data set that is associated with that boot environment in each non-global zone is mounted and activated. Only one non-global zone boot environment that is associated with a global zone boot environment can be active.

When a global zone boot environment is deleted, all corresponding zone-specific boot environments are also deleted.

When a non-global zone is deleted, all corresponding boot environments are deleted. The boot environments for other zones are not deleted.

Displaying Oracle Solaris 11 Boot Environment Details

The Oracle Solaris 11 Boot Environments tab provides you with a large amount of information about the boot environments that are associated with the selected physical or virtual Oracle Solaris 11 operating system. This is particularly valuable since Oracle Solaris 11 boot environments are automatically created and they can quickly consume valuable resources. The following information is available:

- [Displaying Total ZPools Utilization for Oracle Solaris 11 Boot Environments](#)
- [Displaying Oracle Solaris 11 Boot Environments](#)
- [Displaying Snapshots for Oracle Solaris 11 Boot Environments](#)

- [Displaying File Systems for Oracle Solaris 11 Boot Environments](#)
- [Displaying Associated Zone Boot Environments](#)

Displaying Total ZPools Utilization for Oracle Solaris 11 Boot Environments

The first section in the Boot Environments tab is Total ZPools Utilization. It is compressed by default. Use the arrow to expand the table. The amount of resources used by the zpool appears in this section, as shown in [Figure 12–20](#).

Figure 12–20 Oracle Solaris 11 Total ZPools Utilization

Total ZPools Utilization			
ZPool Name ▲	% Space used by all Boot Environments	% Total Space Used	Total Space
rpool	<div><div></div></div> 17%	<div><div></div></div> 20%	72.80 GB

Displaying Oracle Solaris 11 Boot Environments

All boot environments that are associated with the physical or virtual operating system that is selected in the Assets tree appear in the Boot Environments table. As shown in [Figure 12–21](#), the active status, size, and when the boot environment was created appear in this section. A green check mark icon next to the boot environment name identifies the boot environment as active. A green check mark with a green circle and white x indicates that this is the boot environment that is active upon reboot. When you see two green check marks in the Active column, the boot environment is active and is the active boot environment upon reboot.

Figure 12–21 Oracle Solaris 11 Boot Environments Tab

Boot Environments				
<div> <div></div> <div></div> <div></div> </div>				
<div>Search ▾</div>				
Active	BE Name	Size	Description	Created/Last Synchronized on
✓✓	solaris	19.14 GB		11/15/2011 6:58:06
	B1	165.00 KB		11/16/2011 9:46:29
	solaris-backup-1	242.00 KB		02/27/2012 4:36:59
	solaris-backup-2	576.00 KB		02/28/2012 3:57:07
<div> <div>Snapshots</div> <div>File Systems</div> <div>Associated Zone BEs</div> </div>				
Snapshots of selected BE: solaris				
<div> <div></div> <div></div> </div>				
<div>Search ▾</div>				
Name	Size	Created on		
solaris@2012-02-28-10:57:06	357.00 KB	02/28/2012 3:57:06 am MST		
solaris@2012-02-27-11:36:58	7.91 MB	02/27/2012 4:36:58 am MST		

The lower section of the page has three tabs that provide details about the boot environment that you select in the Boot Environments table.

Displaying Snapshots for Oracle Solaris 11 Boot Environments

A snapshot is a point-in-time image of a Volume. It is a non-bootable copy of a boot environment that uses much less disk space than a boot environment. You can create a boot environment from a snapshot.

Non-bootable snapshots are available beginning with Oracle Solaris 11. Select a boot environment in the Boot Environments table to see all associated non-bootable snapshots in the Snapshots tab. You can select a snapshot in this tab and click the icon to create a boot environment from the snapshot.

Figure 12–22 Oracle Solaris 11 Boot Environment Snapshots

Name	Size	Created on
solaris@2012-02-22-19:45:03	21.12 MB	02/22/2012 12:45:03 pm MST
solaris@2012-01-27-23:41:50	195.01 MB	01/27/2012 4:41:50 pm MST

Displaying File Systems for Oracle Solaris 11 Boot Environments

The File Systems tab provides file system details for the boot environment, or alternate boot environment, that you select in the Boot Environments tab. As shown in [Figure 12–23](#), the file system name, type, size and mount location appear in this table. You can also see if the file system is shared or not.

Figure 12–23 Oracle Solaris 11 Boot Environments File Systems

Name	Type	Size	Mounted on	Shared
rpool	ZFS	13.80 GB	/rpool	true
rpool/ROOT/solaris	ZFS	11.73 GB	/	false

Displaying Associated Zone Boot Environments

The Associated Zone BEs tab is populated with the boot environment details for a selected zone's boot environment. The zone might have multiple boot environments. The table shows when the boot environment selected is active. As shown in [Figure 12–24](#), you can see the boot environment name, the zone name, the size and when the boot environment was created. When you select an operating system that is not a zone, No Data appears in the table.

Figure 12–24 Oracle Solaris 11 Boot Environments Associated Zone BEs

Active	BE name	Zone Name	Size

Zone Name: sol11mar1s11Tes1 (1)

About Boot Environment Profiles and Plans for Oracle Solaris 11

A Boot Environments plan and profile for Oracle Solaris 11 operating systems identifies and defines how the boot environment is updated and activated. You can use an agent-managed or agentless-managed operating system.

Use the profile for Oracle Solaris 11 to perform the following tasks:

- Create a Boot Environments profile that defines the creation policies
- Create an OS update deployment plan, specifying a Boot Environments profile and an OS update profile

You can create profiles with different policy options. The following options are available for Oracle Solaris 11:

- Create a new boot environment only when needed, such as when the OS update operation requires a reboot.
- Always create a new boot environment.
- Never create a new boot environment.
- Activate and reboot.

Note: Oracle Solaris 11 boot environments cannot have spaces in the name. When you name a boot environment, do not use spaces.

To Create a Boot Environments Profile for Oracle Solaris 11

1. Click **Plan Management** in the Navigation pane, then scroll down to the Profiles and Policies directory and click **Boot Environments**.

The Boot Environments Profiles page appears in the center pane.

2. Click **Create Profile** in the Actions pane.
3. Enter a unique profile name, without spaces, and provide a description to identify the profile. Select the **Oracle Solaris 11** subtype.
4. Choose the boot environment policy for this profile:
 - Create a new boot environment when needed.
 - Always create a new boot environment.
 - Never create a new boot environment.
5. Click the check box to **activate boot environment and reboot** when the job is completed.
6. Choose to have the software automatically create a unique name or enter a specific boot environment name in the field provided, then click **Finish**.

To Copy an Oracle Solaris 11 Boot Environments Profile

Copying a boot environment profile makes a copy of an existing profile.

1. Click **Plan Management** in the Navigation pane, then scroll down to the Profiles and Policies directory and click **Boot Environments**.

The Boot Environments Profiles page appears in the center pane.

2. Select the profile to copy, then click the **Copy Profile** icon in the center pane.
3. Enter a unique profile name, without spaces, and a description.

4. Change the boot environment parameters, as needed, then click **Finish**.

Creating an Oracle Solaris 11 Boot Environment

A new alternate boot environment is automatically created when you install or update the operating system's kernel or system packages.

You can create a bootable or non-bootable copy of an existing boot environment. A non-bootable copy is called a snapshot. You can create a boot environment from a snapshot.

To Create an Oracle Solaris 11 Boot Environment

You can create a bootable or non-bootable (snapshot) copy of an Oracle Solaris 11 boot environment.

1. Click an Oracle Solaris 11 OS, a global or non-global zone, in the Assets tree, then click the **Boot Environments** tab.
2. Select the boot environment to copy, then click the **Create** icon.
3. Select either the **Create a Snapshot** or **Create a BE** option, then select the boot environment to use as the source. Enter a unique name for the boot environment, without spaces, and provide a description that includes the purpose of the boot environment and the name of the source boot environment. Click **Create**.

Figure 12–25 Create Boot Environment for Oracle Solaris 11

Create Boot Environment(BE)

☒ Create a snapshot from solaris

☐ Create a BE from solaris

* Name:

Description:

4. When the job finishes, refresh the Boot Environment table to see the new boot environment.

If you have a group that contains only Oracle Solaris 11 operating systems and you select the OS group as the target, a new boot environment is created for every active boot environment in the group. You can assign a common name for all of the boot environments. If you do not assign a common name, default names are assigned.

Overview of Oracle Solaris 10 Boot Environments

Boot environments in Oracle Solaris 10 are managed through the Live Upgrade feature. A boot environment is created by using a script that contains the `lucreate` command and options.

The Live Upgrade feature is available on the following operating systems:

- Oracle Solaris 10 OS for x86 platform versions through Oracle Solaris 10 5/09: Update alternate boot environments for physical systems.
- Oracle Solaris 10 OS for SPARC through Oracle Solaris 10 5/09: Update alternate boot environments for physical and virtual machines, including Oracle Solaris Zones and Oracle VM Server for SPARC (formerly known as Logical Domains).

Note: Do not use Oracle Solaris Live Upgrade on your Enterprise Controller or Proxy Controllers. It does not synchronize all of the files that are required for these Oracle Enterprise Manager Ops Center components.

To create and update alternate boot environments, install the latest Oracle Solaris Live Upgrade packages and patches. This ensures that you have all the latest bug fixes and features. In addition, verify that the disk space and format are sufficient for an alternate boot environment.

Live Upgrade packages and patches are available for each Oracle Solaris software release beginning with Oracle Solaris 10 OS. Use the packages and patches that are appropriate for your software instance.

If you installed Oracle Solaris 10 using any of the following software groups, have the required packages:

- Entire Oracle Solaris Software Group Plus OEM Support
- Entire Oracle Solaris Software Group
- Developer Oracle Solaris Software Group
- End User Oracle Solaris Software Group

When you install one of these Software Groups, verify that you have all of the required packages:

- Core System Support Software Group
- Reduced Network Support Software Group

For detailed requirements to install and use the Live Upgrade feature, see the [Oracle Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning](#). For information on other releases, see the [Oracle Solaris Documentation Library](#). Review and verify that all the packages and patches that are relevant to your system are installed and that the disk is formatted properly before creating a new boot environment.

For Oracle Solaris 10 alternate boot environments, file systems are categorized into the following types:

- **Critical File Systems:** Non-sharable file systems that are required by the Oracle Solaris OS, such as root (/), /usr, /var, and /opt. These file systems are separate mount points in the `vfstab` of the active and inactive boot environments and are always copied from the source to the inactive boot environment.
- **Sharable File Systems:** User-defined files, such as /export, that contain the same mount point in the `vfstab` in both the active and inactive boot environments. Updating shared files in the active boot environment also updates data in the inactive boot environment. When you create a boot environment, sharable file systems are shared by default. If you specify a destination slice, also known as a partition, the file systems are copied.

- **Swap:** Depends on the type of file system:
 - For UFS file systems, swap is a special sharable volume. Like a sharable file system, all swap slices are shared by default. If you specify a destination directory for swap, the swap slice is copied.
 - For ZFS file systems, swap and dump volumes are shared within the pool.

When creating a new boot environment in Oracle Solaris 10, the entire contents of a slice is copied to the designated new boot environment slice. You might want some large file systems on that slice to be shared between boot environments rather than copied to conserve space and copying time. File systems that are critical to the OS such as root (/) and /var must be copied. File systems such as /home are not critical file systems and could be shared between boot environments. Sharable file systems must be user-defined file systems and on separate swap slices on both the active and new boot environments. You can reconfigure the disk several ways, depending your needs.

You can reslice, or partition, the disk before creating the new boot environment and put the sharable file system on its own slice. For example, if the root (/) file systems, /var, and /home are on the same slice, reconfigure the disk and put /home on its own slice. When you create any new boot environments, /home is shared with the new boot environment by default. To share a directory, the directory must be split off to its own slice.

Requirements for Oracle Solaris 10 Live Upgrade and Oracle Solaris Zones

The following requirements are needed to use the Live Upgrade feature with zones:

- Agent managed operating system
- At least Oracle Solaris 10 5/09 (update 7) operating system
- Storage library used to house the zones cannot be part of the root pool; you must create a separate pool on a shared file system
- You cannot use the -p option to create alternate boot environments

The -p option, which copies between two root pools on ZFS configuration, is not supported with the lucreate command.

Note: If you plan to use alternate boot environments with zones, you must designate sufficient zone storage space. When you create the zones and configure the zone storage, specify twice the size of the zone file system for the root file system of the zone. For example if your zone root file system was configured as 8 GB, the storage used to back up the zone must be at least 16 GB.

See [Updating Zones](#) for how to update Oracle Solaris Zones without using a dual boot environment.

Displaying Boot Environment Details for Oracle Solaris 10

The Oracle Solaris 10 Boot Environments tab provides you with a large amount of information about the boot environments that are associated with the selected Oracle Solaris operating system. A Boot Environments tab is not available for Oracle Solaris 10 non-global zones. The following information is available:

- [Displaying Total ZPools Utilization for Oracle Solaris 10 and Earlier Boot Environments](#)

- [Displaying Oracle Solaris 11 Boot Environments](#)
- [Displaying File Systems for Oracle Solaris 10 and Earlier Boot Environments](#)

Displaying Total ZPools Utilization for Oracle Solaris 10 and Earlier Boot Environments

The first section in the Boot Environments tab is Total ZPools Utilization. It is compressed by default, use the arrow to expand the table. The amount of resources used by the zpool appears in this section, as shown in [Figure 12–26](#).

Figure 12–26 Oracle Solaris 10 and Earlier Total ZPools Utilization

Total ZPools Utilization			
ZPool Name ▲	% Space used by all Boot Environments	% Total Space Used	Total Space
rpool	<div><div></div></div> 17%	<div><div></div></div> 20%	72.80 GB

Displaying Boot Environments for Oracle Solaris 10 and Earlier

All boot environments that are associated with the physical or virtual operating system that is selected in the Assets tree appear in the Boot Environments table. As shown in [Figure 12–27](#), the active status, size, description, and when the boot environment was created or synchronized appear in this section. A green check mark icon next to the boot environment name identifies the boot environment as active. A green check mark with a green circle and white x indicates that this is the boot environment that is active upon reboot. When you see two green check marks in the Active column, the boot environment is active and is the active boot environment upon reboot.

Figure 12–27 Boot Environments for Oracle Solaris 10 and Earlier

Boot Environments				
<div> </div>				
<div>Search <input type="text"/></div>				
Active	BE Name	Size ▼	Description	Created/Last Synchronized on
	newbemythilis	150.82 GB		-
	cherokeeBE	123.97 GB		-

The lower section of the page the File Systems tab that provide details about the boot environment that you select in the Boot Environments table.

Displaying File Systems for Oracle Solaris 10 and Earlier Boot Environments

The File Systems tab provides file system details for the boot environment, or alternate boot environment, that you select in the Boot Environments tab. As shown in [Figure 12–28](#), the file system name, type, size and mount location appear in this table. You can also see if the file system is shared or not.

Figure 12–28 Oracle Solaris 10 Boot Environments File Systems

File Systems				
File Systems of selected BE: EranSol10ABE				
Search <input type="text"/>				
Name	Type	Size	Mounted on	Shared
shripool/ROOT/EranSol10ABE	zfs	15.57 MB	/	false
/dev/zvol/dsk/shripool/swap	swap	4.00 GB	-	false

About Boot Environment Profiles and Plans for Oracle Solaris 10 and Earlier

You can use an alternate boot environment that was created outside of Oracle Enterprise Manager Ops Center; however, the preferred method is to create a Boot Environment profile and use the associated plan to create a boot environment. In either case, the operating system must be an agent-managed asset in the Oracle Enterprise Manager Ops Center software.

Use the profile for Oracle Solaris 10 and earlier releases to perform the following tasks:

1. Create a Boot Environments profile to specify how to create boot environments across all your managed system, using `lucreate` scripts
2. Create your boot environments using the action in the action panel for your assets. This is typically a one time operation, as you'll probably reuse the alternate boot environment.
3. Create a Boot Environments profile to always sync (and possibly) activate. Use this profile for OS update deployment plans.
4. Create an OS update deployment plan specifying the boot environment in step 3 and any OS update profile

Depending on the Oracle Solaris version that you are using, several methods are available for creating alternate boot environments. The Oracle Enterprise Manager Ops Center software uses Boot Environments profiles to help standardize and simplify the process.

Perform the following tasks before creating a Boot Environments profile:

1. [Determining Your Boot Environment Policy](#)
2. [Determining Your Live Upgrade Feature Requirements and Options](#)

Determining Your Boot Environment Policy

Each Boot Environments profile can have a different boot environment policy. The following options are available for Oracle Solaris 10:

- Activate Boot Environment and reboot on job completion.
- Synchronize existing Boot Environment before submitting a job.

Determining Your Live Upgrade Feature Requirements and Options

For Oracle Solaris 10 only, determine the file system and swap requirements and the `lucreate` command options that you want to use. If you are using ZFS, create a zpool. See the Oracle Solaris documentation for more information about the requirements and command options for the Live Upgrade feature.

You can create an operational plan that includes your script, or you can enter the `lucreate` options directly in the Boot Environments profile without creating a script.

Overview of Boot Environments Profiles and Plans

The following is a high-level overview of how to create an alternate boot environment for Oracle Solaris 10 and earlier releases.

1. Create a Boot Environments profile.

The profile identifies and defines how the boot environment is created and activated, including defining the target type, the options used to create the boot environment, and the activation parameters. Review the Oracle Solaris 10 Live Upgrade documentation for more details on the requirements:

- Define your file system and swap requirements.
- Determine the `lucreate` command options. You can enter the options in the Boot Environments profile, or you can create an operational profile that contains a script with the `lucreate` command and options. If you are specifying a variable boot environment name in the script passed to the operational profile, make sure the variable is called `BENAME`. When you use `BENAME` as the variable name, the operational plan passes on the BE name entered during plan execution to the `lucreate` script.
- Determine your boot environment policies.

2. Create a deployment plan.

The Boot Environment deployment plan defines the failure policy and is associated with a single Boot Environments profile. You can add a Boot Environments profile as a task, or step, within a complex deployment plan. See Plan Management for information about complex plans.

3. Execute the deployment plan.

After the deployment plans are created, a user with the appropriate role and privileges can choose from a list of plans to quickly and consistently create an alternate boot environment.

4. Create and deploy an OS Update plan to update the inactive boot environment. See [Updating an Oracle Solaris Boot Environment](#) for the steps to create and deploy the OS Update plan.

For information about system administration tasks such as managing file systems, mounting, booting, and managing swap space, see the *Oracle Solaris System Administration Guide: Devices and File Systems* in the Oracle Solaris 10 Documentation Library.

Creating an Oracle Solaris 10 Boot Environment

Oracle Solaris Live Upgrade scripts are used to create an alternate boot environment. To create a replica of your boot environment, run the script in Oracle Enterprise Manager Ops Center.

The following methods are available to create the alternate boot environment:

- Copy an existing boot environment.
- Execute a Boot Environments deployment plan.
The deployment plan must reference an Boot Environments profile.
- Run an operational plan that contains the `lucreate` script.
- Upload an Oracle Solaris Live Upgrade script as Local Content in Enterprise Manager Ops Center.

- Run an OS update job and specify a pre-action which runs the script. You can select multiple compatible targets and create an alternate boot environment for each target using the same script at the same time.
- Create an OS profile, and then run an OS Update job. The profile enables you to define the components and the actions to be performed every time you use the profile to create an alternate boot environment.
- Run an Oracle Solaris Live Upgrade script from the command-line interface. With this method, you must log in to each agent and then run the script to create the boot environment.

When you create the alternate boot environment with an OS Update job, you can choose to run the job immediately, or you can schedule the job to run during your maintenance window. In all methods, the new boot environment is automatically discovered and a new Boot Environments tab appears in the center pane for OS management.

This task describes how to run a New OS Update job to create the alternate boot environment. Although it is a New OS Update job, the sole purpose of the job is to create an alternate boot environment. The job uses the Live Upgrade script that you uploaded to Local Content to create a duplicate of your boot environment.

Defining Deployment Options for Oracle Solaris 10

To create boot environments for Oracle Solaris 10, you can define the options in an Boot Environments profile or you can add a script to an operational profile. The software enables you to store the scripts and parameters and provides a couple of methods of deploying the scripts.

- As a profile
- As part of a plan
- As part of an update job

The scripts use the `lucreate` command to create an alternate boot environment. To see a complete list of options, see the Oracle Solaris 10 OS Live Upgrade `lucreate(1M)` documentation. The following are some commonly used options:

- `-m` option to create a boot environment and split a directory off to its own slice.

To share a directory, the directory must be split off to its own slice. The directory is then a file system that can be shared with another boot environment. You can use the `lucreate` command with the `-m` option to create a boot environment and split a directory off to its own slice. But, the new file system cannot yet be shared with the original boot environment. You must run the `lucreate` command with the `-m` option again to create another boot environment. The two new boot environments can then share the directory.

- `-s` option to force a synchronization.
- `-p` specifies the ZFS pool in which a new boot environment resides.

The `-p` is not required when the source and target boot environments are within the same pool. The `-p` option does not support the splitting and merging of file systems in a target boot environment, instead use the `-m` option.

Existing boot environments appear in the hierarchy in the Boot Environments tab for a global zone. You can create the boot environment with Oracle Enterprise Manager Ops Center or outside of the software.

Creating an Oracle Solaris 10 Boot Environment From a Deployment Plan

You can only create a boot environment from a global zone in Oracle Solaris 10. The boot environments on Oracle Solaris non-global zones are tightly coupled with the corresponding boot environment on the global zone and cannot be managed independently.

You can use either of the following deployment plans to create an Oracle Solaris 10 boot environment:

- **Boot Environment:** A simple deployment plan that only creates an alternate boot environment. Use the default alternate boot environment profile and then override the applicable parameters to create a boot environment.
- **Software Deployment / Update:** A multi-step plan where one of the steps creates a boot environment. See the *Oracle Enterprise Manager Ops Center Updating Your Oracle Solaris 10 Operating System* for an example of how to use this plan.

Perform the following steps to create a boot environment using the Alternate Boot Environment profile and Boot Environment deployment plan:

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.
Existing boot environments for the OS appear in the center pane.
3. To create a copy, or clone, of the existing active boot environment, click the **Create** icon.
4. Enter a unique name for the boot environment and provide a description. The text describes what the boot environment and from what source it was created. Edit the applicable parameters in the Alternate Boot Environment profile, then click **Confirm**.
5. When the job finishes, refresh the Boot Environment table to see the new boot environment.

Creating an Oracle Solaris 10 Boot Environment from an Operational Plan

An operational profile contains a single script and can define variables. The associated plan launches the profile. An operational plan is associated with a specific version of an operational profile. By default, creating a profile also creates an operational plan.

You can create the boot environments using this script from a few different contexts:

- Simple BE Operational profile
- Boot Environments profile using BE Operational profile from the plan context
- Create Boot Environment from the asset context

You can save a shell script on the Enterprise Controller and download it into the plan, or you can enter the script in a field when you create the plan. Both types of shell scripts are executed by the user. They differ in the location, either on the Enterprise Controller or on the remote Agent Controller, and the user credentials needed to execute the profile. See [Operational Profiles and Plans](#) for more information about operational profiles and plans and the variables you can use.

1. Click **Plan Management** in the Navigation pane, then click **Operational Profiles**. The Operational Profiles appear in the center pane.
2. Click **Create Operational Profile**. The Create Profile - Operation page appears.

3. Enter a name for the new plan and a description of its purpose or role.
4. Select an asset type from the Subtype list.
5. Click **Next**.
6. Select the Operation Type from the drop-down menu, either **EC Shell** or **Remote Shell**.
 - If you select EC Shell, browse to the location of the script in the Script File field, then click **Load Script**.
 - If you select Remote Shell, enter your script in the Script field.
7. Enter a numeric value in the Timeout field, then select Minutes or Seconds. The default timeout is 60 minutes.
8. (Optional) Click **View System Variables** to view the default variables.
9. Click **Next**.
10. (Optional) Specify Additional Variables. You can specify any variable you want. For example, add Alarm_ID\$ to add the incident identifier number for easier incident management.

Creating an Update Profile for Oracle Solaris 10 Boot Environment

You can use an Update profile to create and update a boot environment.

1. Expand **Plan Management**, then click **Update Profiles** in the Navigation pane.
2. Click **New Profile** in the Actions pane.
3. Enter a name and description for the profile.
4. Select **Script** from the Type drop-down menu.
5. Select an OS channel from the Distribution drop-down menu.
6. Select **Local** from the Category drop-down menu.
7. Choose the options that you want from the View and the Version drop-down menus.
8. Highlight Local RPMs in the Available Packages / Patches table, select the **Apply to all applicable distributions** check box.
9. Click **Profile Contents** and select an item, then click **Create OS Update Profile**.

Creating an Oracle Solaris 10 Boot Environment From an Update Profile

You can use an Update profile to create and update a boot environment.

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.

Existing boot environments for the OS appear in the center pane.
3. To create a copy, or clone, of the existing active boot environment, click the **Create** icon.

The plan for execution that uses the default Boot Environments profile for Oracle Solaris 10 appears.
4. Enter a unique name for the boot environment and provide a description. The text describes what the boot environment and from what source it was created. Edit

the applicable parameters in the Alternate Boot Environment profile, then click **Confirm**.

5. When the job finishes, refresh the Boot Environment table to see the new boot environment.

Creating an Oracle Solaris 10 Boot Environment With an OS Update Job

Oracle Solaris Live Upgrade contains a suite of script commands. To create an alternate boot environment with Enterprise Manager Ops Center, use the `lucreate` command to write one or more Oracle Solaris Live Upgrade scripts, add the scripts to the Local Content library in Enterprise Manager Ops Center, then run an OS Update job and select the ABE options.

When you use Enterprise Manager Ops Center to create the alternate boot environment, the scripts must meet the following requirements:

- The script cannot contain parameters.
- The alternate boot environment name must be hard-coded into the script itself or otherwise be provided outside of Enterprise Manager Ops Center.
- The alternate boot environment name defined in the script must match the alternate boot environment name that you use when you run the update job to create the alternate boot environment.
- The script must return 0 on success and non-zero on failure.

For detailed instructions and examples for using the `lucreate` command to create a boot environment, see *Oracle Solaris 10 9/10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* available at:

<http://www.oracle.com/technetwork/indexes/documentation/index.html>

1. Expand Libraries in the Navigation pane.
2. Click **Local Content** in the Solaris/Linux OS Updates library.
3. Click **Upload Local Action** in the Actions pane.
4. Enter a name for the file.
5. Enter a brief description of the purpose of the action.
6. In the Action list, click the **Pre-action** type of action to run the script on the managed host before job tasks are carried out.
7. Click the name of the distribution that uses the action in the Distribution list. The Parent field shows the category, based on the type of Action.
8. Click **Browse** to locate and select the file.
9. Click **Upload** to upload the file to the selected distribution.

To Create a Boot Environment With an OS Update Job

1. Click the OS in the Assets section of the Navigation pane.
2. Click **New Update OS Job** in the Update section of the Actions pane. The New Update OS Job Wizard appears.
3. Complete the following Job Information parameters:
 - Enter a job name.
 - Select **Actual Run**, which creates the alternate boot environment when you specify in Step 5.

- Select the Sequential task execution order.
 - Select the Target Setting: Use the same Targets for all tasks in the job.
 - Select a Task Failure Policy:
 - Complete as much of the job as possible
 - Stop at failure and notify
 - Select the **Boot Environment Type** check box.
 - Select the **Run ABE Pre-action Script** check box.
4. (Optional) To add tasks to the job, click the **Add Task** icon. To edit, click the Profile and Policy fields. Click **Next**.
 5. Enter the name of the alternate boot environment, as defined in the script. Select a script, then click **Next**.
 6. (Optional) Complete the Boot Environment Workflow, then click **Next**.
 - To synchronize the alternate boot environment with the current boot environment before mounting the alternate boot environment, select the **Sync ABE** check box.
 - To edit the description to describe the state of the Boot Environment, click **Modify Current BE**, and enter text in the Description field.
 - To edit the description to describe the state of the alternate boot environment, click **Modify Alternate BE**, and enter text in the Description field.
 - To switch boot environments after update, select the **Activate and Reboot ABE** check box.
 7. Schedule the job, then click **Next**.
 - Run Now: Starts the job immediately after you click Finish in the Job Summary.
 - Start Date: Select a date and time to start the job.
 - On a recurring schedule: Enables you to run the same job on a monthly or daily scheduled time.
 8. Review the Job Summary, then click **Finish** to run the job as scheduled in the previous step.

When the job completes, the new alternate boot environment is associated with the operating system. To verify that the alternate boot environment is created, click the operating system in the Assets pane. The Boot Environments tab appears in the center pane. Click the Boot Environments tab to display the new boot environment, as specified in the Live Upgrade script. An OS can have several associated alternate boot environments.

Note: The Boot Environments tab only appears when at least one alternate boot environment associated with the operating system.

Synchronizing Oracle Solaris 10 Boot Environments

You can synchronize, or sync, an active Oracle Solaris 10 global zone boot environment with an inactive boot environment on the same system. Synchronizing boot environments makes the inactive boot environment the same as the currently

running boot environment. After you sync the boot environments, you can activate the inactive, or alternate, boot environment.

Activating a Boot Environment

Activating a boot environment switches a new boot environment to become the currently running boot environment when the system reboots.

1. Click an Oracle Solaris 10 OS global zone in the Assets tree to view the Boot Environment in the center pane.
2. Click the **Boot Environment** tab.
Existing boot environments for the operating system appear in the center pane.
3. Click **Activate Boot Environment and Reboot** in the Actions pane.
4. To schedule the activation for a later date or time, select the check box. Click **Next**.
5. If you selected the schedule option, complete the schedule, then click **Next**.
6. Review the Summary, then click **Finish**.

The new boot environment activates when you reboot the system.

Related Resources

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

- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm
- See [Chapter 13, "Operating System Updates"](#) for information about patching, or updating, your operating systems.
- See [Chapter 14, "Oracle Solaris Zones"](#) for information about zones and how you can use Oracle Enterprise Manager Ops Center to efficiently manage all phases of zones lifecycle.
- See [Chapter 4, "Monitoring Rules and Policies"](#) for information on how monitoring rules and policies work in the software.

For in-depth information about these products, see the following Oracle documentation:

- For a list of the Oracle Linux documentation available in HTML and PDF formats, visit the Oracle Linux Documentation website at <http://www.oracle.com/us/technologies/linux/index.html>.
- For a list of the Oracle Solaris 11 documentation available in HTML and PDF formats, visit the Oracle Solaris 11 Documentation website at <http://www.oracle.com/technetwork/documentation/solaris-11-192991.html>.
- For a list of the Oracle Solaris 10 documentation available in HTML and PDF formats, visit the Oracle Solaris 10 Documentation website at <http://www.oracle.com/technetwork/documentation/solaris-10-192992.html>.
- For a list of the Oracle Solaris 8 and 9 documentation, visit the Legacy Solaris Documentation website at <http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html>.

- For more information about JET resources and documentation, see *Solaris 10 10/09 Installation Guide: Custom JumpStart and Advanced Installations* available at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.
- For JET documentation and to download additional modules, see <http://www.oracle.com/technetwork/systems/jet-toolkit/index.html>.

Operating System Updates

This chapter describes the operating system update features that are available in the software.

The following information is included:

- [Introduction to Operating System Updates](#)
- [Roles for Operating System Updates](#)
- [Actions for Operating System Updates](#)
- [Location of Operating System Updates in the User Interface](#)
- [Using System Catalogs](#)
- [About Operating System Update Reports](#)
- [Creating Update Policies](#)
- [Creating Update Profiles](#)
- [Updating Oracle Solaris 8, 9, and 10 and Linux Operating Systems](#)
- [Updating Oracle Solaris 11 Operating Systems](#)
- [Updating an Oracle Solaris Boot Environment](#)
- [Updating Microsoft Windows Operating Systems](#)
- [Related Resources](#)

Introduction to Operating System Updates

Oracle Enterprise Manager Ops Center reduces the complexity of updating a large number of systems, standardizes the update installation process, minimizes downtime, and enables you to choose the level of automation.

You can maintain your Oracle Solaris, Linux, and Microsoft Windows operating systems to the recommended and latest updates and perform complex update tasks in a consistent manner. For most platforms, the update features help you to perform the following tasks:

- Manage different operating system update conditions that exist for installing an update
- Identify dependencies
- Download update packages or updates from the appropriate vendor sites
- Run an update simulation to test the update in your environment

- Rollback your systems to a previous state if an update is not stable in your environment
- Maintain consistent component configuration of your systems to the latest security updates

The list of supported operating system releases and functionality is available in the *Oracle Enterprise Manager Ops Center Certified Systems Matrix Guide*.

The update features include:

- Catalogs
- Reports
- Update Profiles
- Update Policies
- Deployment Plans

Oracle Enterprise Manager Ops Center provides one stop solution for all the requirements for updating your operating systems. You can use Update Profiles and Plans to define which components must be installed and the level of automation during the installation.

You create update jobs with operating system update profiles and policies to update an operating system. OS update profiles and policies define which updates to install, and how the update job proceeds after determining the update dependencies and user interaction. A set of system-defined update profiles is available in the UI. You can copy the profiles and edit them to create your own customized profiles.

A variety of operating system reports are available to give you insight into the operating system compliance status, the state of your operating system update levels, and provide information about the recommended updates and packages.

You can view your managed operating systems with the All Assets view or the Operating Systems view, as shown in [Figure 13-1](#) and [Figure 13-2](#).

Figure 13-1 All Assets View

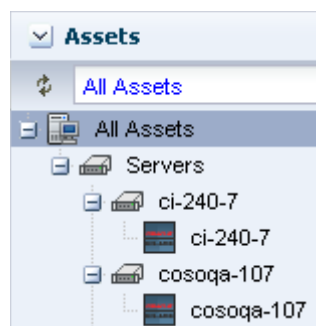



Figure 13–2 Operating Systems View

Requirements for Updating Operating Systems

The Enterprise Controller obtains information about latest updates from the Knowledge Base, Oracle Solaris 11 parent repositories on Oracle.com, and vendor sites. You must have an Internet connection to obtain the updates and packages from the various locations. In the absence of an Internet connection to the Enterprise Controller, you can get the latest updates by using a special script, called the harvester, to create local versions of the Knowledge Base and Oracle Solaris 11 Software Update Library. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about connection modes and how to use the harvester script.

The operating system update feature has the following requirements

- **Agent managed:** Linux and Oracle Solaris 8, 9, and 10 operating system must be agent-managed to perform most operating system update tasks. You can update Oracle Solaris 11 and Microsoft Windows with an agentlessly managed operating system. See [Using Agent Management for Operating Systems](#) for details and how to switch management modes.
- **Access to updates:** The software is designed to use a secure Internet connection to obtain operating system updates as they become available. If you cannot use an Internet connection from within your datacenter, you can run the software in disconnected mode and download the latest updates to an external storage device and use that to update the software libraries.
- **Established Software Update Library:** For Linux or Oracle Solaris, you must create an update library to store the update information. See [Chapter 5, "Software Libraries"](#) for information on establishing and using software update libraries.
- **Compatible configuration for Oracle Solaris 11:** To update an Oracle Solaris 11 operating system, the Enterprise Controller and the Proxy Controller must be running on an Oracle Solaris 11 operating system.
- **Valid update credentials:** You must have update credentials for each platform vendor. You must provide a valid My Oracle Support (MOS) account for Oracle Linux and Oracle Solaris. Valid vendor credentials are required for SuSE Linux and Microsoft Windows. Typically, you supply the credentials when you install Oracle Enterprise Manager Ops Center. See [Authentications](#) in the *Oracle Enterprise Manager Ops Center Administration Guide* for how to add or edit your update credentials.

Manage the operating systems with an agent for the most robust feature set, including reports, monitoring, and analytics. See [Discovering and Managing Assets](#) for more about the discovery and management process. See [Using Agent Management for Operating Systems](#) for the differences between the two modes and how to change the agent management mode.

Methods of Running an Update Job

You can use the following methods to run an update job:

- OS Update deployment plan that uses the update profiles and policies
- System catalogs
- Reports
- Update profile

Options Available When Running an Update Job

You have the following options available when running an update job:

- Select update profiles and policies.
- Select different targets for each task in the job.
- Select job simulation mode. Simulating a patching job helps to estimate the time required to run the job, to know the patch dependencies and the expected job result. In the simulation mode, you can select to download the required updates.
- Failure policy to determine the action when a task fails.

Roles for Operating System Updates

[Table 13–1](#) lists the tasks that are discussed in this section and the role required to complete the task. An administrator with the appropriate role can restrict privileges to specific targets or groups of targets. Contact your administrator when you do not have the necessary role or privilege to complete a task.

See [Asset Management](#) for details on using the discovery feature to add assets and for more information on creating user-defined groups. Contact your administrator if you do not have the necessary role or privilege to complete a task. See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 13–1 Operating System Management Roles and Permissions

Task	Role
New Update OS Job	Update Admin
Deploy or Update Software	Update Admin
Simulate an OS Update	Update Admin or Update Sim Admin
Compare System Catalog	Update Admin
Create Catalog Snapshot	Update Admin
View and Modify Catalog	Update Admin

Table 13–1 (Cont.) Operating System Management Roles and Permissions

Task	Role
Update Management Credentials	Security Admin
Any Actions related to changing credentials	Security Admin
Import image	Storage Admin
Upload image	Storage Admin
Upload image	Storage Admin
Unconfigure, SCCM Configuration	Ops Center Admin
Reboot, upgrade Agent Controller	Asset Admin
Edit Tags	Asset Admin
Edit Attributes	Asset Admin

Actions for Operating System Updates

Two management modes are available, agent managed and agentlessly managed. To perform updates to your Linux or Oracle Solaris 8, 9, or 10 operating system, the operating system must be a managed asset. Oracle Solaris 11 and Microsoft Windows do not require an agent managed operating system.

You can manage your operating systems by installing an Agent Controller on the OS or by using SSH to perform tasks. You must have an agent managed operating system to use the OS Update features, including the system catalogs, OS update jobs, and reports. See [Using Agent Management for Operating Systems](#) for more details, including when an agent is required, how to determine the management mode, and how to change the management mode.

After you manage your assets, you can perform the following actions:

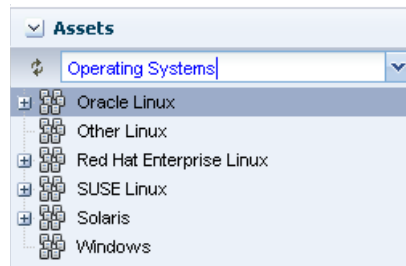
- Monitor your physical and virtual operating systems
- View OS utilization for Oracle Solaris and Linux operating systems
- Provision Oracle Solaris and Linux operating systems
- Manage Oracle Solaris boot environments

After you manage your assets, you can perform the following actions:

- Create a System Catalog
- Create an OS Report
- Update Oracle Solaris and Linux operating systems
- Update Microsoft Windows operating systems

Location of Operating System Updates in the User Interface

To see operating system information, expand Assets in the Navigation pane, then select the operating system. You can use the Operating Systems filter to just display the operating systems, as shown in [Figure 13–3](#).

Figure 13–3 Operating System Filter

Using System Catalogs

The software automatically takes a snapshot of the OS after executing any job on the OS, including when you discover and manage the OS. A snapshot is stored as a catalog with the time stamp and job details after every update job that you run on a system.

You can create a new catalog at any time and use it to record the state of a system. Catalogs enable you to rollback your system to any previous configuration or to create a profile that you can use to apply a consistent configuration throughout your data center.

The Catalog List contains all of the snapshots. When you create a historical catalog, the current state of the selected system is identified and stored as the previous catalog of the system. The saved previous catalog is the most recent system catalog

Note: You can create a historical catalog only for the current state of the system.

The catalog list always provides the listing of the most recent catalog. The software updates the catalog list whenever you update a system or create an historical catalog. You can identify the current catalog by the time stamp. You can use an historical catalog to create a profile and apply it to configure other systems.

Viewing and Modifying a Catalog

When you are using dual boot environments for Oracle Solaris Live Upgrade (Oracle Solaris 10), the catalog displays the inventory of the active boot environment of the operating system. To view the catalog of an alternate boot environment (ABE), you must first activate the ABE from the UI, and then wait for the job to finish. The software updates the current catalog and contains the ABE catalog information and OS software components. This automatically updates the catalog of any zones.

When you have an alternate boot environment, you cannot create and compare catalogs until you activate the ABE. By default, only the catalogs of the active boot environment are compared.

Comparing System Catalogs

You can compare two managed systems or two system catalogs for differences in the installed update components. You can also compare the current system catalog and saved snapshots of the same managed system to examine the differences in the components that are installed and uninstalled after executing a job.

Use the Compare Catalogs option to change the software components of a particular operating system to that of the source system.

The following options are available when you compare catalogs:

- Differences Between Systems: Displays the difference between the source and the target systems update components. The difference appears in the Compare Catalog window.
- Tasks to Make Target Like Source: Creates the list of components that must be installed on the target system. Select Include for the components to install on the target system.

About Operating System Update Reports

To ensure that your managed systems are up-to-date, you must determine which updates (packages and patches) and actions to apply to your system. The OS update reports help you to determine the updates that are applicable to your systems and how many of the applicable updates are compliant or not compliant for the selected systems.

The OS Update reports enable you to check for new updates and update your systems. You can get a general report, or test a system for available fixes.

When you create a report, you select the criteria that are relevant to you, such as a list of hosts that have a specific patch or a list of hosts that do not have a specific patch. [Table 13–2](#) lists the report name and the operating systems for which you can pull the report. See [Creating an Operating System Report](#) in [Chapter 10](#) for detailed information about these reports.

Table 13–2 Compatibility of Reports on Operating System

Report Name	Linux OS	Oracle Solaris 8, 9, and 10	Oracle Solaris 11	Microsoft Windows
CVE Compliance Report	Yes	Yes	Yes	No
Recommended Software Configuration Report (RSC) Report	Yes	Yes	Yes	No
Oracle Solaris Update Compliance Report	No	Yes	Yes	No
Baseline Analysis Report	No	Yes	Yes	No
Incident Compliance Report	Yes	Yes	Yes	Yes
Host Compliance Report	Yes	Yes	Yes	Yes
Profile Analysis Report	Yes	Yes	Yes	No
Package Compliance Report	Yes	Yes	Yes	No

Creating Update Policies

An Update policy defines the level of interaction required during an OS update job for Oracle Solaris and Linux operating systems. Policies define how to answer any questions that are raised during installation, uninstallation, or upgrade of Oracle Solaris and Linux updates.

The following system-defined policies are available in the software:

- Ask for All: The software stops the update job for each action and consults you on all action to take.
- No to All: The software denies all actions.
- Yes to All: The software confirms all actions.

Note: All operating system update plans use the "yes to all" policy by default.

You can create customized policies that answer differently depending on the specific OS update component. For example, when you want to review the questions and manually supply answers for patch 121288, which is part of the Oracle Solaris 10 Recommended updates from February 23, 2012, you can choose to set a specific policy for that patch, as shown in [Figure 13–4](#).

Figure 13–4 OS Update Policies and User Defined Policy Details

OS Update Policies				
<div> </div> <div> <div>Search</div> <div>X</div> </div>				
Policy Name	Description			Defined By
121288 Update Policy	Use this policy when updating 121288			User-defined Policies
Ask For All	System policy that consults the user on all actions			System-defined Policies
No To All	System policy that denies all actions			System-defined Policies
Yes To All	System policy that confirms all actions			System-defined Policies

User-defined Policy Detail				
<p>Reboot Policy: Do not ignore patches that require reboot</p> <p>Default Answer to Questions: Ask me</p>				
Exceptions to the Default Answer to questions				
Component	Description	Attribute	Distribution	Answer
10 Recommended [Feb/23/12]	10 Recommended [Feb/23/12]		SOLARIS_10_0_SPARC	Ask me

Policy settings are hierarchical. When there is not a policy setting for a component, the policy for that component's parent applies. For example, it is possible to create a policy that allows the system to install a given component but prohibits installation of certain specific versions of that component.

Note: The policy only applies to actions that are implicitly generated by the dependency resolver. If a conflict occurs between a profile and policy, the profile overrides the policy.

The update policy is not applicable when updating Microsoft Windows operating systems.

To Create a User-Defined Policy

1. Expand **Plan Management**, then click **Update Policies** in the Navigation pane.
2. Click **Create Policy** in the Actions pane.
3. Identify the policy by providing a name and description.

4. Click the **Available Packages / Patches** tab to view all available packages and updates. Use the search criteria to refine the list:
 - **Distribution:** The operating system distribution, such as Oracle Solaris 9 for SPARC, Oracle Solaris 10 for SPARC, or Oracle Solaris 10 for x86
 - **Category:** Package category, such as Cluster, Recommended Software Configuration, or Hardware
 - **View:** View all, available, withdrawn, modified packages/updates, or reboot
 - **Version:** View All Versions or filter for the Latest Versions

Figure 13–5 Create OS Update Policy

Component	Description	Attr
10 Recommended [Feb/23/12]	10 Recommended [Feb/23/12]	
103346-30	Hardware/PROM: Sun Enterprise 3x004x005x006x00 fl	
103451-05	Hardware, 2.1GB Disks: Download program and ST3255C	
104268-09	Hardware/PROM: Ultra 1E Standalone Flash PROM Update	

5. Select the component, then click an icon for the type of policy you want for the component, either **Answer Questions Yes**, **Answer Questions No**, or **Answer Questions Ask User**.
6. (Optional) To view details of the component, click the **View** icon.
7. Click **Create OS Update Policy**.

Creating Update Profiles

An update profile defines the component configuration of the systems that you want to manage. Update profiles specify which components are to be installed and which are prohibited, and any additional actions to be performed on an Oracle Solaris or Linux OS.

Use profiles to accomplish the following:

- Manage multiple systems in a consistent manner
- Automate repetitive administration jobs
- Record the requirements of your enterprise
- Automatically configure servers and workstations
- Manage dependencies and ensure consistency

The profile settings Required, Not Allowed, and Upgrade affect a managed host only during the actual deployment of that profile. At any time, you can run a job that contradicts the settings of a previously used profile; therefore you must understand your system settings and requirements thoroughly.

Predefined profiles are provided to perform common system-wide checks and to automate the operating system updates. These profiles cannot be edited or deleted.

You can identify the profile type when you create the profile. The profile type is a tag that is used to filter the required profiles when you create a deployment plan. The following are profile types:

- **Install:** Indicates that new components are added, or installed.
- **Upgrade:** Indicates that existing components are upgraded.
- **Script:** Indicates that action scripts are executed.

Note: You can create profiles that perform all of the actions for the profile type. The profile tag filters the required profiles in deployment plans. See [About Complex Deployment Plans](#) for more information.

To Create a New Profile

1. Expand **Plan Management**, then click **Update Profiles** in the Navigation pane.
2. Click **Create New Profile** in the Actions pane.
3. Enter a profile name and brief description of the profile.
 - a. Enter a name and description for the profile.
 - b. Select a Profile Type tag, either **Upgrade**, **Install**, or **Script**, to categorize and filter the profiles later.
 - c. Select a distribution from the drop-down list. For example, SOLARIS10_SPARC.
 - d. (Optional) You can further define the criteria by choosing a category, view, and version from the drop-down lists.

Figure 13–6 OS Update Profile Search Criteria

The screenshot shows a 'Search Criteria' dialog box with the following settings:

- Distribution:** SOLARIS10_SPARC
- Category:** all
- View:** Available
- Version:** Latest Versions

Figure 13–7 Create an OS Update Profile

The screenshot shows the 'Create an OS Update Profile' form with the following details:

- Name:** S10 - Recommended Feb 27, 2009
- Description:** Solaris 10 SPARC recommended package from February 27, 2009
- Type:** Install
- Available Packages/Patches:**

Component	Description	Attributes
10 Recommended [Feb/27/09]	10 Recommended [Feb/27/09]	
10 Recommended [Feb/20/09]	10 Recommended [Feb/20/09]	
10 Recommended [Feb/23/09]	10 Recommended [Feb/23/09]	

4. Locate and select a Component from the Component tree.
5. If required, select the check box to specify that the component is added to all applicable distributions.

Note: This only applies to distributions that are active at the time the profile is created. As new distributions are activated you must edit the profile to explicitly add any components for those distributions.

6. Specify whether the action is Required, Upgrade, or Uninstall.

Note: Some actions might not apply. For example, a component cannot be Required if the system does not have the information about how to obtain the component.

7. (Optional) You can repeat the preceding actions to select multiple components for the same or different operating systems.
8. Click **Save as Named Profile**. If a profile of that name exists you are asked to confirm that it is to be replaced.

Note: You cannot replace system-defined profiles.

Updating Oracle Solaris 11 Operating Systems

Oracle Solaris 11 uses a different update mechanism than earlier versions of the operating system. Oracle Solaris 11 uses packages to update the operating system and any non-global zones. The packages are part of an Image Packaging System (IPS) that is integrated with the ZFS file system.

When you install Oracle Enterprise Manager Ops Center on an Oracle Solaris 11 operating system, you create a local software package repository. The repository, called the Oracle Solaris 11 Software Update Library, is either on a file system on the same system as the Enterprise Controller, or on an NFS server share that the Enterprise Controller can access. You will populate the local library with packages from the parent repository instead of the Knowledge Base.

Note: To update an Oracle Solaris 11 operating system, the Enterprise Controller and Proxy Controller must be running on an Oracle Solaris 11 operating system. When Oracle Enterprise Manager Ops Center is not running on Oracle Solaris 11, the Oracle Solaris 11 update actions are not available.

The ZFS integration automatically creates an alternate boot environment every time an operating system is installed or updated. You can quickly and easily create an alternate boot environment when needed, and manage existing boot environments. Using an alternate boot environment provides a safe method of testing an update before deploying it to your live environment.

Updating Oracle Solaris 11

Unlike earlier versions of Oracle Solaris, you cannot run an ad-hoc operating system update job or browse package contents for Oracle Solaris 11. The best method of updating an Oracle Solaris 11 operating system is with a deployment plan.

To upgrade the operating system to the latest version of the Oracle Solaris 11 package, choose Upgrade all components.

Before you run an update job, verify that the Oracle Solaris 11 Software Update Library contains the package, or latest version of the package, that you want to use. See [Libraries for Oracle Solaris 11](#) for more information.

You will complete the following actions to update your Oracle Solaris 11 operating system:

1. Create an Oracle Solaris 11 update profile.
2. Optionally, create operational profiles that contain scripts to perform preinstall and postinstall actions.
3. Create an Oracle Solaris 11 update plan.
4. Select the target asset, then select the Deploy /Update Software action.

All update plans use the "yes to all" policy by default. If you do not want to automate the update by answering yes to all questions, you can create one or more customized update policies.

Create an update profile to define the images to use. After you have created the profile, you can create a deployment plan and choose the update and operational profiles to perform the tasks that you want for that plan. You can copy and edit plans to create customized plans for different purposes or targets.

About Oracle Solaris 11 Update Deployment Plan

An Oracle Solaris 11 update deployment plan provides a framework of steps and actions that you can perform to update your Oracle Solaris 11 operating systems.

The Oracle Solaris 11 Update deployment plan is a multi-step plan. The Update OS step is required, all other steps are optional. You cannot add steps to this type of plan.

Figure 13–8 Oracle Solaris 11 Deployment Template

Solaris 11 Update		
Target Type: Solaris 11		Version: 1
Template Name: Update Solaris 11 OS		Release Date: 03/22/2012 9:54:24 am MDT
Description: Updates a Solaris 11 OS		
Deployment Template Composition		
Step	Required Input	Required Profile/Plan
Update Configuration	Not Required	Not Required
Execute Pre-Install script	profile	Operational Profile
Create a Boot Environment	profile	Boot Environment Profile
Update OS	profile	Software Update Profile
Execute PostInstall script	profile	Operational Profile
Monitoring	profile	Monitoring Policy

The Oracle Solaris 11 Update deployment plan includes the following steps:

1. Update Configuration: A profile is not attached to this step. The configuration is not required. At runtime, you can choose the type of update job, either an update simulation or an actual job.
2. Preinstall script: Optionally, you can associate an operational profile that contains a script that performs an action before you apply the update.

3. Create an alternate boot environment: You can create an alternate boot environment before the update job. This is optional.
4. Update the operating system: The profile that applies the update packages.
5. Post install script: Optionally, you can associate an operational profile that contains a script that performs an action after you apply the update.
6. Monitoring: Optionally, you can enable monitoring.

You can choose a failure policy for the plan, either to stop the update when a failure occurs, or to complete as many of the steps as possible.

Updating Oracle Solaris 8, 9, and 10 and Linux Operating Systems

You will use a similar methodology and set of procedures to update your Linux and Oracle Solaris 8, 9, and 10 operating systems.

You can use the following update methods and options to update these operating systems:

- Use predefined or custom profiles and associated deployment plans to update a system or group of systems.
- Use a system catalog to create a simple update job without creating a profile. Use this method to apply a single patch quickly.
- Use the compliance reports output to update your OS. Use this method to make your systems compliant with newly released updates.
- Use compare catalogs to roll a system back to its previous state.

Table 13–3 *Methods of Creating an Operating System Update Job*

Update Method	Oracle Solaris 8, 9, and 10	Linux
Create and deploy an Update Plan	Yes	Yes
Create a new Update OS Job	Yes	Yes
Create Update Profile and Policy	Yes	Yes
Modify or compare a System Catalog	Yes	No
Create an Oracle Solaris Update Compliance report	Yes	No
Create a Compliance Report	Yes	No

In addition to the methods described previously, you can use Live Upgrade and alternate boot environments to update your Oracle Solaris OS with a minimum of downtime.

Updating an Operating System From a Deployment Plan

Deployment plans enable you to control how the update is performed and apply updates consistently across many systems. With this method, you can use custom or predefined profiles to perform complex update scenarios or to apply updates consistently across many systems.

You can use the following deployment plan templates to update a supported operating system:

- Software Deployment/Update: Use this plan to apply script based update profiles.

- **Configure Server Hardware and Install OS:** Use this plan to configure a service processor or a chassis, provision OS and update the OS.
- **Configure and Install Dynamic System Domain:** Use this plan to create dynamic system domains, provision and update OS on the domains.
- **Install Server:** Use this plan to provision and update the OS.

You will complete the following actions to use plans to update your Oracle Solaris or Linux operating system:

1. Create an OS update profile.
2. Optionally, create operational profiles that contain scripts to perform preinstall and postinstall actions.
3. Create a deployment plan.
4. Complete the plan, select the target asset, then select the Deploy/Update Software action.

The settings and values in the profiles bound to each step are defaults. You can modify the settings and values when you apply the plan. The profile settings and values are constrained by the target systems to which the plan is applied. All update plans use "yes to all" policy.

About Operating System Update Jobs

Creating a new update job enables you to use custom or predefined profiles. Use this method for complex update scenarios or to apply updates consistently across many systems. You can run the OS update plan in simulate or deploy mode. In simulate mode, you can choose whether to download the updates.

The New Update OS Job option enables you to create customized update jobs. When creating a job, you define how the software performs the job, set the automation level of the job, and select a policy from the list of available policies. You can run the job in simulation mode or run the actual job. Simulation mode determines the actions and results of a job, and estimates the amount of time required to complete the job. You can use a job simulation to determine if your job can succeed based on your policy and profile responses. You can run a simulation with or without downloading updates

Note: To use an alternate boot environment (ABE) and run ABE pre-action scripts for Solaris OS, see the procedures in Oracle Solaris Boot Environments.

Updating an Operating System by Modifying a System Catalog

A system catalog contains a list of operating system software components that are installed on a managed system. Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems.

Updating an operating system by modifying a system catalog provides the following advantages:

- Enables you to create a quick ad hoc job
- Provides an easy method of applying a single patch, baseline, or package
- Enables you to update an operating system without creating a profile for a one-time job

Updating an Operating System From an Operating System Report Result

You can generate compliance reports for an operating system from which you can create an operating system update job.

The Oracle Solaris Update Compliance report is similar to a Recommended Software Configuration report. The report uses the Oracle Solaris update patch bundles as the recommended software configurations. You can use this report to check how compliant a system is with a particular Oracle Solaris update and bring the operating system into compliance.

See [Chapter 10, "Reports"](#) for information about generating these reports. You can generate these reports for non-compliant components. The report result appears with the option to install the updates, packages, updates, and incidents.

The report results are stored in the database associated with the Enterprise Controller. The software maintains a history of the reports for analysis purposes.

From the report result, you can initiate a job to install the non-compliant component updates. The New Update OS Job Wizard starts, enabling you to enter job information and to schedule the job. The required data for profiles, policies, and targets are automatically pre-populated in the New Update OS Job Wizard.

Note: Updating a version of Oracle Solaris is not the same as upgrading to a new version.

For example, you can run the Oracle Solaris Update Compliance reports for Oracle Solaris update releases and bring systems into compliance with those bundles. Oracle Solaris update release bundles contain the equivalent set of updates to the corresponding update. You can use them to bring pre-existing packages up to the same software level as the corresponding update. However, this feature does not perform a full Oracle Solaris upgrade from one release to another. The update release bundles do not contain additional packages that are in the update releases and they do not change the first line of `/etc/release` to specify an upgrade has taken place, although they do append a line to `/etc/release` to specify that the update bundle is applied.

Using a System Catalog

A system catalog is a list of operating system software components that are installed on a particular managed system. An initial catalog is created after the system is discovered and managed.

After an operating system is available and selected, you can view and modify the catalogs and create historical catalogs (snapshots of the system).

Modifying a catalog is an alternate way to run an operating system update job to install, uninstall, or upgrade a component. Modifying a catalog does not require an update profile to run the update job and is a quick way of changing the component configuration of a system.

You can compare the system catalogs of two managed systems, view the summary of the comparison, and you can choose to make the target system the same as the source system.

Catalogs provide the capability to directly manipulate the installed software components on a single operating system or a group of operating systems. Alternatively, a catalog can be saved as a profile, and then an operating system update job can be run using this profile.

You can run an operating system update job, or you can use the simulate feature to run an update simulation before you apply updates.

Create an Update Profile From a System Catalog

You can save the catalog of a system as a profile. Using this profile, you can create the systems with the required configuration in your data center.

Updating an Oracle Solaris Boot Environment

You always need an update profile to update an Oracle Solaris Boot environment. You can use the update profile in an update deployment plan or the Update Job Wizard.

You can update an alternate boot environment as part of a Software Deployment / Update deployment plan by selecting the alternate boot environment as the target. See the *Oracle Enterprise Manager Ops Center Updating Your Oracle Solaris 10 Operating System* for an example of how to use this plan.

Updating a Boot Environment Using an Update Job

You can create a customized update job, including the option to use an alternate boot environment (ABE) to perform a live upgrade of your Oracle Solaris 10 operating system. With Live Upgrade, you create an inactive ABE, update and patch the ABE, synchronize the ABE and BE, and then switch boot environments. When you switch boot environments, the patched and tested ABE becomes the active boot environment.

Note: Do not use Live Upgrade on your Enterprise Controller or Proxy Controllers. Live Upgrade does not synchronize all of the files that are required for these components.

You must run a separate update job for systems that use an ABE from those that do not use an ABE. When creating a job, you must define the following job parameters:

- Name and description of the update job.
- Alternate Boot Environment: Whether or not you want to use an alternate boot environment.
- Profile: Defines what updates are to be installed, uninstalled, or updated on an operating system. Select a profile from the list of predefined and customized profiles.
- Policy: Defines how a job is performed and sets the automation level of the job. Select a policy from the list of available policies. You can also create your own policies.
- Target Settings: Defines whether the target is different or similar for each task in the job.
- Actual Run: Defines whether this job is in simulation mode. You can choose to deploy the job, or to run a job simulation. A job simulation determines the actions and results of a job, and estimates how much time is required to complete the job. A job simulation also indicates whether your policy and profile responses will enable the job to succeed.
- Task Execution Order: Specifies whether the tasks is run in parallel or sequentially.
- Task Failure Policy: Specifies the action to take if a task fails.

- **Targets:** Select one or more target hosts for this job.

To create an ABE as part of this job, you must write at least one script that uses the `lucreate` command and then upload the script to the Local Content.

Note: The ABE name defined in the script must match the ABE name that you use when you run the update job to create the ABE.

Perform the following steps to update a boot environment:

1. Click **Assets** in the Navigation pane.
2. Expand All Assets, or use the All Assets filter to locate the Oracle Solaris 10 operating system instance.
3. Click **New Update OS Job** from the Actions pane. The New Update OS Job Wizard is displayed. The Job Information window is displayed first.
4. Complete the following Job parameters:
 - Type a job name.
 - Select the Run Type:
 - Simulation. To download the required updates as part of the simulation, select the Download check box.
 - Actual Run. Updates the operating system.
 - Select the task execution order:
 - Sequential
 - Parallel
 - Choose the Target Setting:
 - Use the same Targets for all tasks in the job
 - Use different Targets for each task in the job
 - Choose the Task Failure Policy:
 - Complete as much of the job as possible
 - Stop at failure and notify
 - Select the ABE check box.
 - (Optional) To create an alternate boot environment during this job by running an ABE Pre-Action Script, click the Enable check box.

Note: You must create the script and upload it to the library before you can use this option.

5. Define the profile, policy and target for each task, or edit the profile and policy.
6. (Optional) To edit the profile or policy of the default task, click the Profile or Policy cell for the task to display a drop-down menu. Select the profile or policy from the menu.
7. (Optional) To add a new task, click the **Add (+)** icon.

- A second row appears. Click the Profile cell for that row to display a drop-down menu. Select the new profile that you want to add.
 - To change the policy for the new profile, click the Policy cell and select a new policy from the drop-down menu.
 - When you chose the parameter to use a different target for each task, click the Targets cell to display the Select Targets page. Select one or more target from the list of Available Items, then click Select to include the asset in the Target List. Click **Add to Target List** to close the page.
 - Click **Next**.
8. If you selected the option to create an ABE as part of the job, the Create ABE page appears.
 9. When you have only one ABE, the Boot Environment Workflow page appears, go to step 10. When you have multiple alternate boot environments, the ABE Selection page appears.
 - One or more of the targets has more than one possible associated ABE. Select the ABE from the drop-down menu for each of the Targets. You can use the **Select ABE** field to filter for the ABE name.
 - Click **Next**. The Boot Environment Workflow page is displayed.
 10. If you selected Simulation in the job parameters, the boot environment workflow cannot be edited. Skip to step 12.
 11. If you selected Actual Run in the job parameters, you can edit the pre-actions and post-actions in the workflow.
 - Pre-actions by default will unmount and then mount the ABE. To synchronize the ABE with the BE before mounting, click the **Sync ABE** check box.
 - Post-Actions by default will unmount the ABE.
 - Click **Modify Current BE** to edit the description of the current boot environment. You might use this to describe the state of the current BE. For example, Boot environment running Oracle Solaris 10 5/08 operating system before applying the Oracle Solaris 10 operating system September baseline.
 - Click **Modify Alternate BE** to edit the description of the ABE. You might use this to describe the state of the ABE. For example, boot environment running Oracle Solaris 10 5/08 operating system after applying the Oracle Solaris 10 operating system September baseline.
 - Click **Activate and Reboot ABE** to switch boot environments after update.
 12. Schedule the job, then click **Next**.
 - Run Now starts the job immediately after you click Finish in the Job Summary.
 - Start Date enables you to select a date and time to start the job.
 - On a recurring schedule enables you to run the same job on a monthly or daily scheduled time.
 13. Review the Job Summary, then click **Finish** to run the job as scheduled in the previous step.

Updating Microsoft Windows Operating Systems

You can update your managed Microsoft Windows operating systems by using the Microsoft System Center Configuration Manager (SCCM) and Windows Management Instrumentation (WMI) software.

Oracle Enterprise Manager Ops Center uses the Microsoft SCCM 2007 and WMI software to update your managed Windows operating systems. The Windows Update function depends on the SCCM's agent installed on the managed systems. You can configure SCCM to install agents on your managed Windows systems either automatically or through a manual process.

You must have access to SCCM software that is configured for software updates. The Enterprise Controller connects to the SCCM software to get the latest updates and packages. The SCCM software connects to the Microsoft website through the Internet and downloads the metadata that is used for compliance analysis. You can connect Oracle Enterprise Manager Ops Center to the Microsoft website to download updates that are then handled by SCCM for installation.

You do not need any authentication to access the Microsoft website. However, you must provide authentication information to access the SCCM server.

About Windows OS Update Jobs

Oracle Enterprise Manager Ops Center contains the following options in an update job to maintain control and consistency across your data center:

- Groups: Help you to organize your assets in the user interface and act as targets for many types of jobs.
- Roles: Determine the tasks that you can perform on a specific piece of an asset or a group of assets.
- Reports: Enable you to run compliance reports and create update jobs from the compliance reports.

You can define the following job parameters while creating a windows update job:

- Name and Description: Identify the name of the report against which you want to create a Windows operating system update job. Provide a detailed description that clearly identifies the job in the historical record.
- Reboot behavior: Enables you to select the reboot behavior when a reboot is required after running the new update job. You can choose to reboot the system immediately following the update operation or to reboot the system at the default setting of the SCCM server.
- License Terms: Enables you to review the license terms and either accept or decline them. The License Terms window appears only when the updates in the report require you to review the License Terms.
- Schedule: Enables you to schedule when the update job runs.

Modify the Registry

Due to some changes that Microsoft introduced for registry key ownership, you must manually modify the registry and change the ownership permissions for the Administrators group.

Note: You must modify the registry on a Windows Server 2008 R2. Other Windows servers, such as 2008 Server SP2, do not require you to modify the registry.

Configure Oracle Enterprise Manager Ops Center for Updating the Windows Operating System

Oracle Enterprise Manager Ops Center uses the `j-Interop` command to access the Windows Management Instrumentation (WMI) and get Windows update information. Oracle Enterprise Manager Ops Center uses the software update capability of the Microsoft System Center Configuration Manager (SCCM) to update any managed Windows operating systems.

Before you can use the software to update your Windows systems, configure it to interact with the identified Microsoft System Center Configuration Manager (SCCM). In addition, you might need to modify the WMI registry.

To configure Oracle Enterprise Manager Ops Center to interact with the identified SCCM, you must have the following credentials:

- SCCM Server
 - Server Name
 - Domain Name
 - Site Name
 - User Name
 - Password
- SCCM Share
 - URL
 - Domain Name
 - User Name
 - Password

The configuration information appears in the Configuration tab of the Windows Update window.

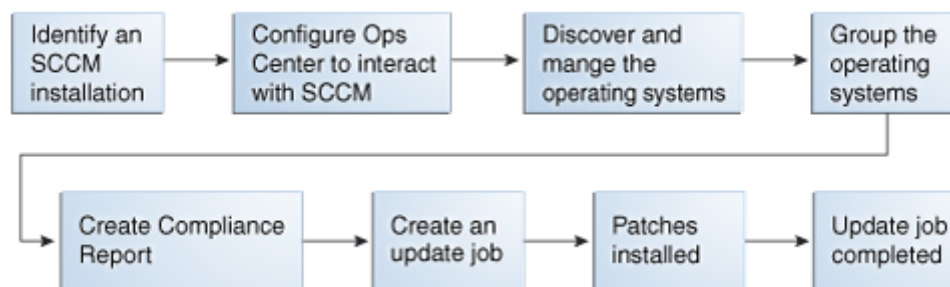
Note: Oracle Enterprise Manager Ops Center uses the same SCCM credentials to access the SCCM server and enable the SCCM share. Use the `<domain>` format for the Domain Name field. Do not use the `<domain>\<username>` format. Entering an incorrect format for the Domain Name field returns a configuration task. In this case, unconfigure the SCCM and configure the SCCM again with the correct format for the credentials.

Creating an Update Job for the Windows Operating System

You can use the output from compliance reports to update your Windows operating system to comply with the newly released updates.

From the results of the Windows Host Compliance Report and the Windows Incident Compliance Report, you can make your systems compliant by initiating an update job for the Windows operating system.

Figure 13–9 Process for Updating Windows Operating System



The Create New Windows Update Job Wizard enables you to create an update job. When creating a new update job, you must define the following job parameters:

- Name and Description for the new Windows software update job.
- Reboot behavior: Lets you select whether you want the system to reboot immediately following the update operation or at the default setting of the SCCM server.
- License Terms: Lets you review the license terms and either accept or decline them. The License Terms window appears only when the updates in the report require license terms that must be reviewed.
- Schedule: Lets you decide how you want to schedule the execution of the new update job.

Related Resources

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

- See [Chapter 5, "Software Libraries"](#) for how to add and update operating system packages and images.
- See [Chapter 12, "Operating Systems"](#) for details on managing operating systems, including [Using Agent Management for Operating Systems](#), [Overview of Oracle Solaris Boot Environments](#), [Overview of Oracle Solaris 11 Boot Environments](#), and [Overview of Oracle Solaris 10 Boot Environments](#).
- See [Chapter 10, "Reports"](#) for details on the operating system Update reports.
- See [Chapter 14, "Oracle Solaris Zones"](#) for information about zones and how you can use Oracle Enterprise Manager Ops Center to efficiently manage all phases of zones lifecycle.
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm

For in-depth information about these products, see the following Oracle documentation:

- Oracle Linux at <http://www.oracle.com/us/technologies/linux/index.html>
- Transitioning From Oracle Solaris 10 to Oracle Solaris 11 Guide at http://docs.oracle.com/cd/E23824_01/html/E24456/docinfo.html

- Oracle Solaris 11 Information Library at http://docs.oracle.com/cd/E23824_01/index.html
- Oracle Solaris 10 documentation at <http://www.oracle.com/technetwork/documentation/solaris-10-192992.html>
- Oracle Solaris 8 and 9 documentation at <http://www.oracle.com/technetwork/documentation/legacy-solaris-192993.html>

Part IV

Virtualize and Cloud Management

Part IV contains the following chapters:

- [Chapter 14, "Oracle Solaris Zones"](#)
- [Chapter 15, "Oracle VM Server for SPARC"](#)
- [Chapter 16, "Oracle VM Server for x86"](#)
- [Chapter 17, "Server Pools"](#)
- [Chapter 18, "Virtual Datacenters"](#)

Oracle Solaris Zones

The following information is included:

- [Introduction to Oracle Solaris Zones](#)
- [Roles for Zone Management](#)
- [Zone Management Functions](#)
- [Preparing Your Global Zone](#)
- [Discovering and Managing Existing Zones](#)
- [Outline of Zone Creation](#)
- [Determining Zone Requirements](#)
- [Zone Configuration Parameters](#)
- [Creating a Zone Profile](#)
- [Creating and Deploying Zone Plans](#)
- [Managing Zones](#)
- [Migrating Zones](#)
- [Recovering Zones](#)
- [Zones Server Pool](#)
- [Updating Zones](#)
- [Related Zone Operations](#)
- [Related Resources](#)

Introduction to Oracle Solaris Zones

Oracle Solaris Zones, also known as Oracle Solaris Containers, are used to virtualize operating systems and provide an isolated and secure environment for running software applications. A zone is a virtualized operating system environment created within a single instance of the Oracle Solaris operating system.

Think of a zone as a box with flexible, software-defined walls. One or more applications can run in this box without interacting with the rest of the system. Because zones isolate software applications or services, applications that are running in the same instance of the Oracle Solaris OS can be managed independently of each other. For example, you can run different versions of the same application in separate zones.

Zones require a machine that is running an Oracle Solaris 10 or later release.

Global and Non-global Zones

The global zone is the default operating system and has control over all of the processes and has system-wide administrative control. A global zone always exists, even when no other zones are configured.

Non-global zones, or simply zones, are configured inside the global zone. Zones are isolated from the physical hardware by the virtual platform layer. A zone cannot detect the existence of other zones.

Types of Non-Global Zones

You can create different types of non-global zones for different purposes:

- **Sparse Root Zone:** Contains a read/write copy of a portion of the file system that exists on an Oracle Solaris 10 global zone. Other file systems are mounted read-only from the global zone as loop-back virtual file systems. As part of creating a sparse root zone, the global administrator selects which file systems to share with the sparse root zone and the default read-only file systems: `/usr`, `/lib`, `/sbin`, and `/platform`. All packages that are installed on the global zone are available to the sparse root zone; a package database is created and all files in the mounted file system are shared with the zone.

Note: Sparse root zones are not available beginning with Oracle Solaris 11. You can create sparse root zones only in Oracle Solaris 10.

- **Whole Root Zone:** Contains a read/write copy of the entire file system that exists on the global zone. When a whole root zone is created, all packages that are installed on the global zone are available to the whole root zone; a package database is created and all files are copied onto the whole root zone for the dedicated and independent use of the zone.
- **Branded Zone:** The non-global zone runs the same operating system software on the global zone. The branded zone facility is used to create non-global branded zones that contain operating environments different from that of the global zone. For example, you can install Oracle Solaris 8, 9, or 10 in a branded zone.

Roles for Zone Management

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

Table 14–1 Oracle Solaris Zones Roles and Permissions

Task	Role
Create, manage, update, and delete zones	Virtualization admin
Provision and manage virtualization host	Virtualization admin
Discover and manage virtualization hosts	Asset admin
Create and manage zone profiles and plans	Profile and plan admin
Create and manage IPMP groups	Network admin
Create and manage Link Aggregation	Network admin

Table 14–1 (Cont.) Oracle Solaris Zones Roles and Permissions

Task	Role
Set monitor thresholds	Asset admin

Zone Management Functions

Oracle Enterprise Manager Ops Center provides the following solution for managing your zones environment in a data center:

- Create zones using profiles and deployment plan.
- Discover and manage existing zone environments.
- Perform zones functions such as booting, rebooting, shutting down, cloning, migrating and delete zones from the software UI.
- Manage zone configuration such as file systems, storage, and networks from the UI.
- Aggregate your NICs or create IP Multipathing (IPMP) groups.
- Create zones server pool for balancing available resources and provide the infrastructure support for virtual datacenter creation.
- Use UI and native CLI interchangeably to create and manage zones.
- Complete support for managing Oracle Solaris 11 zones.
- Options to upload scripts to manage the unmanaged file system and zone dependencies.

Preparing Your Global Zone

The global zone is the default Oracle Solaris 10 or 11 operating system installed on a system. When you provision systems, you can use the **Add Assets** or **Find Assets** option in Oracle Enterprise Manager Ops Center to discover the global zone. See [Chapter 12, "Operating Systems"](#) for information on provisioning a system with an Oracle Solaris operating system.

You must prepare your global zone with the network and storage resources that the non-global zones will use, including associating the required storage libraries with the global zone and attaching networks to the global zone. You can also aggregate the interfaces or create IPMP groups in the global zone before you connect them to the zones. The use of IPMP and link aggregated interfaces results in enhanced network availability for the zones.

This section covers the following topics:

- [Associating Storage Libraries](#)
- [Managing Global Zone Networks](#)
- [Creating IPMP Groups](#)
- [Creating Link Aggregation](#)
- [Modifying and Detaching a Network from the Global Zone](#)

Associating Storage Libraries

You can associate or disassociate libraries from a global zone. When you associate a library with the global zone, the libraries are available to all the zones in the hierarchy of the global zone.

You can associate the following types of storage libraries to a global zone:

- Filesystem storage: NAS libraries
- Block storage: SAN and Dynamic storage libraries

All the libraries associated with the global zone are available for association with the zones.

When you disassociate a library from the global zone:

- The libraries are not available to new zones.
- The libraries remain associated to existing non-global zones in that global zone.

To Associate a Library With the Global Zone

1. Select the global zone in the Assets section.
2. Click **Associate Libraries** in the Actions pane.

The Associate Library window lists the libraries that are not associated with the global zone.

3. Select one or more libraries to add.
4. Click **Associate Libraries**.

A job is submitted to associate the libraries with the global zone.

To Disassociate Libraries From Global Zone

1. Select the global zone in the Assets section.
2. Click the **Libraries** tab in the center pane.

The associated libraries with the global zone are listed.

3. Select a library from the list.

The Disassociate Library icon is enabled.

4. Click the **Disassociate Library** icon to display the Disassociate Library window.
5. Click **Disassociate Library** to confirm unmounting the library.

Managing Global Zone Networks

You can attach one or more networks to global zone using the Attach Networks option. When your networks are grouped as network domains in Oracle Enterprise Manager Ops Center, select the domain and choose the corresponding networks to be attached.

When you assign a network to a stand-alone global zone, you must define the mode of the network, either Shared IP or Exclusive IP:

- In Shared IP mode, the global zone shares its network interface with one or more zone. You must define the network interface when you assign the network to the global zone.
- In Exclusive IP mode, a dedicated network interface is allocated to the zone. You can choose the network interface when you assign the network to a zone.

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.

While you attach networks to a global zone, you can deploy IP Multipathing (IPMP) to obtain better network performance or link aggregation to provide increased reliability, availability, and network performance for systems with multiple physical interfaces.

For stand-alone global zones with Oracle Solaris 10 OS, network can be attached in shared or exclusive IP mode. Whereas with Oracle Solaris 11 OS, the network is always attached in exclusive IP mode. This is because a virtual NIC or VNIC is created when the zone boots and deleted when the zone is halted.

This scenario is different for attaching networks to global zones that are in a server pool. See [Attaching Networks](#) for more information about attaching networks in a zones server pool.

IPMP Groups

IPMP groups provide network failover for your global and non-global zones. You can configure one or more physical interfaces into an IPMP group in the global zone and extend that functionality to the non-global zones. After configuring the IPMP group, the system monitors the interfaces in the IPMP group for failure. If an interface in the group fails or is removed for maintenance, IPMP migrates, or fails over, the failed interface's IP addresses. The failover feature of IPMP preserves connectivity and prevents disruption of any existing connections. The network access changes from the failed interface to the standby interface in the IPMP group and the data address of the failed interface migrates to the standby interface. See [IP Multipathing Groups](#) and [Creating IPMP Groups](#) for more information about IPMP groups.

Note: You can create IPMP groups when you attach network in exclusive IP mode on Oracle Solaris 10 and 11 operating systems or you can create an IPMP group without using the attach network option. IPMP groups are not available for shared IP mode networks.

When you attach networks to a global zone in a shared IP mode, you can create an IPMP group and configure the following characteristics:

- Define the number of members for the IPMP group.
- Select the NICs that are part of the IPMP group. The NICs selected are placed in the IPMP group.
- Select the type of interface:
 - Active interface enables you to provide the data address and choose whether failover must be enabled for the interface for uninterrupted access to the network. This is applicable for Oracle Solaris 10 OS and is inherent in Oracle Solaris 11 OS.
 - Standby interface enables you to provide a test address when you want probe-based detection.

Multiple shared IP zones can use the IPMP group. When you create a zone or connect a network to a zone, the IPMP group appears in the NIC list. Select the IPMP group from the list to connect to the selected network.

Link Aggregation

Several interfaces in a system can be aggregated into a single logical link. The aggregation is done as per the standard defined in IEEE802.3ad. When interfaces have been aggregated, they are treated as a single network interface. Oracle Enterprise Manager Ops Center displays the link aggregation in the list of available NICs as if it were an individual interface. You can assign a network with a link aggregation to a non-global zone, and select the link aggregation from the NIC list.

When you attach a network to a global zone, you can aggregate the physical interfaces and attach to the network. Link aggregation is available for both shared and exclusive IP network stack.

To Attach a Network to a Global Zone

When you attach a network to a global zone, you can choose to create an IPMP group or aggregate link. See [Creating IPMP Groups](#) to create IPMP groups on the global zone without using the attach network option.

The following network options are available:

- **Exclusive network:** For zones that use an exclusive IP network, each zone has a dedicated network interface. You can configure the interfaces when you connect a zone to the exclusive network. An IPMP group requires an exclusive network.
- **Shared network:** The global zone shares the network interface with the zones that use a shared IP network. You must configure the network interfaces of the global zone.

Perform the following steps to attach a network to a stand-alone global zone:

1. Select the global zone in the Assets section.
2. Click **Attach Network** in the Actions pane.
The Attach Network Wizard is displayed.
3. Select the network domain and the corresponding networks that are not yet connected to the asset are displayed in the network list. Select a network from the list.

You can attach one or more networks to the global zone. Click **Next** to configure the networks.

4. For each selected network, select the following information:
 - Select the IP stack as **shared** or **exclusive**.
 - Select the network configuration as **IPMP**, **Link Aggregation**, or **None**. For an exclusive IP stack, only the Link Aggregation network configuration is available.
 - When you select IPMP network configuration, you are taken through the steps for IPMP configuration.
 - When you select Link Aggregation configuration, you are taken through the steps to configure the aggregation.
 - When you select None, you are taken to Step 8 to configure the network interfaces for shared IP stack or to the Summary step.
 - Enter the number of connections for each network.

The number of connections depends on the Oracle Solaris OS version. For Oracle Solaris 10 OS, you cannot make multiple connections. For Oracle Solaris 11 OS, you can make multiple network connections.

Click **Next** to specify the link aggregation.

5. Provide the following information for link aggregation:

- The link aggregation name.
- Select the load balancing policy.

Click **Next** to configure the link aggregation.

6. Specify the following information for link aggregation:

- LACP mode and timer
- MAC address policy and the MAC address if required.

7. Specify the IPMP group details.

- IPMP group name
- Number of members of the IPMP group
- Select whether the probe-based failure detection must be enabled.

Click **Next** to configure the interfaces.

8. For IPMP group, provide the following information:

- The NICs that are member of the group.
- Select the interface for the standby interface. You must have at least one active interface.

Click **Next** to specify the data addresses.

Configure the interfaces for networks that do not have network configuration:

- Specify the NIC and the IP address for the network connection.
- You can select **System Allocated for the NIC** and **Assign by DHCP** for the system to take care of the NIC and IP address allocation.
- Select **Do Not Allocate IP** for the IP address, when you do not want to allocate IP address for the network connection.

9. Enter the data addresses that must be used for the active network interfaces. Also, specify whether failover must be enabled for the interface. You must have more than one interface to enable failover.

10. Review the information provided for attaching a network to the global zone and click **Finish** to attach the network.

To prevent a failure of the attach network job, ensure that you enter the correct information for different versions of Oracle Solaris OS.

Creating IPMP Groups

You can directly create IPMP groups on the global zone without using the attach network option. From the network tabs, you have the option to create and manage the IPMP groups in the selected Oracle Solaris OS.

Figure 14–1 shows the options that are available to create and manage IPMP groups.

Figure 14–1 IPMP Group Option

IPMP Group Name	Failure Detection	Network
PSG	Link-Based	10.168.89.0/24

IPMP provides physical interface failure detection, transparent network access failover, and packet load spreading for systems with multiple interfaces. Oracle Enterprise Manager Ops Center provides option to create IPMP groups. You can configure one or more interfaces into an IPMP group. The group functions like an IP interface with data addresses to send and receive network traffic. When an underlying interface in the group fails, the data addresses are redistributed among the remaining underlying active interfaces in the group. Thus, the group maintains network connectivity despite an interface failure. With IPMP, network connectivity is always available, provided that a minimum of one interface is usable for the group. IPMP also provides load spreading for the outbound network traffic across the network interfaces in the group.

To create an IPMP group, you must define the following parameters for the group:

- The active and the standby interfaces of the group. By default, an interface added to an IPMP group is active. You can configure as many standby interfaces as you want for the group.
- The link-based failure detection is enabled by default. You must select whether you want to enable Probe-Based failure detection. For probe-based failure detection, you must provide the test address to track the interface status.
- You must assign the data addresses for the physical interfaces in the IPMP group. Data traffic flow uses the data addresses that are hosted on the IPMP interface and flow through the active interfaces of that group.

To Create an IPMP Group

1. Select the Oracle Solaris OS in the Assets section.
2. Click the **Network** tab in the center pane.
3. Click the **IPMP Groups** subtab in the pane.
Existing IPMP groups in the OS are listed.
4. Click the **Create IPMP Group** icon to create an IPMP group.
The Create IPMP Group Wizard is displayed.
5. Enter the following details for the IPMP group:
 - Provide a name for the IPMP group.
 - Select a network from the list.

The list of available network interfaces in the global zone are displayed.

- The Link-Based failure detection is always enabled by default. Select whether you want to enable Probe-Based failure detection.
- Select the interfaces that are part of the IPMP group.

Figure 14–2 shows the Specify IPMP Group Wizard step.

Figure 14–2 Specify IPMP Group Details

Click **Next** to specify the NIC settings.

6. When you have enabled probe-based failure detection, then enter the test address for the NICs in the group.
7. Select the interfaces that are in standby mode.
You must have at least one active interface in the group. Click **Next** to specify the data address.
8. Enter the data address for the active interfaces of the group and select whether the interface has a failover.
Click **Next** to view the summary.
9. Review the information and click **Finish** to create the IPMP group.

Note: For Oracle Solaris 11 OS, you cannot create IPMP groups. Instead, you can aggregate the links.

Creating Link Aggregation

You can also create link aggregation directly on the Oracle Solaris OS. to create link aggregation, you must define the following details:

- Load balancing policy
- LACP mode and timer
- MAC address policy and if required, the MAC address

To Create a Link Aggregation

1. Select the global zone in the Assets section.
2. Click the **Network** tab.
3. Click the **Link Aggregation** subtab.
4. Click the **Create Link Aggregation** icon.

The Create Link Aggregation Wizard is displayed.

5. Enter the name of the link aggregation.

The name starts with "aggr" by default. Append a number to make the name unique.

6. Select the NICS for the link aggregation.

Click **Next** to configure the link aggregation.

7. Specify the following information for link aggregation:

- Select the load balancing policy.
- LACP mode and timer.
- MAC address policy and MAC address if required.

Click **Next** to view the summary.

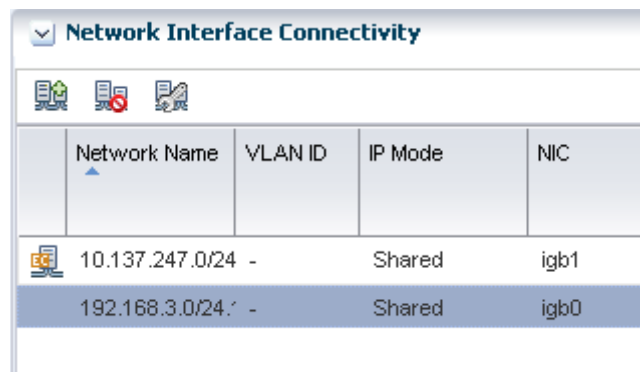
8. Review the information and click **Finish** to create the link aggregation.



Modifying and Detaching a Network from the Global Zone

You can modify the network connection that are attached to the network except for the management network. The modify and unbind network options are available in the Network tab.

Figure 14–3 shows the options that are available to manage networks of a global zone.

Figure 14–3 Network Options



Network Interface Connectivity				
	Network Name	VLAN ID	IP Mode	NIC
	10.137.247.0/24 -		Shared	igb1
	192.168.3.0/24 -		Shared	igb0

To Modify Physical Connectivity

1. Click the **Network** tab of the selected global zone in the Assets section.
2. Click the **Modify Physical Connectivity** icon

Figure 14–4 shows the Modify Physical Connectivity window that displays the network details that can be modified.

Figure 14–4 Modify Physical Connectivity

Oracle Enterprise Manager Ops Center - Change Virtual Host Network/NIC Connection

Change Network-NIC Connection on

Network Name: 192.168.3.0/24.1

Physical NIC: igb0

IP Address Allocation Method: Use Static IP

IP Address: 192.168.3.58

DHCP ID:

Submit Cancel

You can change the permanent IP address or change the allocation method.

Note: When the network interface is an IPMP group, you cannot modify the network interface.

Unbind Network connection

When you unbind a network from a global zone, the global zone's non-global zones are also disconnected from the network. Select the global zone and click the **Unbind Network** icon to remove the network.

Discovering and Managing Existing Zones

When you discover a global zone that has existing zones, these zones are also automatically discovered and displayed on the UI. The option to deploy the agent on the global zone is inherited to the zones.

The following steps provide an outline of the procedure to discover and manage existing zones:

1. Use the option **Add Assets** to discover a global zone.
2. In the Add Assets Wizard, select whether you want to deploy the agent on the global zone or you want agentless management.
3. All the non-global zones in the selected global zone are automatically discovered.
4. When you want to deploy the agent on the global zone, the agent is installed on all the zones that are in running state. When you want to deploy the agent for the zones in shutdown state, boot them and use the option Switch Management Access. Using this option you can change the state from agentless management to managed by agent mode.
5. You can always use the option **Switch Management Access** to switch between managed by agent and agentless management mode.

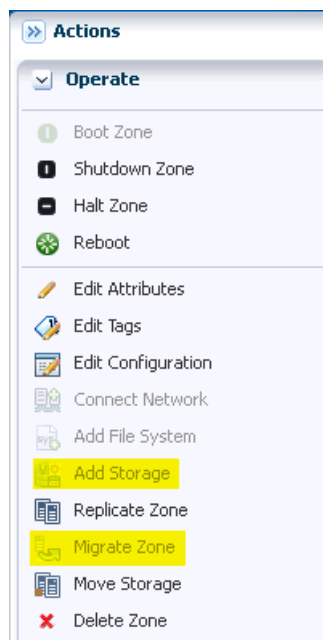
You cannot create zones on a global zone that is managed agentless. For an agentless managed zone, you can boot, shutdown, halt and delete the zone. Agent is required on the zone when you want to use the full range of OS update actions on the OS.

When you use the native CLI to create a zone, the new zone is automatically reflected in the UI. You can interchangeably use the CLI and the UI to perform zone functions. The UI reflects the change in the zone configuration and state.

You can also use the **Find Asset** to discover an asset with Service Tags. See [Chapter 2, "Asset Management"](#) for more information about discovering and managing an asset in Oracle Enterprise Manager Ops Center.

As shown in [Figure 14–5](#), the discovered zones appear in the UI and the following actions are enabled for the zone: Shutdown Zone, Halt Zone, Reboot, Edit Attributes, Edit Tags, Edit Configuration, Replicate Zone, Move Storage, and Delete Zone.

Figure 14–5 Enabled Actions



In [Figure 14–5](#) the Migrate Zone and Add Storage options are not enabled because the zone storage source is unmanaged. To enable these actions, you must move the storage source to managed using the option Move Storage.

See [Zone Management Functions](#) for more information about performing zone management operations.

Deleting or Unmanaging a Global Zone

When you delete or unmanage the global zone, all the non-global zones under it are also automatically unmanaged.

You do not need to unmanage the non-global zones before unmanaging the global zone. See [Chapter 2, "Asset Management"](#) for more information about deleting an asset.

Unmanaged Zone Storage Information

When you discover and manage existing zones, Oracle Enterprise Manager Ops Center handles the zone metadata and networks except for the zone storage. The metadata of the existing zones is stored in the local storage of the global zone. Whereas, the zone's storage is considered to be unmanaged storage source.

Even for the zones that are created using Oracle Enterprise Manager Ops Center, when you add a storage or file system manually using the native CLI, the storage becomes an unmanaged storage or file system.

See [Moving Zone Storage](#) for more information about how to move the unmanaged storage to managed storage.

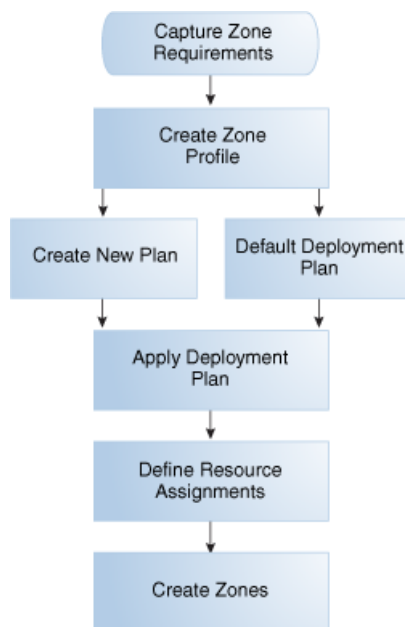
Outline of Zone Creation

Using Oracle Enterprise Manager Ops Center, you create a zone profile that captures the zone configuration. Use the profile in a deployment plan to create one or more zones simultaneously. The benefit of using a profile is that you can create multiple zones with consistent configuration.

Note: When you create or migrate a zone using Oracle Enterprise Manager Ops Center, you cannot use Oracle Solaris Live Upgrade and an alternate boot environment (ABE) to update the zone.

Create a zone profile that captures all the zone configurations. Then, create a deployment plan with the zone profile. The deployment plan is a single step plan which covers only the zone profile. During the application of the deployment plan, select the number of zones to create. In the plan deployment, you can correct the storage and network details as required.

Figure 14–6 Zone Creation Process



Determining Zone Requirements

Before you create a zone profile, determine the requirements for the zone. The requirements varies according to the Oracle Solaris OS version and the type of zone. Refer to the following section according to the Oracle Solaris OS version on which you create the zones:

- [Requirements for Zones on Oracle Solaris 10 OS](#)

- [Requirements for Zones on Oracle Solaris 11 OS](#)

Requirements for Zones on Oracle Solaris 10 OS

You can create sparse root, whole root, and branded zones on Oracle Solaris 10 OS. While creating the zone profile, select the appropriate options and provide the required resources for a successful zone creation.

Requirements for a Whole Root Zone

A whole root zone contains a read/write copy of the global zone's file system and has the following requirements:

- Minimum size of the file system is 5 GB.
- Minimum size of the virtual disk for the storage library is 6 GB.

Requirements for a Sparse Root Zone

A sparse root zone contains a read/write copy of a portion of the global zone's file system and shares the global zone's `/usr`, `/lib`, `/sbin`, and `/platform` directories in read-only mode. A sparse root zone has the following requirements:

- Minimum size of the file system is 1 GB.
- Minimum size of the virtual disk for the storage library is 1 GB.

Requirements for a Branded Zone

A branded zone emulates the user environment of earlier versions of Oracle Solaris. For example, you can create a branded zone to run Oracle Solaris 8 applications on your Oracle Solaris 10 system. Use the following procedure to prepare to create a branded zone:

1. Download the images for the operating system, as shown in [Table 14-2](#).
2. Import the images into one of the software libraries associated with the global zone.

During the process of creating a branded zone, you install the images in the global zone. For instructions, see the README files in the Oracle Solaris 8 or 9 Containers bundle.

Table 14-2 Packages for Branded Zones

Global Zone OS	Branded Non-Global Zone
Oracle Solaris 10 10/08 OS or later	For Oracle Solaris 8, install SUNWs8brandk from Oracle Solaris 8 Containers 1.0.1. For Oracle Solaris 9, install SUNWs9brandk from Oracle Solaris 9 Containers 1.0.1.
A version of Oracle Solaris before Oracle Solaris 10 10/08 OS	For Oracle Solaris 8, install SUNWs8brandr and SUNWs8brandu from Oracle Solaris 8 Containers 1.0.1. For Oracle Solaris 9, install SUNWs9brandr, SUNWs9brandu, and SUNWs9brandk from Oracle Solaris 9 Containers 1.0.1.

Requirements for Zones on Oracle Solaris 11 OS

You can install zones in Oracle Solaris 11 OS either using the IPS packages in the default Oracle Solaris 11 Package Repository or using an image of an installed system running the Oracle Solaris release.

To install zones using images, you must have an Oracle Solaris 10 image in flash archive format or from an existing Oracle Solaris 11 OS image in gzip format. See <http://www.oracle.com/technetwork/documentation/solaris-11-192991.html> for how to prepare your installed Oracle Solaris 10 or 11 systems, and create the `flar` or `gzip` image. Upload these images to the storage libraries and install branded zones in Oracle Solaris 11 OS.

Requirements for Installing Zones Using Repository

To install zones from the repository, configure your Oracle Solaris 11 Software Update Library in the Enterprise Controller to synchronize with the Oracle Solaris 11 Package Repository. When you install from repository, the default software group, `solaris-small-server` group is used. This occupies less space. The zones are whole root type only.

Requirements for Oracle Solaris 10 Branded Zone

You can migrate an Oracle Solaris 10 OS into an Oracle Solaris 11 environment. Create an archive of the Oracle Solaris 10 instance that you would like to migrate. In the Enterprise Controller software library, import an ISO image of Oracle Solaris 10 OS that has the same architecture of the instance to be migrated, for example, SPARC or x86 architecture. Then, import the `flar` archive with the parent as the ISO image imported previously.

For Oracle Solaris 11 OS, create an archive in the format of `gzip` file using `cpio` command. Import the archive into Enterprise Controller software library with the reference to the parent ISO image.

You can migrate only systems that have Oracle Solaris 10 10/09 or later versions. To migrate earlier versions, install the kernel patch 141444-09 (SPARC) or 141445-09 (x86/x64), or later version, on the original system.

Since zones do not nest, existing zones in the original system is detected and a warning is issued that nested zones are not usable and that the disk space can be recovered.

To use the Oracle Solaris 10 package and patch tools in your Oracle Solaris 10 Container, install patches 119254-75 (SPARC) and 119255-75 (x86/x64) on your source system before the image is created. The P2V process works without the patches, but the package and patch tools do not work properly within the `solaris10` branded zone.

To use Oracle Solaris 10 zones on your system, the `system/zones/brand/s10` package must be installed on the system running Oracle Solaris 11 OS.

Zone Configuration Parameters

When you create a zone profile, you can provide zone configuration details such as CPU model, memory caps, priority value, network mode, and storage resources for the zone.

- **Zone name:** Specify a unique name for the zone. Do not use names that start with `global` or `SUNW`.
- **Autoboot:** Specify whether the zone must boot immediately after you create it and whenever the global zone boots.
- **Shared CPU:** A zone with a shared CPU gets its CPU resources by the number of shares you allocate to it from the resource pool, which is also used by other zones. The new zone is added to the Fair Share Scheduling automatically. You have the option to set a maximum value for the CPU resources. The CPU cap limits the

amount of CPU resources that can be used by one zone. For example, a CPU cap value of 1 means 100% of a CPU.

- **Dedicated CPU:** A zone with a dedicated CPU gets exclusive use of the available CPU resources. You specify the minimum and maximum number of CPUs available to the new zone. A temporary resource pool is created and dedicated to the zone.

Note: This parameter is available when you select dedicated CPU. Set the priority of the zone. Assign an importance value for the zone so that when there are not enough CPU resources to satisfy all zones, the zone with the greater importance value receives a larger share of the available CPU resources.

- **Priority of recovery:** When you enable automatic recovery for the zone, the priority of recovery value decides which zone must be migrated first during a global zone failure in a server pool. Set the priority of recovery between 0 to 100.
- **Memory Caps:** Set the maximum value for physical, swap and lock memory resources.
- **Naming Service:** Specify the name service that the zone uses to communicate with network objects. You can select the DNS, NIS, NIS+ or LDAP naming service. To specify the name service, you require the domain name and the IP address of the name server.
- **Shared IP Address or Exclusive IP Address**
 - A zone with a shared IP address uses its global zone's IP layer configuration and state. The zone has a logical network interface to the IP address.
 - A zone with an exclusive IP has its own dedicated IP layer configuration and state. The zone has its own set of network interfaces. You must configure the network interfaces using the same network configuration methods applied to all Oracle Solaris OS configurations.

For Oracle Solaris 11 OS, only exclusive IP mode is supported for the network.

Creating a Zone Profile

Oracle Enterprise Manager Ops Center provides option to create profile that captures the zone configuration. Use the zone profile in a deployment plan and apply to create one or more zones simultaneously.

To Create a Zone Profile

1. Expand the **Plan Management** section in the Navigation pane.
2. Expand Profiles and Policies and click **Oracle Solaris Zone**.
3. Click **Create Profile** in the Actions pane.

The Create Profile – Oracle Solaris Zone Wizard is displayed.

4. Enter a name and description for the profile identification.

Select whether you want to create a deployment plan automatically using this profile.

5. Select the OS version of the zone.

Ensure that you apply the plan on correct target to create the zone. See [Table 14–3](#) for more information.

Table 14–3 Oracle Solaris Zones Targets

Zone OS Version	As	On Target
Oracle Solaris 10	Whole root zone	Oracle Solaris 10
	Sparse root zone	
Oracle Solaris 10	Branded zone	Oracle Solaris 11
Oracle Solaris 11	Whole root zone	Oracle Solaris 11
Oracle Solaris 8	Branded zone	Oracle Solaris 10
Oracle Solaris 9	Branded zone	Oracle Solaris 10

Click **Next** to specify the zone identity.

6. You can create one or more zones using the profile. To identify the zones, provide a zone prefix name and a number to start the series.

Each zone created uses the prefix name appended with the number that is incremented. For example, if the prefix name is Myzone and the number to start from is 1, then the zones are created with names Myzone1, Myzone2, and Myzone3.

Enter the description and tags for the zones. This is common for all the zones created using this profile.

Click **Next** to specify the zone installation source.

7. According to the OS version selected, the page displays the following information:

- Oracle Solaris 10

Select whether you want to create whole root, sparse root, or branded zone. For branded zone, provide the following information:

- Select the architecture as x86 or SPARC
- Select the branded zone image. You must have created a flash archive image of an installed Oracle Solaris 10 OS and uploaded to the software library in Oracle Enterprise Manager Ops Center. These images are listed in the Branded Zone Image list.
- Enter the prefix for the Host ID and the starting number.
- Select the machine type as sun4u or sun4v.

- Oracle Solaris 11

You must select the installation source for installing Oracle Solaris 11 zones. You have the following options to select:

- Install from repository. You must have configured your Oracle Solaris 11 Software Update Library in Oracle Enterprise Manager Ops Center in synchronization with the Oracle Solaris 11 Package Repository. This option installs the `solaris-small-server` software group by default. This is also referred to as the whole root zone for Oracle Solaris 11 OS.
- Install from selected image. You must have created `gzip` archive image of an installed Oracle Solaris 11 OS and uploaded to software libraries in Oracle Enterprise Manager Ops Center. Select an image from the list. You can create the `gzip` archive formats using `cpio` or `zfs` command options.

- Oracle Solaris 8 and 9

You can install Oracle Solaris 8 and 9 as branded zones on Oracle Solaris 10 OS only.

Download the images of Oracle Solaris 8 or 9 and upload them to a software library. Refer to [Requirements for Zones on Oracle Solaris 10 OS](#) for more information about requirements for branded zones.

Click **Next** to specify the zone configuration details.

8. Specify the CPU type, either Shared or Dedicated. Set the memory thresholds and verify that the locked memory threshold value is less than or equal to the physical memory threshold.

You can enable automatic recovery option for the zone. Set the priority of recovery value between 0 to 100. This value decides which zone is recovered first when the global zone fails.

Click **Next** to specify the zone file system.

9. The zone is created with a default root file system which is the zone path. You cannot delete this file system or change its read and write access. You can add more file systems from original zone's list of file systems. Set the size and access to the file system.

For each file system added, specify whether it is managed or unmanaged. When it is unmanaged, enter the mount point for the file system.

The Reserved size is the size of the file system that the user can reserve. The Quota size is the maximum size that the file system can utilize.

Click **Next** to configure the zone's storage.

10. Select the storage library for the zone and its metadata.

You can store the zone metadata in the local or NAS library only. To migrate a zone, it is required that you store the zone metadata in a NAS storage library.

All the libraries that are available in the Oracle Enterprise Manager Ops Center are listed. Select the library type and the corresponding libraries are listed. Select the library and the virtual disk for the zone. You can select library types such as NAS, SAN, Local, Local Devices, and Dynamic Storage.

Note: When you specify a local storage library, you cannot migrate the zone in the future.

Caution: An Oracle Solaris Zone running with a zone path on a NFS share is not a supported configuration. When you specify this type of storage, do not use the zone for production or non-experimental workload.

11. When the library is local or NAS storage, specify the virtual disk name and size of the disk. For SAN library, select a LUN from the available list. You cannot change the size of the LUN and the size of the local devices.

For Dynamic Storage, select the Dynamic library that is available for the discovered storage servers in Oracle Enterprise Manager Ops Center. You can select the existing LUNs or create new LUN. Select Create LUN in the drop-down

list of the column LUN/Virtual Disk Name and select the volume group. The size of the volume group is automatically displayed.

Click **Next** to specify the zone networks.

12. Select the networks that you want to connect to the zone. Also, specify the number of connections to the zone for each network. The actual binding of the networks takes place during deployment of a plan with this profile.

Click **Next** to specify the zone setup parameters.

13. Specify the following setup parameters:

- Language, time zone, terminal-type and root password for the zone.
- Provide a domain name for the NFSv4 Domain Name or accept the default value dynamic to allow the naming service that you specify in Step 14 to determine the NFSv4 domain at run time.
- Set the boot properties for the zone. You can set the zone to boot after it is created or whenever the global zone boots.
- For Oracle Solaris 10 version, you can select whether to install the agent on the zone or not. When the agent is not installed, the global zone access the zone through zlogin.

Click **Next** to specify the naming services.

14. Specify the naming service for the zone: DNS, NIS, NIS+ or LDAP naming service. You can also choose not to specify a naming service.

- **DNS:** Enter the domain name of the DNS server and the IP address of the DNS server. You can enter up to three IP addresses in the Name Server field. To specify additional domains to search, enter up to six domain names for the Domain Name Search List. The total length of each entry cannot exceed 250 characters.
- **NIS and NIS+:** Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the Specify an NIS Server option to provide the NIS server host name and its IP address. When you do not have the NIS server information, select the Find an NIS Server option.
- **LDAP:** Enter the domain name of the LDAP server. Specify the name of the LDAP Profile. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.
- **NONE:** Select this option so that no name server is configured.

Click **Next** to view the summary of the selected parameters for creating a zone profile.

15. Review the information and click **Finish** to create the zone profile.

The zone profile is created with version 1 and a corresponding deployment plan also with version 1. Apply the deployment plan to create one or more zones of consistent configuration.

Creating and Deploying Zone Plans

Apply the zone deployment plans on the required number of targets. When you apply the deployment plan, you are required to provide the resource assignments for storage and network. When there are resources that are provided from the profile are not

available or not accessible, it is flagged in red and you must re-assign the resources to continue further.

Use zone deployment plans to modify the parameters such as storage, networks, and zone name. The zone deployment plan is a single step plan which collects details of the number of zones to be created. You can also create your own zone deployment plans.

To Create a Zone Deployment Plan

1. In the Plan Management section, expand **Deployment Plans** and click **Oracle Solaris Zone**.
2. Click **Create Plan from Template** in the Actions pane.
3. Enter a name and description for the plan.
4. Select the failure policy.
5. In the Deployment Plan Steps, select the **Oracle Solaris Zone** profile.
6. Enter the number of zones to create.
7. Click **Save** to save the deployment plan.

When you modify the zone profile, you can choose to update the deployment plan with the correct version of the zone profile.

Zone deployment allows you modify the parameters such as storage, networks, and zone name. When you apply the zone deployment plan, resource assignments that are not available is marked in red and it must be corrected.

To Apply Zone Deployment Plan

1. Select the zone deployment plan and click **Apply Deployment Plan** in the Actions pane.

2. Select the target asset from the list and click **Add to Target List**.

You can add more than one asset to the list.

3. Select how to apply the plan.

You can either apply the plan with minimal interaction or override the profile values. When you select to override the profile values, you are taken through each step of the profile. Otherwise, you are directed to provide only the required resource assignments. In this procedure the minimal interaction is taken into consideration. Click **Next**.

4. In the Specify Storage Step, correct the storage resources defined in the profile if required.

For example, when the selected storage libraries in the profile are not associated with the target global zone., then you must modify the storage resources accordingly.

Click **Next** to specify the zone networks.

5. Designate the IP stack for the zone as Shared or Exclusive. Select the IP stack and the network list is updated accordingly.

For Shared IP network:

- Select a network from the list.
- The address allocation method is Use Static IP only. This is fixed for shared IP networks.

- Enter the IP address. When you create more than one zone, enter the IP addresses either in the form of range, separated by comma or both. For example, 192.168.1.2 - 192.168.1.5, or 192.168.1.10, 192.168.1.210, 192.168.1.167.

For Exclusive IP network:

- Select a network from the list.
- Enter the number of times the zone connects to the network. This is applicable only for exclusive IP networks.

Note: For Oracle Solaris 11 OS, you can connect to networks configured as exclusive IP stack only.

- Select the type of address allocation as Use static IP or DHCP allocated.
- When the allocation is static, then the IP range for the zones is proposed. The range depends on the number of zone and number of connections for each zone. You can modify the IP range as required.

Click **Next** to define the network resource assignment.

6. Modify zone host name and network assignment as required. You can also add or remove the network assignment to a zone.

For each network connection of the zone, the NIC and the individual IP address is displayed. Each zone's host name is also editable. Modify the zone resource assignment as required.

Click **Next** to schedule the job.

7. Select to run the zone creation job now or schedule to a later time.

Click **Next** to view the summary.

8. Review the information provided for creating zones. Click **Apply** to start the zone creation tasks.

Zpool and File System of Zones

Creating zones with Oracle Enterprise Manager Ops Center results in the following operations which run automatically in the background:

- The file systems of the zones are implemented as ZFS file systems.
- A zpool is created for each zone with the name of the zone metadata.
- The storage that is allocated to the zone is pooled in a zpool and used by all the file systems.

For each deployed zone you can view the following information: file systems, zpools with reservations and quotas, the storage library used, and the file system usage. The Storage tab for a zone shows the zpools with reservations and quotas. The Analytics tab for a zone contains information on file system usage. Use the **Move Storage** action to change the reservation and quotas.

To View the File System and Storage Added to a Zone

1. Select the zone in the Navigation pane.
2. Click the **Storage** tab.

As shown in [Figure 14–7](#), the Storage tab has two tables: a File Systems table and a ZPool and Storage table. The file system table displays the file system, the real path of the file system on the global zone, the amount of Reserved GB, the quota in GB, and the Access permissions for the zone.

Figure 14–7 Storage Tab for a Zone

File Systems

File System	Real Path on Global Zone	Reserved (GB)	Quota (GB)	Access
ZPool: nfss10zones0 (1 Item)				
/	/var/mnt/oc-zpools/c1f2a2fa-d606-4ba3-8791-a...	8	8	Read/Write

ZPool and Storage

The following table shows virtual disks that make up the storage of each ZPool or directly attached raw device.

Storage

Virtual Disk Name	Library	Size (GB)
Storage Type: ZPool (1 Item)		
nfss10zones0-vdisk1	NFS-library	20

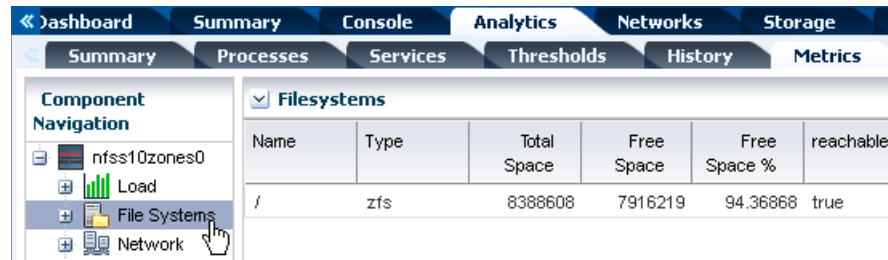
To Change the Default Reservation and Quota

1. Select the zone in the Navigation pane.
2. Click the **Storage** tab.
3. Click **Move Storage** in the Actions pane.

To View the File System Usage for a Zone

1. Select the zone in the Navigation pane.
2. Click the **Analytics** tab, then click the **Metrics** subtab.
3. Click **File Systems** in the Navigation pane.

Figure 14–8 Zone File System Metrics



Modify Zone Configuration

You can modify the configuration of a zone to change the CPU or memory resources, scheduler, and the identity of the zone. When you modify the CPU and memory configuration, you must reboot the zone for the changes to take effect.

You can modify the following properties of a zone:

- **CPU Model:** The CPU model can be shared or dedicated. For a shared CPU, you can modify the number of CPU shares allocated to the zone. For a dedicated CPU, you can edit the minimum and maximum dedicated CPU numbers. You can also set the relative importance of the zone compared to other zones when contending for CPUs.
- **Memory Caps:** You can specify caps on various aspects of memory used by the zone. You can modify the caps set for physical, swap, and locked memory.
- **Scheduler:** For shared CPU model, the scheduler is assigned to Fair Share Scheduler (FSS). For a dedicated CPU model, you can set the following scheduler attributes:
 - Fair Share Scheduler (FSS)
 - Fixed Priority (FX)
 - Interactive (IA)
 - Real-time (RT)
 - Timer Sharing (TS)
- **LightWeight Processes (LWP):** You can set the maximum number of LWPs simultaneously available to a zone.
 - **Message IDs:** Set the maximum number of message queue IDs.
 - **Semaphore IDs:** Set the maximum number of semaphore IDs.
 - **Shared Memory IDs:** Set the maximum number of shared memory IDs.
 - **Shared Memory:** Set the maximum amount of shared memory.
- **Automatic Recovery:** Set the value of priority of recovery. When the server fails, the zone with highest priority is recovered first.

Modify Zone Attributes

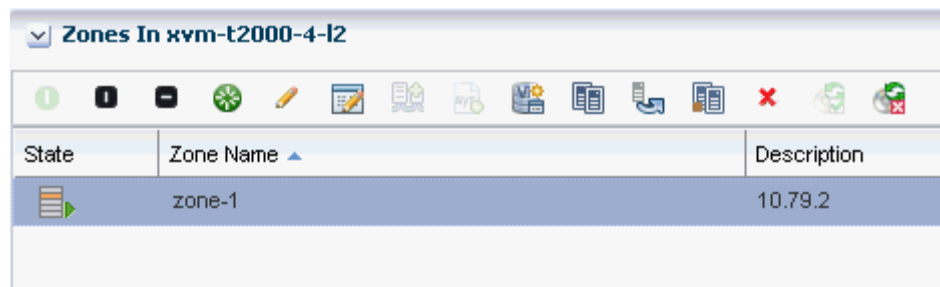
Use the option **Edit Attributes** to modify the description and tags of the zone. You can also add new tags to the zone.

Managing Zones

After creating the zones, you can perform various operations such as boot, reboot, shut down, and halt on the zone. Oracle Enterprise Manager Ops Center provides all these options to be performed from the UI. The UI updates the status of the zone when you perform these operations from the CLI.

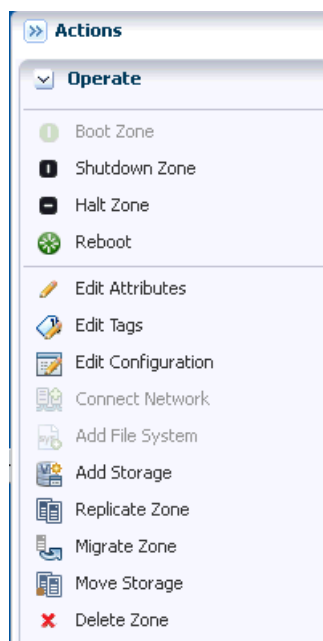
Select the global zone Summary tab in the center pane. The actions available for the zone are displayed as in [Figure 14–9](#).

Figure 14–9 Zone Management Functions



The same actions are available in the Actions pane when you select the zone.

Figure 14–10 Zone Actions Pane View



Simple zone operations that can be performed from the UI are as follows:

Boot Zone

Booting a zone places the zone in the running state, using the current configuration. This option is enabled only when the zone is in the ready or installed state. The zone boots whenever the global zone boots depending on the autoboot properties set during zone creation.

Reboot Zone

You can reboot a zone that is in the running state. The zone is shut down and then booted. This is different from the `zoneadm reboot` command in which the zone is first halted and then booted.

Halt Zone

When you halt a zone, it removes both the application environment and the virtual platform of the zone. Halting a zone changes the zone's state to Installed, all processes are killed, devices are unconfigured, network interfaces are unplumbed, file systems are unmounted, and the kernel data structures are destroyed.

Shutdown Zone

Shut a zone down in a graceful manner so that it is in a state that can be restarted.

Delete Zone

When you delete a running zone, the zone is halted, uninstalled, then deleted from the global zone. The following changes are also made:

- Zone root file system is deleted.
- Other file systems that were added to the zone are deleted.
- Zone metadata is deleted from the storage library.
- The zpool for the zone is deleted and the storage is made available.
- Exclusive IPs that were assigned to the zone are available for re-use.

Replicating Zones

Use this option to copy an existing zone so that you can provision a new zone on the same system efficiently. The process of cloning a zone is similar to the process of creating a zone because you can accept each of the original zone's specification or change it before you create the copied zone.

To Clone a Zone

1. Select the zone that you want to clone in the Assets section.
2. Click **Replicate Zone** in the Actions pane.
The Replicate Zone Wizard is displayed.
3. Enter a different zone name and description for the zone.
4. Select a library from the list of libraries that are associated with the global zone to store the cloned zone's image and metadata. Click **Next**.
5. The CPU shares that are allocated in the original zone are displayed. You can edit the changes for the cloned zone. Click **Next**.
6. Accept or change the attributes of the original zone: language, 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. When the cloned zone's file system is smaller than the original zone's file system, the clone operation cannot complete and the job fails. Do not modify the root file system of the new zone to a size less than the source zone root file system. Do not modify the system file system if it is defined in the source zone.

9. Accept or change the storage library. The library can be either a NAS storage or Fibre Channel library. When the library is NAS storage, specify the virtual disk name and size of the disk. For SAN library, select a LUN from the available list of LUNs in the library. The size of the selected LUN is displayed. You cannot change the size of the LUN.
10. The accumulated size of the storage is displayed as Currently Accumulated Storage. The required storage is displayed as Recommended Storage Size. Click the **Add** icon to configure more storage resources to the zone if the Currently Accumulated Storage is less than the Recommended Storage Size. When you have defined the Currently Accumulated Storage size as at least equal to the Recommended Storage size, click **Next**.
11. Assign at least one network to the zone. Select the zone IP type as Shared or Exclusive.
 - For Shared IP networks:
 - a. Select a network from the list of networks that use the Shared IP mode and are assigned to the global zone.
 - b. Select a NIC from the list of shared Network Interface Cards (NIC).
 - c. Specify the management interface for the NIC. When the network has a defined IP range, the Auto Allocate IP option is displayed with the zone's IP address from the range. When the network does not have a defined IP range, the Use Static IP option is displayed. Enter an IP address for the zone in the Zone IP field.
 - d. Click **Next**.
 - For Exclusive IP networks
 - a. Select a network from the list of networks that are assigned to the global zone and are not used by other zones.
 - b. Select a NIC from the list of the selected network's NICs that are not bound or assigned to other networks.
 - c. Specify the management interface for the NIC. When the selected network has a defined IP range, the Auto Allocate IP option is displayed in the Management Interface and the zone's IP address is populated with an IP address from the defined range.

When the network does not have a defined IP range, you must either provide the IP address or specify a DHCP server to provide one. To provide an IP address, select the **Static IP** option and enter the IP address in the Zone IP field. To designate a DHCP server, select the **Assigned by ext.**

DHCP option. The Zone IP field contains the Automatically Allocated value.

- Click **Next**.

12. Click **Finish** to launch the job for replicating a zone.

Adding Storage to Zones

You can add storage to zones dynamically. The storage is added to the zpool that is created for the zone. You cannot delete the storage from the zpool.

The storage libraries that are associated with the global zone are available for the zone. You can assign the following types of libraries to a zone:

- File system storage: Use NAS libraries
- Block storage: Use SAN or Dynamic storage libraries
- Local storage: Add the local storage on the global zone and local devices that are attached to the system

See [Chapter 6, "Storage Libraries"](#) and [Chapter 5, "Software Libraries"](#) for more information about setting up your storage servers and software libraries.

To Add Storage to Zones

1. Select the zone in the Assets section.
2. Click **Add Storage** in the Actions pane. As an alternative, click the **Storage** tab and click the **Add Storage to Zone** icon under Zpool and Storage.

The Add Storage to the Zone window is displayed.

3. Click the **Add** icon to add storage.

The storage libraries associated with the global zone are displayed in the Library list. The library can be a local storage, NAS, local device, Dynamic Storage, or SAN library.

4. Select a library from the list.
5. The selection of virtual disks varies according to the library selected.
 - Local and NAS library: Enter a virtual disk name and specify the size of the disk.
 - For SAN and iSCSI libraries: Select a LUN from the list. The size of the LUN displays automatically.
 - For Dynamic Storage Library: Either select a LUN from the list or create new LUN. When you create a new LUN, select the volume group from the list and specify the size of the new LUN.

The Additional Storage Specified shows the additional storage added to the zone.

6. Click **Add Storage** to add the specified storage to the zone.

Moving Zone Storage

Some management and monitoring functions of the zone might be disabled depending on the managed resources of the zone. For example, the zone migration option is disabled when the zone file system is not on a shared storage.

To enable the storage management and monitoring functions of the zone, you must move the zone's existing storage to a shared storage. This provides the following options to you to manage the storage and zpools:

- Convert the local storage to shared storage so that you can migrate zones.
- Convert the experimental NFS file systems into local or shared storage.
- Combine small ZFS pools per zone into a single ZFS pool.
- Split a single ZFS pool hosting zones into ZFS pools dedicated to each zone.

Oracle Enterprise Manager Ops Center provides the Move Storage option to convert the storage of the zones. You can use this option to manage your unmanaged storage resources of the zone.

Note: You cannot move managed storage to unmanaged storage and you cannot alter the root file system of the zone.

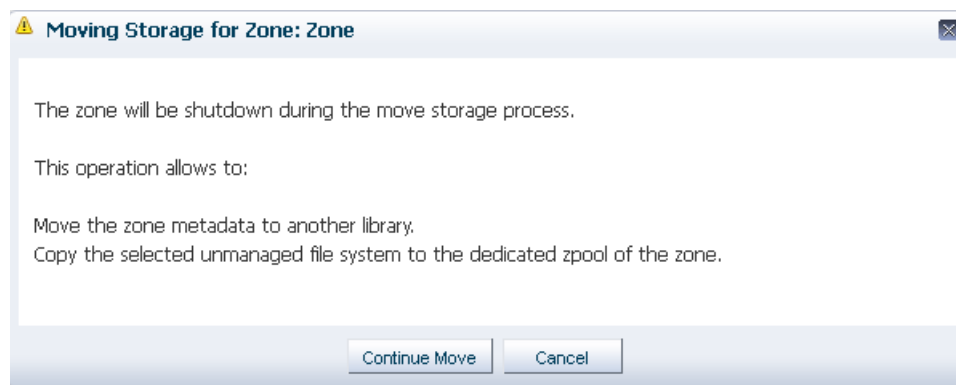
When you discover and manage existing zone environments in Oracle Enterprise Manager Ops Center, the storage is considered to be unmanaged. You can move this storage from unmanaged state to managed using this option.

To Move Zone Storage

1. Select the zone for which you want to move the storage.
2. Click **Move Storage** in the Actions pane.

A warning message is displayed that the zone is shutdown during this process. You can use this option to relocate the zone metadata to a shared storage and copy the selected unmanaged file system to the dedicated zpool of the zone.

Figure 14–11 Warning Message for Move Storage



3. Click **Continue Move** to proceed with the continue the process.
The Move Storage Wizard is displayed. The zone details are displayed.
4. You can edit the description and tags of the zone.
5. If required, you can change the library in which you have stored the zone metadata. All the libraries associated with the global zone are listed. Select the library from the list. Click **Next**.
6. Select the file systems that you want to be managed.

- When there are unmanaged file systems in the zone, you can select the option **Managed** and provide the **Reserved** and **Quota** size of the file system. These are added as new ZFS file systems on the existing zpool for the zone.
- When you want to relocate the root file system to a managed storage, a dedicated zpool is created for the zone and you must configure the virtual disks for the file systems.
- You can also modify the size of the file systems except for the root file system.
- You can keep some unmanaged file systems and relocate only those you require.
- You cannot relocate managed storage to unmanaged storage.

Click **Next** to configure the virtual disk storage when you move the root file system from unmanaged to managed state.

7. Select the library and the virtual disk for the zone's storage.

The libraries associated with the global zone are listed. Select **NAS**, **SAN** or **Dynamic Storage** library from the list.

For **NAS** library, provide a virtual disk name and enter the size of the disk.

For **SAN** library, select a LUN from the list.

For **Dynamic Storage** library, either select a LUN or create new LUN from the list. When you create a new LUN, select the volume group and enter the size of the LUN.

Click **Next** to view the summary.

8. Review the information and click **Finish** to change the zone storage from unmanaged to managed.

Adding File Systems to Zones

You can add file systems to zones. The zone must be in a shut down state to add file systems. The storage source for the file system can be managed or unmanaged storage source. For unmanaged storage source, you must provide the mount point of the storage.

Provide the **Reserved** and **Quota** size for the file system. Boot the zone for the changes to take effect. The file system is added to the existing zpool of the zone.

To Add a File System to a Zone

1. Select the zone in the **Assets** section.
2. Click **Add File Systems** in the **Actions** pane.
The **Add File Systems** window is displayed.
3. Click the **Add** icon to add file system.
4. Enter the file system.
5. Select whether the storage is managed or unmanaged.

When it is not managed, enter the mount point of the storage source.

6. Enter the **Reserved** and **Quota** size for the new file system.
7. Click **Add File Systems**.
8. Boot the zone for the changes to take effect.

Connect and Disconnect Networks

You can connect and disconnect networks from non-global zones. The networks attached to the corresponding global zone are available for the zones. The shared IP zones can connect to only networks that are specified for shared IP mode. The exclusive IP zones can connect to only exclusive IP networks.

Select the network and click the connect or disconnect icon as required.

Automatic Recovery

Use the options Enable Automatic Recovery and Disable Automatic Recovery to set the recovery option of created zones. To set automatic recovery for the zone, select the option Enable Automatic Recover. Edit the zone configuration to set the priority of recovery. The zone with highest value is recovered first. See [Recovering Zones](#) to manually recover the zones.

Migrating Zones

In Oracle Enterprise Manager Ops Center, zone migration is a cold migration because the zone is shut down, all applications are stopped, the migration occurs, and then the zone is restarted.

To enable migration action for a zone, the zone storage must be on a shared storage. When a zone uses local storage, use the [Moving Zone Storage](#) option to change the storage from local to shared.

This section describes the procedure for the following types of migration:

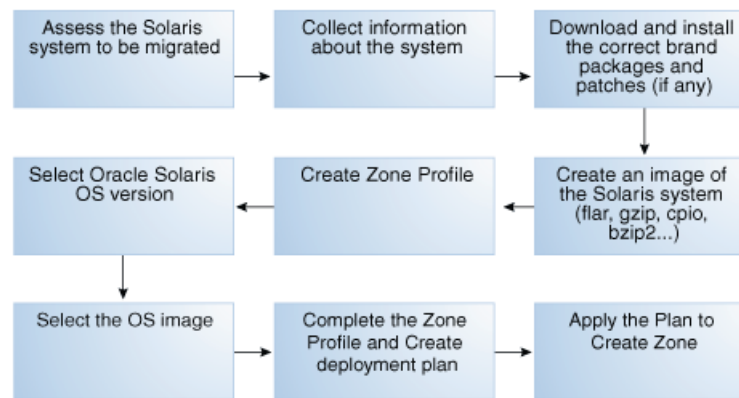
- Migrating an Oracle Solaris system into a new non-global zone. This is referred to as physical to virtual (P2V) conversion. See [Migrating a Physical Oracle Solaris System into a Zone](#).
- Migrating a non-global zone from one global zone to another global zone, when all components are managed by Oracle Enterprise Manager Ops Center. See [Migrating Zones to a Different Machine](#).
- Migrating a non-global zone that has dependencies that are not managed by Oracle Enterprise Manager Ops Center. This process uses a script to perform the migration. To migrate unmanaged file systems, see [Script to Migrate a Zone With Dependencies](#).

Note: When you migrate a zone, the Alternate Boot Environment (ABE) is not supported.

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 14-12](#) shows the steps for a P2V conversion.

Figure 14–12 Workflow for Migrating an Existing Oracle Solaris System Into a Zone

Ensure the following before migrating the zone:

- Assess the system to be migrated and collect information.
See <http://docs.oracle.com/cd/E19683-01/817-1592/> for information about collecting information about the source system.
- The system image to be installed in the non-global zone must not be newer than the target global zone's operating system release or the installation fails.
- The destination global zone must be running at least Oracle Solaris 10 8/07 OS.
- The supported branded zones are Oracle Solaris 8 and 9. For the branded zones to be supported on the destination global zone, remove the following brand packages from the global zone:
 - SUNWs8brandu
 - SUNWs8brandr
 - SUNWs9brandu
 - SUNWs9brandr

For instructions to download and install the correct brand packages, refer http://docs.oracle.com/cd/E22645_01/index.html

Migrating Zones to a Different Machine

Using Oracle Enterprise Manager Ops Center you can migrate one or more zones simultaneously. You can either migrate the zone to an individual global zone or to a zones server pool. When you migrate zones to a server pool, the target global zone depends on the server pool placement policy.

When you migrate a zone, perform a trial run before submitting the migration job to verify that the target global zone has the correct configuration to host the non-global zones.

The target global zone must have same or later versions of the following operating system packages and patches that are installed on the non-global zone.

- Packages that deliver files under an `inherit-pkg-dir` resource
- Packages where `SUNW_PKG_ALLZONES=true`

Other packages and patches, such as those for third-party products, can be different.

When the source and target global zone do not have the same patches and packages during migration, either update the zones patches and packages to match the target global zone or migrate without updating the zone patches and packages.

When the target global zone has later versions of the zone-dependent packages or patches, update those packages in the non-global zone before the migration to match the target global zone. When the target global zone supports the update on attach capability, it checks the non-global zone for packages that must be updated and only those packages are updated. The rest of the packages, and their associated patches, can vary from zone to zone.

You cannot downgrade the patches and packages of the zones to a lower version. In such cases, the update option fails.

The patches that must be backed out of the zone before the update are also listed. You must remove the patches manually and run the migration job again.

You can change the Name, Description, Tags, and NIC details of the non-global zone but, when you change other configuration, a warning message indicating that to change the zone configuration you must execute the `sys-unconfigure` command on the zone.

After you migrate a zone, you cannot use an alternate boot environment (ABE) to upgrade the zone.

Compatible Global Zones for Migration

When you click the Migrate Zone option in the Actions pane, Oracle Enterprise Manager Ops Center checks all global zones for compatibility with the source global zone. When there are compatible global zones, the Migrate Zone Wizard is displayed.

Otherwise, the following pop-up window is displayed.

Figure 14–13 Message Displayed When No Global Zones are Compatible

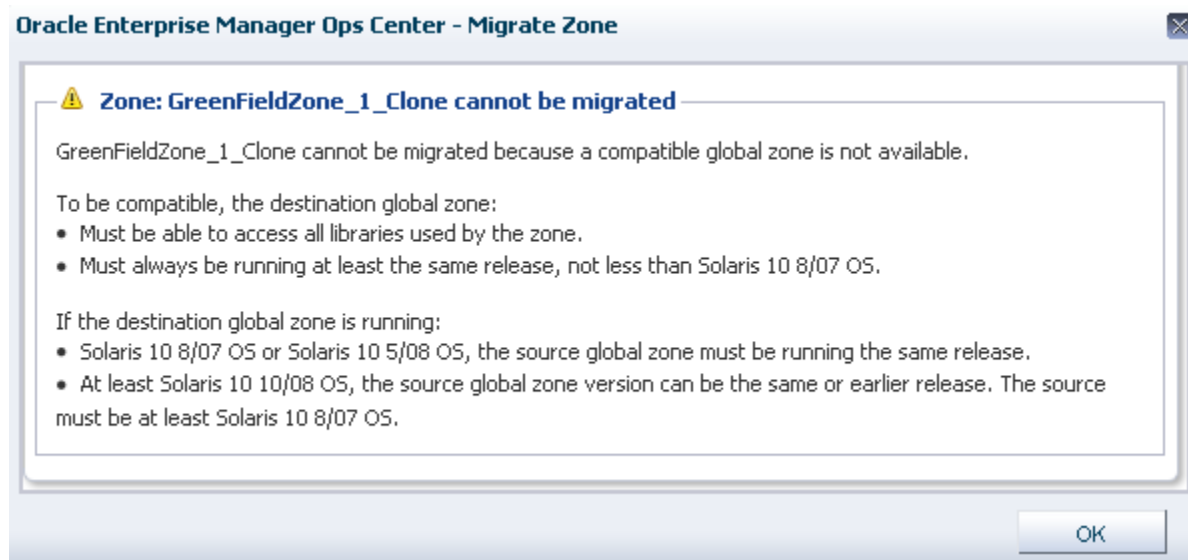


Table 14–4 shows the compatible global zones, depending on the operating system release.

Table 14–4 Compatible Global Zones

Source Global Zone	Compatible Target Global Zone
Oracle Solaris 10 8/07	Oracle Solaris 10 8/07
	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 5/08	Oracle Solaris 10 5/08
	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 10/08	Oracle Solaris 10 10/08
	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 5/09	Oracle Solaris 10 5/09
	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 10/09	Oracle Solaris 10 10/09
	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 9/10	Oracle Solaris 10 9/10
	Oracle Solaris 10 8/11
Oracle Solaris 10 8/11	Oracle Solaris 10 8/11
Oracle Solaris 11	Oracle Solaris 11
	Oracle Solaris 11 Update 1
Oracle Solaris 11 Update 1	Oracle Solaris 11 Update 1

Verify the following conditions:

- The source and target global zones are compatible.
- The source and target global zones have access to all the libraries associated with the non-global zone.
- The non-global zone's metadata is not stored in the source global zone's local library.
- The non-global zone's data is not stored in the source global zone's local library. The zone must use a shared storage library.
- The non-global zone is in the running state.

- For a server pool, it must have at least one compatible global zone for migration.
- When you migrate the zone, a warning message indicates that when you change the zone configuration except for Name, Description, Tags, and NIC details, you must execute the `sys-unconfigure` command on the zone. Also, when the NIC names are changed for a zone that uses exclusive IP mode, the `/etc/hostname.itf` and `/etc/dhcp.itf` file are renamed accordingly.

When the target global zone does not support backout on attach capability, you must remove or downgrade the patches and packages manually before continuing with the migration. Create an update profile that includes the patches that must be removed. Run an update job with this update profile. Repeat the migration job on the zone after the removal of the patches and packages.

To Migrate a Zone

1. Select the zone in the Assets section.

2. Click **Migrate Zone** in the Actions pane.

A warning message indicating that when you change the zone configuration except for Name, Description, Tags, and NIC details, you must execute the `sys-unconfigure` command on the zone. Also, when the NIC names are changed for a zone that uses exclusive IP mode, the `/etc/hostname.itf` and `/etc/dhcp.itf` file are renamed accordingly.

3. Click **Continue Migration**.

The Migrate Zone Wizard is displayed.

4. Select an individual global zone or server pool as the target.

The compatible global zones list the number of zones that are running, total CPUs, and available dedicated CPUs. The server pools list the average usage of CPU and memory.

Click **Next**. The zone migration test starts.

5. Review the migration test result. Select an update option to continue the migration on the target global zone or server pool:

- Update the patches and packages of the zone to match the target and then migrate the zone.
- Migrate the zone without updating its patches and packages.

Click **Save Test Result As** to save the migration test result. When you want to change the target, click **Previous** and select another target to run the migration test.

Click **Next** to review the zone identification.

6. If the zone name exists in the target global zone, you must change the zone name.

Click **Next** to specify the zone setup.

7. Select the language, time zone, and terminal type for the zone. The host name is the zone name as defined.

The dynamic value for NFSv4 domain name enables the domain name to be derived dynamically from the naming service configuration. To hard code the value for NFSv4 domain, provide a domain name.

Leave the password fields empty to use the existing password. Click **Next** to define the network interfaces.

8. You must have at least one network interface for the migration to continue. The network interfaces that are not accessible to the target global zone are displayed in yellow. Specify a new network interface for the inaccessible networks or click **Do Not Connect** to that network.
 - a. Select the network which is marked in yellow color. The network interface details are displayed under Network Interface.
 - b. Select a new network interface from the Network list or click **Do Not Connect**. The selected network can be either shared or dedicated. Enter the required network interface information for the selected network.

Click **Next** to specify the naming service.
9. Specify the naming service for the zone: DNS, NIS, NIS+ or LDAP naming service, or choose to not specify a naming service. Click **Next**.
 - **DNS:** Enter the domain name of the DNS server and the IP address of the DNS server. You can enter up to three IP addresses in the Name Server field. To specify additional domains to search, enter up to six domain names for the Domain Name Search List. The total length of each entry cannot exceed 250 characters.
 - **NIS and NIS+:** Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the Specify an NIS Server option to provide the NIS server host name and its IP address. When you do not have the NIS server information, click **Find an NIS Server**.
 - **LDAP:** Enter the domain name of the LDAP server. Specify the name of the LDAP Profile. Enter the IP address of the LDAP Profile Server. You can also provide the Proxy Bind Distinguished Name and Password.
 - **NONE:** Select this option when you do not want to configure a naming service.
10. Review the summary of the migration job. When there are no patches and packages to be backed out, continue with the migration. When there are patches or patches that must be removed, a warning is displayed.
 - When the target global zone supports backout on attach capability, a warning is displayed that the list of patches have been removed or downgraded.
 - When the target global zone does not support backout on attach, the patches and packages on the zone must be removed or downgraded manually so that the source zone matches the destination global zone. Cancel the migration or select a new target global zone.
11. Click **Finish** to submit the migration job.

To Migrate Multiple Zones

1. Select the global zone from which you want to migrate the zones.
2. Click **Migrate Zones** in the Actions pane.

The Migrate Zones Wizard appears. The list includes the zones running in the global zone.
3. Select one or more zones from the list. Click **Next**.
4. Select an individual global zone or server pool to be the destination for the zone migration.

The table displays the list of eligible global zones and server pool to which you can migrate the zones. The target global zone in the server pool depends on the server pool placement policy.

5. Select an update option to continue with migration.

The source and the target global zones might not be in the same patch level. Either select to update the patches and packages of zone to match the target global zone or continue migration without updating the zone.

6. Review the details and click **Finish** to migrate the zones.

Script to Migrate a Zone With Dependencies

You can migrate zones even when the zone has dependencies that are not managed by Oracle Enterprise Manager Ops Center, such as when the zone uses storage that is not part of the zpool.

Scripts enable you to extend the zone migration feature to include the migration of dependencies, such as storage or other resources that are not managed by Oracle Enterprise Manager Ops Center. For example, unmanaged file systems are lost when you migrate a zone. To avoid this, use a script to migrate the file systems. See [Example 14-1](#) for a sample script.

Develop your own scripts to migrate the dependencies and place them on the source and target global zone before migration.

The script for migrating the dependencies is executed in the following way:

- The migration job checks for a script placed on both the source and target global zones. The migration job is aborted when the script is found only on one of the global zones.
- When you create zones server pool, ensure that you upload the script to handle the unmanaged storage. The script is placed in all the global zones in the server pool and thus the migration of zone dependencies are taken care.
- The job checks whether the script has zero on exit. A non-zero exit is a failure and the migration job fails.
- The script is called on the global zone eight times when the migration job is executed.

Table 14-5 Script Call to Environmental Variables

Call to the Script	Environmental Variables
On the source global zone, to verify the script existence and user dependencies	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

Table 14–5 (Cont.) Script Call to Environmental Variables

Call to the Script	Environmental Variables
On the source global zone, after the zone is shutdown.	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION
On the source global zone, after the zone is detached	OEMOC_ZONENAME=source zonename OEMOC_PHASE=PREOPERATION_SHUTDOWN_ DETACHED OEMOC_OPERATION=MIGRATION
On the target global zone, before the zone is attached	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=POSTOPERATION_SHUTDOWN_ DETACHED OEMOC_OPERATION=MIGRATION
On the target global zone, before the zone is started.	OEMOC_ZONENAME=destination zonename OEMOC_PHASE=POSTOPERATION_NOTRUNNING OEMOC_OPERATION=MIGRATION
On the target global zone, after the zone is started	OEMOC_ZONENAME=destination 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 POSTROLLBACK
- POSTOPERATION becomes PREROLLBACK

Script Requirements

The script must be executable and follow these conventions:

- Good error checking and clean-up within the script
- Standard exit code conventions
- Non-zero exit indicates that an error has occurred
- Informational messages are in stdout
- Error messages are in stderr
- The script must be named as guest-operations and placed in the `/var/opt/sun/oc/public` directory on both the source and target global zone.

See the [Sample Script](#) for migrating the additional inherited file systems of an adopted zone. The sample script is based on the assumption that `/opt/` file system is the additional inherited file system. The script re-configures `/opt/` file system on the target global zone before rebooting the zone. The script reconfigures the user-inherited file systems during POSTOPERATION_NOTRUNNING phase

Example 14–1 Sample Script

```
#!/bin/sh

do_migration_action() {
```

```
echo "executing migration action for zone ${OEMOC_ZONENAME}"

if [ ${OEMOC_PHASE} != "POSTOPERATION_NOTRUNNING" ]; then
    exit 0
fi

#
# for migration action, before starting the zone on the target GZ
# add to /opt inherited filesystem.

/usr/sbin/zoneadm -z ${OEMOC_ZONENAME} detach >/dev/null 2>&1
if [ $? -ne 0 ] ;then
    echo "unable to detach ${OEMOC_ZONENAME}"
    exit 1
fi

echo "add inherit-pkg-dir" >/tmp/tmp-$$$.txt
echo "set dir=/opt" >>/tmp/tmp-$$$.txt
echo "end" >>/tmp/tmp-$$$.txt

/usr/sbin/zonecfg -z ${OEMOC_ZONENAME} -f /tmp/tmp-$$$.txt >/dev/null 2>&1
if [ $? -ne 0 ] ;then
    echo "unable to add inherit filesystem for ${OEMOC_ZONENAME}"
    exit 1
fi

/usr/sbin/zoneadm -z ${OEMOC_ZONENAME} attach -u >/dev/null 2>&1
if [ $? -ne 0 ] ;then
    echo "unable to attach ${OEMOC_ZONENAME}"
    exit 1
fi
}

echo ${OEMOC_ZONENAME}
echo ${OEMOC_PHASE}
echo ${OEMOC_OPERATION}

if [ ${OEMOC_OPERATION} = "MIGRATION" ] ;then
    do_migration_action
fi
exit 0
```

Recovering Zones

When the global zone crashes or must be halted, its non-global zones can be migrated to another global zone. However, this zone recovery procedure is not the same as zone migration because Oracle Enterprise Manager Ops Center cannot get access to the information in the global zone and therefore cannot perform compatibility checks. The procedure to relocate non-global zones from the source global zone to the target global zone is a forced attachment of the non-global zone.

The zone recovery procedure uses the command-line interface for the Oracle Enterprise Manager Ops Center.

See the [Chapter 19, "Command Line Interface"](#) 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 the same libraries associated with the non-global zone.
- Verify that the non-global zone's metadata and operational data is not stored in a local library.
- Verify that JDK version 6 is used. In the Enterprise Controller command prompt, enter the following command:

```
export JAVA_HOME=/usr/jdk/latest
```

To Recover Zones

1. Connect to the Oracle Enterprise Manager Ops Center CLI using the following command:

```
/opt/SUNWoccli/bin/oc
```

2. Connect to the local Enterprise Controller.

```
xvmSh > connect
```

```
localhost >
```

3. Enter the virtualization mode.

```
localhost > virtualization
localhost [virtualization] >
```

4. List the available global zones.

```
localhost [virtualization] >list_hosts
```

Name ObjectName	Type	Health	Reachable
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, when the source global zone is gzhost36, then enter the following command.

Note: Do not use the user-friendly name of the zone in the commands.

```
localhost [virtualization] >list_hosts -C
com.sun.hss.domain:type=xVMServer,name=NORM-NORM-localhost
```

Name ObjectName	Type	State	Migratable
--------------------	------	-------	------------

```

-----
test                ZONE                RUNNING                True
com.sun.hss.domain:type=Server,name=NORM-07e91405-8313-43ec-9671-dc320989866e
6. Select the destination global zone and start the zone on it. For example, when the
   destination global zone is gzhost44, execute then execute the following command
   to start the test zone on it.

localhost [virtualization] >startup -Z <source global zone> -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.

Zones Server Pool

You can create server pools for zones in Oracle Enterprise Manager Ops Center. Pooling your virtualization hosts provides the capability for load balancing the virtualization servers, high availability and minimize power consumption.

See [Chapter 17, "Server Pools"](#) for information about creating and managing zones server pools.

Updating Zones

Oracle Enterprise Manager Ops Center enables you to update the global and non-global zones. You can also patch zones that are running on a supported configuration. The installation of the patches on the zones depend on the package parameters and the attribute set for the patch commands. This section describes the parameters for installation of the packages and patches. The concepts involved in updating global and non-global zones, and the procedures to update the zones are described in this section.

Note: When you use Oracle Solaris Live Upgrade to update the OS in a zone, you cannot use Oracle Enterprise Manager Ops Center to manage the zone. Alternate Boot Environment (ABE) is not supported for zones that were created using Oracle Enterprise Manager Ops Center.

Install Packages and Patches on Zones

A patch is a collection of files and directories that replace existing files and directories that are preventing proper execution of the software.

You can install packages and patches on a zone. The `patchadd` and `pkgadd` commands operate in the background to install a patch and package respectively. However, the installation of packages on zones also depends on the parameters `SUNW_PKG_`

ALLZONES, SUNW_PKG_HOLLOW, and SUNW_PKG_THISZONE. These parameters control whether a package can be installed on global zones or non-global zones. The actions for the parameters are as follows:

- **SUNW_PKG_ALLZONES:** If the value is true, the package is installed on all zones, both global and non-global.
- **SUNW_PKG_HOLLOW:** If the value is true, the package information is propagated to the non-global zones, but the package is not installed.
- **SUNW_PKG_THISZONE:** If the value is true, the package is installed only in that zone.

Configure patchadd and pkgadd Commands

In Oracle Enterprise Manager Ops Center, the patchadd, pkgadd, patchrm, and pkgrm commands are implemented without the -G switch by default. To install updates or packages only on the current zone, enable the -G switch by editing the .uce.rc file.

Note: Two files in the /SUNWuce/agent/bin directory have similar names. Verify that you are editing the .uce.rc file. Do *not* edit the uce.rc file.

Editing the .uce.rc File

1. Open the .uce.rc file in the /SUNWuce/agent/bin directory in the managed system.
2. Add the following lines to the .uce.rc file:


```
( all ) (invisible.__is_patchadd_g_specified, false)
( all ) (invisible.__is_patchremove_g_specified, false)
( all ) ( invisible.__is_pkgadd_g_specified, false)
```
3. Set the -G parameter to true for the action that you want to perform.
4. Save and close the file.
5. For this change to take effect, restart the services using the following commands:


```
svcadm disable -s update-agent
svcadm enable -s update-agent
```

Updating a Global Zone

In Oracle Enterprise Manager Ops Center, when a package or patch is installed, the patchadd and pkgadd commands are implemented in the background as shown in the following example:

```
patchadd <patchid>
pkgadd <pkgname>
```

Change the way that these commands are implemented by enabling the -G switch. You can enable the -G switch to cause the patch or package to be installed to the target zone only if the package parameter SUNW_PKG_THISZONE is set to true. See [Install Packages and Patches on Zones](#) for information about configuring the patchadd and pkgadd commands on the managed systems.

See the following scenarios when you are updating a global zone. The result for each scenario determines whether the update job is successful, depending on the package information.

Table 14–6 *Updating a Global Zone Scenarios*

SUNW_PKG_ALLZONES	SUNW_PKG_THISZONE	SUNW_PKG_HOLLOW	Impact	Impact with -G Configuration
False	False	False	The package is installed on the global zone, and all the non-global zones	The package is installed only on the global zone.
True	False	False	The package is installed on the global zone and all the non-global zones.	The -G switch cannot override the SUNW_PKG_ALLZONES parameter, and the package is installed on all the zones.
True	False	True	The package is installed on the global zone and the package information is made available on all the non-global zones.	The -G switch cannot override the SUNW_PKG_ALLZONES parameter, and the package is installed on all the zones.
False	True	False	The package is installed only on the global zone.	The package is installed only on the global zone.

Patches are sets of updates to packages. When you install a patch, the patch is installed on the global zone and the non-global zones, depending on the package parameters as shown in the previous table.

Note: Use caution while enabling the -G option on a host with sparse zones. Packages that are inherited from the global zone that are not SUNW_ALL_ZONES cannot be patched within a sparse zone.

Updating a Non-Global Zones

A user with the virtualization administrator role can install packages and patches on non-global zones. The `patchadd` and `pkgadd` command must be used without options. Do not configure the -G switch to the commands while updating the non-global zones.

See the following scenarios when you are updating a non-global zone. The results of each scenario determine whether the update job is successful, depending on the package information.

Note: The -G switch does not have an effect on installing packages or patches in a non-global zone.

Table 14–7 *Updating Non-Global Zones Scenarios*

SUNW_PKG_ALLZONES	SUNW_PKG_THISZONE	SUNW_PKG_HOLLOW	Impact
False	False	False	The package is installed only on the target non-global zone.

Table 14–7 (Cont.) Updating Non-Global Zones Scenarios

SUNW_PKG_ALLZONES	SUNW_PKG_THISZONE	SUNW_PKG_HOLLOW	Impact
True	False	False	The package installation fails.
True	False	True	The package installation fails.
False	True	False	The package is installed only on the target non-global zone.

Note: When the patch is installed only on the non-global zone, ensure that autoboot property is set to true for the zone. Otherwise, single user mode patches fail to apply as the zone does not come up after the reboot.

Patches are sets of packages that must be installed. When one of the packages has the SUNW_PKG_ALLZONES parameter set to true, then the patch installation fails. For a successful patch installation, ensure that none of the packages have SUNW_PKG_ALLZONES parameter set to true.

Note: Packages that deliver to read-only inherit directories do not install on sparse root zones. These packages must be installed from the global zone with the -G switch disabled. When a package has the parameter SUNW_PKG_THISZONE=true, it does not appear as installed from the sparse zone and the software might not function correctly. In this case, a whole root zone must be used. Packages with SUNW_PKG_THISZONE=true must not deliver to read-only inherit directories.

Zone Parallel Patching

To view the number of zones that you can patch in parallel when you update the global zone, go to the global zone's Summary page. When the agent is installed on the Oracle Solaris OS, the number of zones that you can patch in parallel is calculated as 1.5 times the number of CPU cores in the server. For example, if you have a 2 CPU core machine, then you can patch three (3) zones in parallel. This is set and displayed on the UI.

Oracle Solaris 10 and 11 OS version displays this information in the Summary page. When the update capability is not enabled on the OS, then the zone parallel patching is also not enabled.

Related Zone Operations

Refer to the following chapters in this guide for operations that are common for many resources:

- See [Chapter 6, "Storage Libraries"](#) for setting up your storage resources.
- See [Chapter 7, "Networks"](#) for setting up your network infrastructure.
- See [Chapter 8, "Plans and Profiles"](#) for managing your zone profiles and deployment plans.

- See [Chapter 12, "Operating Systems"](#) for monitoring your resources and setting up boot environments.
- See [Chapter 2, "Asset Management"](#) for discovering and registering your assets.
- See [Chapter 17, "Server Pools"](#) for creating and managing zones server pools.

Related Resources

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

- See [Chapter 17, "Server Pools"](#) for information about creating and managing zones server pools.
- See [Chapter 7, "Networks"](#) for information about setting up your network infrastructure.
- See [Chapter 6, "Storage Libraries"](#) for information about setting up your storage resources.
- For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm. See the following example workflows for deploying and using zones:
 - Deploy Oracle Solaris 10 Zones Workflow
 - Deploy Oracle Solaris 11 Zones Workflow
 - Operate Zones Workflow
- See the *Oracle Enterprise Manager Ops Center Administration Guide* for information about how to manage the software, including user and role management.
- See <http://docs.oracle.com/cd/E19044-01/sol.containers/817-1592/> for Oracle Solaris Resource Management and Oracle Solaris Zones documentation.
- The complete Oracle Solaris 10 documentation set is located at <http://docs.oracle.com/cd/E19253-01/index.html>.
- See http://docs.oracle.com/cd/E23824_01/html/E24456/transzone-1.html for information about creating Solaris Flash archive images of an Oracle Solaris 10 operating system.
- See <http://docs.oracle.com/cd/E19082-01/819-6990/index.html> for information about network interfaces and virtualization, and administration of your network interfaces.

Oracle VM Server for SPARC

The following information includes:

- [Introduction to Oracle VM Server for SPARC](#)
- [Roles for Managing Oracle VM Server for SPARC](#)
- [Available Management Features](#)
- [Prerequisites for Provisioning Oracle VM Server for SPARC](#)
- [Recommended Minimum Configuration](#)
- [Creating an OS Provisioning Profile](#)
- [Installing Oracle VM Server for SPARC](#)
- [Oracle VM Server for SPARC Management](#)
- [Server Pools](#)
- [About Logical Domains](#)
- [Creating a Logical Domain Profile](#)
- [Creating and Deploying a Logical Domain Plan](#)
- [Management of Logical Domains](#)
- [Manage Existing Oracle VM Server for SPARC Environments](#)
- [Related Resources](#)

Introduction to Oracle VM Server for SPARC

Oracle VM Server for SPARC technology enables server virtualization on SPARC platforms. You can create and manage multiple virtual machine instances simultaneously on a single SPARC machine. Each virtual machine, or guest, can run a different operating system.

Oracle VM Server for SPARC technology is virtualization of SPARC servers. This technology is part of a suite of methodologies for consolidation and resource management for SPARC Chip Multi Threading (CMT) systems. Using this technology, you can allocate the various resources of the system such as memory, CPU threads, and devices, into logical groupings and create multiple discrete systems. These discrete systems have their own operating system, resources, and identity within a single system. By careful architecture, an Oracle VM Server for SPARC environment can help you achieve greater resource usage, better scaling, and increased security and isolation.

When Oracle VM Server for SPARC software is installed, a domain called the control domain is created. From this control domain, you create virtual machines called logical domains that each run an independent OS. A logical domain is a virtual machine with resources, such as 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.

For more information about Oracle VM Server for SPARC, refer to <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>.

Roles for Managing Oracle VM Server for SPARC

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

Table 15–1 Oracle VM Server for SPARC Roles and Permissions

Task	Role
Provision and manage virtualization host	Virtualization admin
Create, manage, update, and delete guests	Virtualization admin
Discover and manage virtualization hosts	Asset admin
Create and manage profiles and plans	Profile and plan admin
Create and manage IPMP groups	Network admin
Create and manage Link Aggregation	Network admin
Set monitor threshold	Asset admin
Create credentials	Security admin

Available Management Features

Using Oracle Enterprise Manager Ops Center, you can perform the following tasks:

- Provision Oracle VM Server for SPARC
- Manage Oracle VM Server, including, rebooting, shutting down, changing name servers and NFS4 domains.
- Monitor the performance of Oracle VM Server
- Create and provision logical domains
- Manage logical domains, including editing, migrating, starting, rebooting, and shutting down domains
- Monitor the performance of logical domains
- Discover and manage existing Oracle VM Server for SPARC systems and logical domains
- Create server pools to maximize capacity and for automatic load balancing
- Migrate logical domains

Prerequisites for Provisioning Oracle VM Server for SPARC

You can provision the following versions of Oracle VM Server for SPARC using Oracle Enterprise Manager Ops Center:

- 1.2
- 1.3
- 2.0
- 2.1
- 2.2
- 3.0

The specific hardware, OS, and firmware requirements that must be met for provisioning Oracle VM Server for SPARC are described in this section.

Operating System

You require Oracle Solaris 10 9/10 OS or later release to provision 1.2, 1.3, 2.0, and 2.1 version.

For 2.2 version, you require Oracle Solaris 10 8/11 OS or Oracle Solaris 11 OS. The minimum version of Oracle Solaris 11 OS for all domain type is Oracle Solaris 11 Support Repository Update 8 (SRU 8).

For 3.0 version, you require Oracle Solaris 10 8/11 OS or Oracle Solaris 11.1 OS. The minimum version of Oracle Solaris 11.1 OS for all domain types is Oracle Solaris 11.1 Support Repository Update 1 (SRU 1).

For Oracle Solaris 10 8/11 OS or later release, when you select to deploy 3.0 version, the installed version on the server is Oracle VM Server for SPARC 3.0.0.2 version. To get all the functionality of 3.0.0.2 version, use Oracle Solaris 10 1/13 OS.

When the Oracle Solaris 11 Software Update Library is synchronized with the parent repository for Oracle Solaris 11.1 Support Repository Update 3.5 (SRU 3.5), the list of Oracle VM Server for SPARC versions is populated with 3.0.0.2.

SPARC T5 Server

You must use Oracle Solaris 10 1/13 OS for provisioning Oracle VM Server for SPARC on a SPARC T5 server. For all other SPARC servers, you can use Oracle Solaris 10 8/11 OS.

Refer to the Oracle VM Server for SPARC 3.0 Release Notes at http://docs.oracle.com/cd/E37707_01/html/E29668/index.html for more information about system requirements.

Hardware and Firmware

Refer to the Oracle VM Server for SPARC Release Notes in <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html> for the supported platforms to install different versions of Oracle VM Server for SPARC.

Ensure that you install the correct system firmware. You can update the firmware using Oracle Enterprise Manager Ops Center. See [Chapter 11, "Hardware"](#) for downloading and updating the required firmware version.

When you use Oracle Enterprise Manager Ops Center to provision the Oracle VM Server for SPARC software, it installs the Oracle Solaris OS, the Control Domain, and

an agent on the target system. This provisioning action removes existing virtualization software, including any previous logical domains installed on the service processor.

To install Oracle VM Server for SPARC software, apply a profile that specifies the values for resources such as CPU threads, crypto units, and memory as described in [Creating an OS Provisioning Profile](#).

When you have existing Oracle VM Server for SPARC environment that you want to discover and manage in Oracle Enterprise Manager, refer to the section [Manage Existing Oracle VM Server for SPARC Environments](#) for more information.

Recommended Minimum Configuration

The recommended minimum configurations for the Control Domain are described in the following sections.

CPU 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 contain at least one thread from the parent core. Crypto units cannot be split as CPU threads are split. For example, you have assigned the crypto unit for the first CPU core to the control domain. When a new logical domain is assigned a thread from the first CPU core and the crypto unit for that core is already assigned, the Control Domain cannot assign that crypto unit to the new logical domain. Allocation of crypto units might not succeed, especially when a core is split between domains. An Oracle VM Server might allocate fewer crypto units or none at all.

You must assign at least one crypto unit to the Control Domain because the crypto unit enables domain migration.

The use of Crypto Units is not mandatory although it might speed the logical domain migration. Allocation of Crypto Units might not be available in all the hardware supported. See the corresponding hardware data sheet and documentation for more information before planning for Crypto Units.

Example 15–1 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 16 GB, set the Control Domain's RAM to 8 GB.

Note: See Oracle VM Server for SPARC 2.1 Administration Guide at http://docs.oracle.com/cd/E23120_01/html/821-2854/index.html for more information about allocating memory size.

Creating an OS Provisioning Profile

Create an OS provisioning profile to install the Oracle VM Server for SPARC software, the OS, and to set the parameters for the Control Domain. The profile collects details such as type of target, OS image, time zone and language setup, required JET modules, disk partitions, naming services and network details.

Automatic booting on a SPARC system uses the default boot device that is defined in the non-volatile RAM (NVRAM). User-defined commands that are executed during start-up are stored in the NVRAMRC file in the NVRAM. When you run an OS provisioning job on a SPARC machine, Oracle Enterprise Manager Ops Center resets the configuration to the factory default configuration and removes the information that is stored in the NVRAMRC file. The Control Domain OS provisioning profile gives you the option to preserve the information in the NVRAMRC file before resetting the server to the factory defaults, and then restore the information after the reset.

Upload the required Oracle Solaris OS image to the software library in Oracle Enterprise Manager Ops Center. See [Chapter 5, "Software Libraries"](#) for more information about uploading OS images.

Oracle VM Server for SPARC Parameters

You can select the following parameters while creating the profile:

- Oracle VM Server for SPARC version: 1.2, 1.3, 2.0, 2.1, 2.2 and 3.0.
- Software groups
- CPU Threads
- Memory
- Crypto Units
- Virtual Console Port Range
- Enable Multiplexed I/O

Enabling multiplexed I/O results in enabling the Fibre Channel ports on the storage system.

- Name service
- NFS Domain value
- Manual netboot option

To Create an Oracle VM Server for SPARC Provisioning Profile

1. Select the Plan Management section and expand Profiles and Policies.
2. Choose OS Provisioning and click **Create Profile** in the Actions pane.
3. Specify the profile as shown in [Figure 15-1](#).
 - a. Enter the name and description of the profile.
 - b. Select the option **Create a deployment plan for this profile**.
 - c. Select Oracle VM Server for SPARC as the Subtype.

Figure 15-1 Profile Identification

Identify Profile * Indicates Required

* **Name:** myldom21

Description:

☒ Create a deployment plan for this profile.

* **Subtype:** Subtype
 Oracle VM Server for SPARC
 Oracle Linux
 Oracle VM Server for x86
 Red Hat Linux
 SUSE Linux
 JET Template
 Solaris SPARC
 Solaris x86

Target Type: Target Type
 V class (sun4v)

Click **Next** to select the OS image and distribution.

4. Select the following parameters in the Specify OSP Parameters step:
 - Oracle VM Server for SPARC Version
 - Software Group

The supported Oracle VM Server for SPARC versions to the available Oracle Solaris OS images are listed in [Table 15-2](#):

Table 15-2 Supported Oracle VM Server for SPARC Versions

Oracle VM Server for SPARC Version	Oracle Solaris 10 OS	Oracle Solaris 11 OS
1.0, 1.2, 1.3, 2.0 and 2.1	Oracle Solaris 10 9/10 OS or later release	N/A

Table 15–2 (Cont.) Supported Oracle VM Server for SPARC Versions

Oracle VM Server for SPARC Version		
	Oracle Solaris 10 OS	Oracle Solaris 11 OS
2.2	Oracle Solaris 10 8/11 OS or later release	Oracle Solaris 11 OS Oracle Solaris 11 Software Update Library synchronized with Oracle Solaris 11 Support Repository Update 8 (SRU8)
3.0	Oracle Solaris 10 8/11 OS or later release Oracle VM Server for SPARC 3.0.0.2 version is provisioned when Oracle Solaris 10 8/11 or later release of OS is used. To get all the functionality of 3.0.0.2 version, use Oracle Solaris 10 1/13 OS.	Oracle Solaris 11.1 OS Oracle Solaris 11 Software Update Library synchronized with Oracle Solaris 11.1 Support Repository Update 1 (SRU 1) When Oracle Solaris 11 Software Update Library is synchronized with Oracle Solaris 11.1 Support Repository Update 3.5 (SRU 3.5), Oracle VM Server for SPARC version list is populated with 3.0.0.2 version. Select 3.0.0.2 version from the list to provision.

Click **Next** to specify the OS setup.

5. Specify the following OS setup parameters:

- Select a language from the list.
- Specify the time zone.
- Specify a terminal type.
- To monitor the installation using a serial connection, select the console serial port device and the baud rate.
- Enter the NFS4 domain name for the target system to use. A dynamic NFSv4 domain name enables the name to be derived at run time, based on the naming service configuration. You can also enter a static domain name.
- Enter the root password for the root user on systems provisioned using this profile. Re-enter the password for confirmation.
- Select the Manual Net Boot option to enable manual control of network boot operations for the target system. You must select this option for a target system that does not have a service processor because Oracle Enterprise Manager Ops Center cannot control the network boot process remotely on these systems.

6. Click **Next**.

If it is Oracle Solaris 11 OS, then you are directed to Step 7 to create a user account to log in.

Otherwise, to Step 8 for Oracle Solaris 10 OS.

7. For Oracle Solaris 11 OS, root login to the system is not enabled. Create a user account through which you can SSH to the OS after provisioning. Enter the username and password for the user account.

Click **Next** to go to Step 9 to specify the Control Domain parameters.

8. You can use JET modules 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.Click **Next** to specify the Oracle VM Server Control Domain parameters.
9. Specify the resources that you want to assign to the Control Domain. The remaining resources are available for the logical domains.
 - **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 Multiplexed I/O (MPxIO):** Select this option to enable SAN library connectivity for the control domain. This action enables the Fibre Channel ports on the storage system.
 - **Save NVRAMC values:** Select this option to preserve your user-defined commands that are executed during start-up.

Note: For Oracle VM Server for SPARC 1.2 version, you can select to install SUNWJass package to harden the system.

Click **Next** to specify the file system layout.

10. 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: `ufs`, `unnamed`, `zfs`, or `swap`.
 - **Mount Point:** Enter a directory to use as a mount point for partitions.
 - **Device:** Enter the `rootdisk` keyword and a slice value to describe a partition on the target system's boot disk, for example, `rootdisk.s0`, or enter the logical device name, for example, `c1t0d0s0`, of the partition that you want to create.
 - **Size (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** to specify the name service.

11. Specify the name service, domain name and the corresponding name server. Select one of the following name services:

- **DNS:** Enter the domain name of the DNS server and enter the IP address of the DNS server in the Name Server field. You can enter up to three IP addresses as the value for the Name Server. Provide the additional domains to search for name service information in the Domain Name Search List. You can specify up to six domain names to search. The maximum length of each search entry is 250 characters.
- **NIS or NIS+:** Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, 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** to specify the network information.

12. Select one of the network options for the target system:

- When you select Use Link Aggregation, refer to the section [Create Link Aggregation in Oracle VM Server for SPARC](#) to define the parameters for Link Aggregation.
- When you select None, you are forwarded to the next step to define the network interfaces and IP address allocation for the selected networks.

Note: You can create Link Aggregation only with Oracle Solaris 10 OS and it is not applicable for Oracle Solaris 11 OS.

Select an option and click **Next** to continue.

13. 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.

Select the Address Allocation Method for each network except the boot interface.

- When 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.
- When you select Assign By DHCP, the IP address is automatically assigned from the designated DHCP server.

Click **Next** to go to the Summary step.

14. Review the Summary of your selections. Click **Finish** to create the profile.

Create Link Aggregation in Oracle VM Server for SPARC

While creating the profile for installing Oracle VM Server for SPARC, follow the procedure to create link aggregation:

1. Select **Link Aggregation** in the Specify Networks step of the profile.
2. Specify and configure the Link Aggregation details:
 - **Link Aggregation Name:** The name of the Link Aggregation is set to "aggr". Add an integer to make the name unique.

- **Network:** Select a network from the list.
- **NICs:** List the physical interfaces of the selected network that must be configured as a single unit. For example, GB_1, and GB_2.

Click **Next** to configure the IEEE 802.3ad Link Aggregation.

3. Define the following parameters for IEEE 802.3ad link aggregation standard:
 - **Load Balancing Policy:** Define the policy for outgoing traffic.
 - **Aggregation Mode and Switches:** When the aggregation topology connects through a switch, determine whether the switch supports the Link Aggregation Control Protocol (LACP). When the switch supports LACP, you must configure LACP for the switch and the aggregation. Define one of the modes in which LACP must operate.
 - **MAC Address Policy:** Define whether the MAC address of the NICs are fixed.Click **Next** to continue.
4. Continue to complete the profile.

Create and Apply Deployment Plan for Oracle VM Server for SPARC

When you create a profile, you can select to create a deployment plan with the profile automatically. You can also create your own deployment plan. Oracle Enterprise Manager Ops Center provides two different types of deployment plan to provision Oracle VM Server for SPARC:

- **Single Step Deployment Plan**

You have the Provision OS deployment plan to provision Oracle VM Server for SPARC. This requires the profile to install Oracle VM Server for SPARC.

- **Multi-Step Deployment Plan**

- **Install Server**

This multi-step deployment comprises the steps to provision Oracle VM Server for SPARC, install updates or software, apply monitoring profiles, and create logical domains.

- **Update Firmware and Install Oracle VM Server for SPARC**

This multi-step deployment plan comprises the steps to update the firmware of the hardware and then install the required version of Oracle VM Server for SPARC.

Select the required method and create the plan. All the plans described here require the OS provisioning profile for Oracle VM Server for SPARC. See [Chapter 8, "Plans and Profiles"](#) for more information about creating and managing the profiles and plans.

Apply a Deployment Plan for Oracle VM Server for SPARC

When you apply a deployment plan to provision Oracle VM Server for SPARC, you must provide the IP address for the OS to be provisioned. As shown in [Figure 15–2](#), you can also select the Oracle VM Server for SPARC to be placed in a server pool after provisioning. You have the option to create a new server pool or place the server in an existing server pool.

Figure 15–2 Placing Oracle VM Server for SPARC in Server Pool

Server Pool * Indicates Required Field

The server that will be installed can be assigned to an Oracle VM Server SPARC Server Pool. Select an assignment choice:

☒ Do not assign to a Server Pool. The new server will execute in stand-alone mode. You may add the server to a pool at a later time.

☐ Assign to one of these compatible Server Pools:

Server Pool ▲	Current Members
MyVPool	xvmt5120-2,xvmt5120-1

☐ Create a new Server Pool based on the attributes of the new server and assign the server, using default pool settings. The pool settings can be changed later after it has been created.

* Server Pool Name:

Storage Library ▲	Type	Description
LDomNAS	NAS	nfs://10.166.74.193/xVMShares/LDomNAS
MyFCLib	SAN	fc://cdab12aa-181c-4a1f-8ebd-281c
MyVNASZoneLib	NAS	nfs://10.166.88.16/xVMShares/ZoneLib

Read through the following sections to provision Oracle VM Server for SPARC in your environment.

Installing Oracle VM Server for SPARC

You can install the software on bare-metal or on systems already configured with logical domains. When you provision Oracle VM Server through Oracle Enterprise Manager Ops Center, any previous configuration is removed and the service processor is reset to its factory defaults.

When you provision Oracle VM Server for SPARC, the instance of the Oracle Solaris OS becomes the Control Domain which is the first domain to be created. The virtualization host on which this instance of Oracle Solaris OS runs is called the Oracle VM Server Host or simply, Oracle VM Server. The Logical Domains Manager runs in the Control Domain 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.

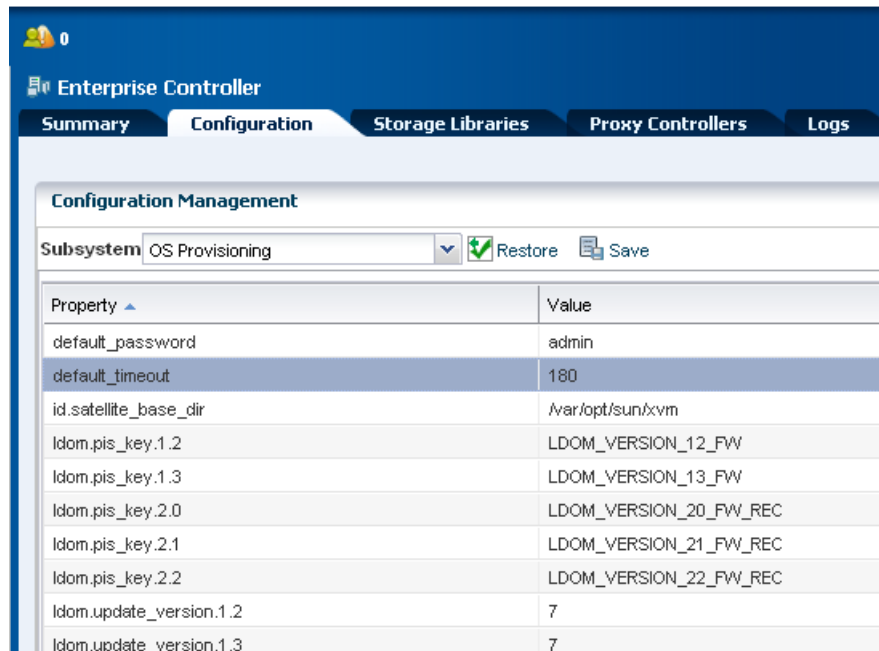
To Install Oracle VM Server for SPARC

This section describes how to install the Oracle VM Server for SPARC software on a service processor:

1. The Oracle VM Server for SPARC software depends on particular Oracle Solaris OS versions, required software patches, and particular versions of system firmware. See [Selection of Oracle VM Server Version](#).
2. Discover the target system on which you want to install Oracle VM Server for SPARC software.
3. Configure DHCP services on Proxy Controllers.
4. Create a profile for provisioning Oracle VM Server for SPARC. See [Creating an OS Provisioning Profile](#).

5. Create a deployment plan with the profile created for provisioning Oracle VM Server for SPARC.
6. (Optional) Adjust the amount of time allowed for the provisioning job. The default time is three hours or 180 minutes. You can edit the `default_timeout` value for OS provisioning in the Enterprise Controller Configuration in the Administration section, as shown in [Figure 15–3](#).

Figure 15–3 Default Timeout Settings



The Oracle VM Server for SPARC provisioning job performs the following major tasks:

- Downloads the appropriate OS image.
- Initiates a net boot action on the Service Processor.
- 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.
- When selected, installs the SUNWJass package to harden the system. For example, you can install the SUNWJass package for version 1.2.
- Enables the Fibre Channel ports on the storage system.
- Installs and configures the Agent Controller. Do not install the Agent Controller manually on the Oracle VM Server host.

Selection of Oracle VM Server Version

Verify that the target system has the correct hardware and firmware configuration to support the installation of selected version of Oracle VM Server for SPARC.

When the supported hardware and firmware versions are not there, the provisioning job might fail. You cannot force set the version of the software to be installed. For the complete list of supported hardware, firmware version, and the supported Oracle VM

Server for SPARC version, refer to the Release Notes in the library <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html> for more information.

Important Notes for Installation

Some important notes for installing Oracle VM Server for SPARC in Oracle Enterprise Manager Ops Center:

- The OS in the Control Domain must have the default locale set to C. The Control Domain provisioning is supported on Oracle Solaris SPARC or x86 Proxy Controllers, but not on Linux Proxy Controllers.
- To provision the Oracle VM Server for SPARC software, the proper system model name must be populated on the service processor. The model name is not available on the service processor of the Sun Blade T6320 hardware.
- Oracle Enterprise Manager Ops Center does not support use of the LDOMs configuration manager in LDOMs 1.2.
- Do not attempt to reconfigure the Control Domain and logical domains using the native CLI.
- You must not unconfigure and reconfigure the agent installed on the Control Domain.

Manual Net Boot Initiation

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

1. Log in to the service processor of the target system.
2. When the target system is running, enter `halt`.
3. Enter the following command

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

Additional Configurations

The following configuration for the Control Domain is performed during installation of Oracle VM Server for SPARC and is not required 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*

Oracle VM Server for SPARC Management

After you have provisioned or discovered the Oracle VM Server for SPARC, you can dynamically manage the Oracle VM Server Control Domain resource configurations such as CPU Threads, Crypto Units, and memory. You can associate storage libraries and attach networks with Oracle VM Server for SPARC, create logical domains, and monitor the performance of logical domains.

You can also create server pools for Oracle VM Server for SPARC for an efficient management of your resources. For creating server pools, see [Chapter 17, "Server Pools"](#).

The following actions are available for managing the Oracle VM Server:

- [Edit Configuration and Tags](#)
- [Associate Storage Libraries](#)
- [Attaching Networks](#)

Edit Configuration and Tags

You can modify the Control Domain configuration that includes CPU Threads, Crypto Units, and memory. You can also modify the Name and Description of the Control Domain. Use the **Edit Attributes** option to modify the configuration.

The Oracle VM Server must have Oracle Solaris 10 10/09 OS or higher version to edit the attributes. When you modify the memory size, Oracle VM Server is rebooted unless it is running Oracle VM Server for SPARC 2.0 or higher version. Use **Edit Tags** to modify the existing tags and add new tags.

Associate Storage Libraries

Storage libraries are required for storing logical domain metadata, ISO images and to provide virtual disks for logical domain storage. You must first associate the storage library to the Oracle VM Server and then it is made available for all the logical domains under it.

You can associate the following libraries for Oracle VM Server:

- **Filesystem storage:** This includes the NAS storage libraries.
- **Block storage:** This includes the Static and Dynamic Storage libraries. The Static Storage library includes the exported LUNs of SAN and iSCSI storage servers that are not managed by Oracle Enterprise Manager Ops Center. The LUNs of storage servers managed in Oracle Enterprise Manager Ops Center are available through Dynamic Storage library. This library provide option to create new LUN when you add storage to the logical domains.

Select the option **Associate Library** to associate libraries with the selected Oracle VM Server for SPARC.

Attaching Networks

Attach networks to Oracle VM Server to provide networking facilities for the logical domains. You can make multiple connections to a network in the Oracle VM Server Control Domain. For each network connection, a virtual switch is created and the switch has a naming pattern. For an example network 1.1.1.0/24, the virtual switches take the name as 1.1.1.0_24, 1.1.1.0_24_1, 1.1.1.0_24_2 and 1.1.1.0_24_3. This ensures that the switches have a unique name. When a network connection is made to the server, the virtual switch created is incremented. When you create and start a logical domain, you define the virtual switch that connects to the logical domain. Each virtual switch must be connected to a NIC.

When you migrate logical domains, the switch name must be identical in the source and the target Oracle VM Server for SPARC.

See [Migrate Logical Domains](#) for more information about migrating the logical domains.

You can create IPMP groups and aggregate links in the Control Domain. Navigate to the Oracle Solaris OS of the Control Domain and create IPMP groups or Link Aggregation. Refer to [Chapter 7, "Networks"](#) for more information about creating IPMP groups and Link Aggregations in an Oracle Solaris OS.

To Attach Networks to Oracle VM Server

This procedure to attach networks is applicable for Oracle VM Server for SPARC in stand-alone mode.

1. Select the Oracle VM Server for SPARC in the Assets tree.
2. Click **Attach Networks** in the Actions pane.

The list of available networks in Oracle Enterprise Manager Ops Center are displayed. The list also displays the existing number of connections with the server.

3. Select one or more networks from the list. You can make multiple connections to a network.

Click **Next**.

4. Specify the number of connections for the selected networks. Increment the total number of connections.

Note: The number of connections does not limit the number of logical domains that can be connected to this network.

5. Specify the NIC and IP address for each network connection.

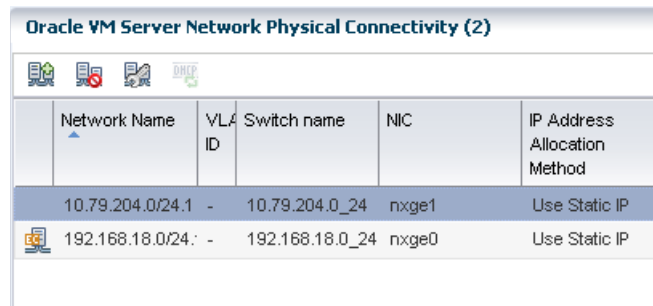
For each network connection, you must provide a NIC as a virtual switch is created for each connection. The virtual switch requires a physical interface to allow communication between the logical domains and external network. Follow these rules while assigning the NIC and IP address:

- Specify the same NIC to different network when the network has different VLAN ID. Otherwise, you cannot assign the same NIC to different networks.
- When supported by the network, select System Allocated to enable the system to allocate the NIC and IP address.
- Select Do Not Allocate IP for IP address when you do not want to assign IP address to the selected network interface.

Click **Next**.

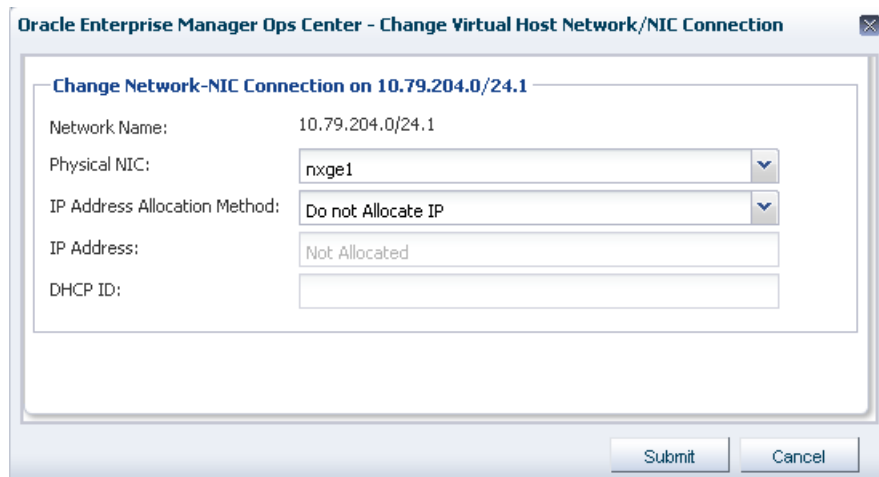
6. Review the summary and click **Finish** to attach the selected network to the Oracle VM Server.

To manage the attached networks, use the icons on the Network tab of the Oracle VM Server for SPARC, as shown in [Figure 15-4](#).

Figure 15–4 Network Options


Network Name	VLAN ID	Switch name	NIC	IP Address Allocation Method
10.79.204.0/24.1	-	10.79.204.0_24	nxge1	Use Static IP
192.168.18.0/24.1	-	192.168.18.0_24	nxge0	Use Static IP

- **Refresh DHCP connectivity** – Reconnects to the DHCP server.
- **Unbind networks** – Disconnects the selected network from Oracle VM Server.
- **Modify physical connectivity** – Changes the connectivity attributes of the selected network, as shown in [Figure 15–5](#).

Figure 15–5 Modify Physical connectivity


Change Network-NIC Connection on 10.79.204.0/24.1

Network Name: 10.79.204.0/24.1

Physical NIC:

IP Address Allocation Method:

IP Address:

DHCP ID:

Maximum Transmission Unit (MTU) Size

The default size of the network's MTU is 1500 bytes. A network can be created with MTU size that varies between 576 and 9216. Oracle Enterprise Manager Ops Center configures the IP configuration of Oracle VM Server for SPARC with the MTU value set at the network level. For a logical domain, the network to be attached must have minimum MTU size as 1500 bytes.

The network interface cards driver configuration must be updated to support MTU size greater than or equal to 9216 in the following scenario:

- You have VLAN tagged networks created over the virtual switches with Link Aggregation as the uplink.
- You are using Oracle Solaris 11 OS.

Otherwise, the tagged networks become unreachable, and Oracle VM Server and logical domains cannot connect to these networks.

Reboot Oracle VM Server

You can reboot the Control Domain regardless of the state of the logical domains in it. During the reboot, the logical domains might suspend temporarily but keep their current states.

Performance Management

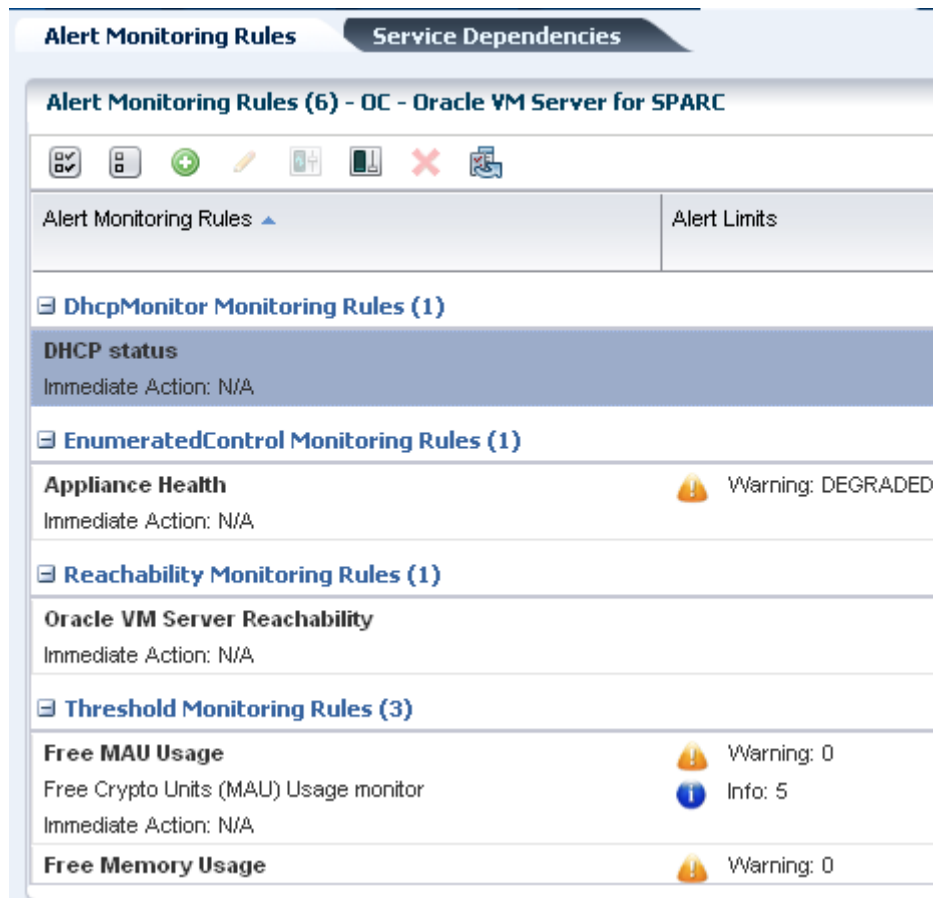
You can view the following information about Oracle VM Server performance:

- The Summary tab displays the Oracle VM Server for SPARC configuration information, and its health status.
- The Dashboard displays the general information about the selected Oracle VM Server for SPARC. A graphical representation of the group members and its relationship are displayed. The unassigned incidents and the list of recent incidents are also displayed.
- The Analytics tab displays the CPU, memory and network utilization by the logical domains in the Oracle VM Server. The consumption is represented in the form of charts.
- The Incidents tab lists the incidents that are reported for the Oracle VM Server. The tab provides the list of unresolved incidents and a graphical representation of the problem composition.

To see the corresponding alerts for each incident, see [Chapter 9, "Incidents"](#) for more detailed information.

- The Monitoring tab, as shown in [Figure 15–6](#), shows the monitoring values and boundaries that are set for Oracle VM Server activity. You can review the rules, monitoring attributes, specific time period for monitoring, and the level of severity for monitoring the activity of Oracle VM Server. To change these values, see [Chapter 4, "Monitoring Rules and Policies"](#).

Figure 15–6 Oracle VM Server for SPARC Monitoring Options



- The Configuration tab displays information such as remote logging, routing, NFS domain name, and name service information for the Oracle VM Server.
- The Charts tab displays the CPU, memory and network utilization of the Oracle VM Server. The data are collected in five-minute intervals and displayed graphically. The utilization data are provided for different time intervals. See [Chapter 12, "Operating Systems"](#) for more information about reading charts.

Server Pools

You can create server pools of your Oracle VM Server for SPARC systems in Oracle Enterprise Manager Ops Center. An Oracle VM Server must not have logical domains in running, shutdown, or suspended states when you want to add it to a server pool.

The detailed description and procedure for creating server pools are described in [Chapter 17, "Server Pools"](#).

About 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.

Using Oracle Enterprise Manager Ops Center, you can create logical domains and install Oracle Solaris OS on them. Using profiles and deployment plans, you can create more than one logical domain simultaneously and then save the configuration for future use.

You can use one of the following methods to create logical domains:

- **Create Logical Domain Profile and Plan**

1. Create a profile which defines the configuration of the logical domain.
2. Using the profile, create a deployment plan.
3. Apply the plan on an Oracle VM Server to create logical domains.
4. The logical domains do not have the OS installed. You must select each logical domain and apply an OS provisioning plan to install the OS. See [OS Provisioning on Logical Domains](#) for provisioning OS on the logical domains.

- **Configure and Install Logical Domains**

This is a complex plan which contains deployment plans that create logical domains and install Oracle Solaris OS on them. You must have the required profiles and deployment plans available to create the complex plan. The outline to create a plan is as follows:

1. Create a logical domain profile.
2. Create an OS provisioning profile to install the OS.
3. Create a deployment plan to install the OS and other updates if any.
4. Create the Configure and Install Logical Domains plan.

You can now manage logical domains that were created manually using the native CLI in Oracle Enterprise Manager Ops Center. See [Manage Existing Oracle VM Server for SPARC Environments](#) for more information.

Creating a Logical Domain Profile

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.

Ensure that you have the following information before you create a profile:

- **CPU and memory requirements**

The physical CPUs of the Oracle VM Server for SPARC are shared among the CPU threads of all the logical domains. Each logical domain requires:

- At least one GB of memory.
- At least one CPU thread.

Each CPU thread can be allocated independently to logical domain. The CPU threads are grouped into cores. For example, there are 4 Threads per core in the Oracle Sun Fire and SPARC Enterprise T1000 CMT processors. Some hardware resources are provided on a per-core basis and therefore, shared between the threads in the core. When the threads in a core are allocated to two or more separate logical domains, it can lead to suboptimal performance

of those threads. To get the best performance, it is best to avoid allocating the threads of a core to many logical domains. The best approach would be create large logical domains first, the logical domains with complete cores, and then the smaller logical domains.

Crypto units are assigned based on CPU thread assignments. You can request the number of crypto units to be assigned to the logical domain. However, the number of crypto units assigned might be different than the amount requested because we can only allocate a crypto unit for every given number of CPU thread allocation, depending on the server hardware. After the creation of logical domain, view the job notification to see the actual number of Crypto Units assigned to the logical domain. You can also edit the logical domain configuration, storage, and network resource configuration later.

- **Storage resources**

The libraries for virtual disk can be local, local device, NAS, SAN, or Dynamic Storage libraries.

- **Network resources and number of connections for a network**

You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain. The number of virtual network devices created depends on the number of virtual switches present in the Control Domain.

When you start a logical domain, you must define the virtual switches through which you connect the logical domain to the external network.

To Create a Logical Domain Profile

1. Select the Plan Management section in the Navigation pane.
2. Expand Profiles and select Logical Domain from the list.
3. Click **Create Profile** from the Actions pane.

The Create Logical Domain Profile Wizard is displayed.

4. Provide a name and description to identify the profile.
5. Select Create a deployment plan for this profile to automatically create a plan using this profile.
6. Click **Next** and provide a name and the starting number for the logical domain.

Using this profile, you can create more than one logical domain. To identify each new logical domain, provide a prefix start name and starting number. For example, for a Start Name defined as TestDomain and a starting number of 10, three new logical domains have the name TestDomain10, TestDomain11, and TestDomain12.

7. Click **Next** to configure the CPU Threads and memory.

The physical CPUs of the Oracle VM Server for SPARC are shared among the CPU threads of all the logical domains. Each logical domain requires:

- At least one Gigabyte of memory.
- At least one CPU thread.

If required, enable the Automatic Recovery option and the enter the value for Priority of Recovery. The value must be between 0 and 100. A logical domain with a higher value is recovered first.

8. Click **Next** to specify the library for storing the domain metadata and storage disks for the logical domain.
9. Select a library to store the logical domain metadata.
Select only the local or NAS library to store the domain metadata.
10. Select one or more libraries that form the logical domain storage disks.
The libraries for virtual disk can be local, local device, NAS, SAN, or Dynamic Storage libraries. Click **Next** to specify the networks for the domains.
11. Select at least one network from the list of networks identified by Oracle Enterprise Manager Ops Center. You can associate a logical domain with more than one network.
12. Enter the number of connections for each network.

You can connect to a network multiple times to the logical domain. When you create a logical domain, the number of connections for a network translate to the number of virtual network devices created for the logical domain. The number of virtual network devices created depends on the number of virtual switches present in the Control Domain.

When you start a logical domain, you must define the virtual switches through which you connect the logical domain to the external network.
13. Click **Next** to view the summary of the details selected for creating a logical domain.
14. Review the information and click **Finish** to save the profile.

When you want to create a deployment plan with this profile in the first step, then a corresponding logical domain plan is also created.

Creating and Deploying a Logical Domain Plan

When you create a logical domain profile, you can select to create a deployment plan with that profile. You can also create logical domain plans. You can either create a simple plan to install only the logical domains or a complex plan to create the logical domains and provision OS on it. The following procedure describes how to create a simple plan to install the logical domains only. You require to have created a logical domain profile to select in the plan.

To Create a Logical Domain Plan

1. Expand Deployment Plans in the Plan Management section and select **Create Logical Domains** plan.
2. Click **Create Plan from Template** in the Actions pane.
The Create a Deployment Plan window is displayed.
3. Enter a name and description for the plan.
4. Select the failure policy.
5. Select the logical domain profile in the Create Logical Domain step.
This is a simple single step plan to select the logical domain profile. The logical domain profiles that are available are listed in the list of Associated Profile/Deployment Plan.
6. Enter the number of logical domains to be created.

7. Click **Save** to create the deployment plan.

The deployment plan is created with the associated profile. Now, apply the plan on a suitable target to create the logical domains.

Applying Logical Domain Deployment Plan

Apply the created logical domain deployment plan on the target server and create the logical domains.

1. Expand Deployment Plans in the Plan Management section and select the created logical domain deployment plan.
2. Click **Apply Deployment Plan** in the Actions pane.
The Select Target Assets window is displayed.
3. Select the target server from the list.
Select one or more targets to apply the plan.
4. Click **Add to Target List** option to add the selected targets on which you want to apply the plan.
5. Select whether you want to run the plan with minimal interaction or you want to override the profile values.
6. Click **Next** to specify the resource assignments in the profile.
7. If required, modify the name of the logical domains starting name and the number appended to it.
Click **Next** to configure the CPU Threads and memory.
8. Edit the CPU Threads, memory, and automatic recovery as required.
Click **Next** to specify the storage resource assignments.
9. When the storage resources from the profile are not available for the selected target, it is flagged in red color. Modify the storage resources accordingly.
Click **Next** to specify the network resource assignments.
10. When the network resource assignments from the profile are not available for the selected target, then it is flagged in red color. Modify the network resources accordingly.
Click **Next** to schedule the job.
11. Schedule the job to run now or at a later time. Click **Next**.
12. Review the summary and click **Apply** to execute the deployment plan on the selected targets.

OS Provisioning on Logical Domains

The Oracle OS version that is provisioned on the logical domains is independent of the OS version on the Oracle VM Server for SPARC Control Domain. You can provision Oracle Solaris 10 or Oracle Solaris 11 OS on the logical domains irrespective of the Oracle Solaris OS version of the Control Domain.

After you create a logical domain, you must provision an OS on it.

- Ensure that you have imported the appropriate ISO image into the library. If it is Oracle Solaris 10, then it must be at least Oracle Solaris 10 8/07.

- Create or identify an OS provisioning profile for SPARC. See [Provisioning Operating Systems](#) for more information.
- Create or identify a deployment plan with the OS provisioning profile and apply the deployment plan.

Note: When you want to monitor the OS provisioning on the logical domains, you must first enable the logical domain console. See [Connect to Logical Domain Console](#) for enabling the console.

You can also use network interfaces other than the default GB_0 to boot the OS. You can specify GB_x as the alternate boot interface while selecting the networks during OS provisioning.

Complex Plan

You have the Configure and Install Logical Domains deployment plan that comprises the step to create logical domains and provision OS on them. This complex plan steps are:

- Create Logical Domains Profile: See [Creating a Logical Domain Profile](#) for creating the logical domain profile.
- Install and Update OS Plan: See [About Multi-Step Deployment Plans](#) for creating the plan.

Create the required profile and plan to create this complex plan.

Using WANBoot

You can use WANBoot to provision OS on logical domains in the following scenarios:

- Using complex plans that creates logical domains and provisions OS on them
- Provision OS on newly created logical domains
- Provision OS on already provisioned logical domains

Management of Logical Domains

The following operations are available for managing the logical domains that were created using Oracle Enterprise Manager Ops Center:

- Edit logical domain attributes
- Start a logical domain
- Shut down a logical domain
- Reboot a logical domain
- Migrate logical domains
- Add storage to logical domains
- Connect networks to logical domains
- Automatic recovery of logical domains

Edit Logical Domain Configuration

Use the option **Edit Attributes** in the Actions pane of a selected logical domain. You can edit the name, description, CPU Threads, memory, Crypto Units and automatic recovery priority value.

Use the option **Edit Tags** to modify the tag values or to add new tags.

Shut 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 on the same Oracle VM Server or on a different target server.

Starting a Logical Domain

You can start a logical domain from a shut down state. When you start a logical domain, you can select to start the logical domain on the same target or on a different target server. Depending on the number of network connections to an Oracle VM Server for SPARC, you have that many virtual switches created for an Oracle VM Server for SPARC.

You must select the virtual switch that connects to the logical domain MAC address. A virtual switch can be connected to only one MAC address. You can also select not to connect to network for a virtual switch.

To Start a Logical Domain

1. Select the shutdown logical domain.
2. Click **Start** in the Actions pane.

The Start Logical Domain Wizard is displayed.

3. Select the server pool in which you want to start the logical domain.

When you shutdown a logical domain, the logical domain is disassociated from the Control Domain and it can be started on a server.

When you have networks attached to a logical domain and is not available in the selected server pool, then that network is not available when the logical domain starts.

4. Select the Oracle VM Server for SPARC in which you want to start the logical domain. You can select the same Oracle VM Server in which the logical domain is running already. Otherwise, select an Oracle VM Server for SPARC from the list.

The current load for all the Oracle VM Server for SPARC are displayed. Use this information to place your logical domain.

Click **Next** to specify the network interfaces.

5. Select the virtual switch that must be associated with the logical domain MAC address.

The number of virtual switched displayed depends on the number of network connections for the server pool. For each network connection, you have a virtual switch created. You can associate a virtual switch with only one MAC address.

6. Select **Do Not Connect** when you do not want the virtual switch to connect to MAC address.

Click **Next** to schedule the job.

7. Schedule the job to run now or at a later time.
8. Review the properties and click **Start** to run the logical domain.

Add Storage to Logical Domain

You can add virtual disks to a logical domain. The associated libraries to the Oracle VM Server for SPARC are available to be added as additional storage to logical domains. Use the option **Add Storage** in the Actions pane to add virtual disks to the logical domain. The logical domain can be in running state for adding storage.

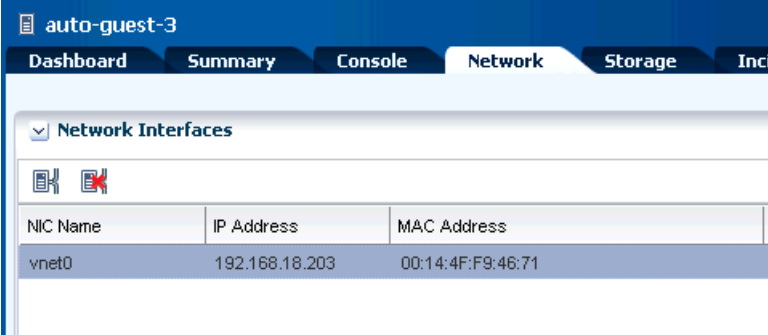
You can add virtual disks from the following type of storage libraries:

- **File system storage:** These are NAS libraries. Specify a virtual disk name and size of the disk.
- **Static block storage:** Select LUNs from SAN or iSCSI storage servers. The LUNs sizes are fixed.
- **Dynamic block storage:** Add LUNs from SAN or iSCSI storage servers discovered and managed in Oracle Enterprise Manager Ops Center. You can also create LUNs by selecting a volume group and specifying the size of the LUN.

Managing Logical Domain Networks

You can connect or disconnect a network and a logical domain. The logical domain can be in running state. You can connect to a network multiple times. Use the **Connect Network** and **Disconnect Network** icons, shown in [Figure 15-4](#), in the Networks tab.

Figure 15-7 Logical Domain Network Options



NIC Name	IP Address	MAC Address
vnet0	192.168.18.203	00:14:4F:F9:46:71

1. Click the **Connect Network** icon to display the Connect Guests to Network window. Select the network.
2. Select the virtual switch through which you want to connect the logical domain to the network. Depending on the number of network connections, the virtual switches are listed. The virtual switches can be reused. Therefore, you can attach multiple vnets on the same network on the same virtual switches.
3. You can create IPMP groups or aggregate the NICs allocated to the logical domain. See [Chapter 7, "Networks"](#) for more information about creating IPMP groups and link aggregation.

Migrate Logical Domains

You can migrate a logical domain from an Oracle VM Server to another Oracle VM Server. The type of migration is 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. To migrate domains, shared storage resources is required.

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 later versions, you can migrate with any number of MAUs.

From Oracle VM Server for SPARC 2.1 version onwards, live migration of the logical domains is supported.

Note: For Oracle VM Server for SPARC 2.1 and 2.2, you can migrate the logical domains only when the source and target machines are running at least Oracle VM Server for SPARC 2.1 version.

Setting User Accounts for Migration

Temporary user accounts are created dynamically when a logical domain migration is initiated. The user account is deleted after the migration. If your datacenter environment does not support dynamic creation of user accounts, then you can set the logical domain migration to use an existing user account instead of creating a temporary one. You must have the Security Admin role to create the credentials. If not, contact the Ops Center Administrator for creating a user account for logical domain migration.

To create a user account, follow the procedure:

1. Select the **Administration** section in the Navigation pane.
2. Select **Credentials** in the Administration section.
3. Click **Create Credentials** in the Actions pane.

The Create Credentials window is displayed.

4. Select the protocol as OVM For SPARC Migration and enter the following details:
 - Name of the credential
 - User name for the account
 - Password for the user account
5. Click **Create** to create the user account.

Set the following system property in the configuration management of the Enterprise Controller:

1. Select the **Administration** section in the Navigation pane.
2. Select the **Configuration** tab in the center pane.
3. Select **Virtualization** in the Subsystem.

The property Virtualization.Ldoms.Migration.Username is displayed.

4. Enter the user name of OVM For SPARC Migration user account created previously to be used for all logical domain migration.

5. Click **Save**.

The migration job proceeds to use the user account as defined in the system property. If the system property is set default to null, then the migration job dynamically creates a temporary user account. Ensure to set this property before initiating any logical domain migration, if required.

Migrating a Logical Domain

1. Select the logical domain to be migrated from the Assets section.

2. Click **Migrate** in the Actions pane.

The Migrate Logical Domain Wizard is displayed.

3. The eligible Oracle VM Servers are displayed in the decreasing order of preference. Select an Oracle VM Server from the list.

Click **Next**.

4. Review the information and click **Finish** to migrate the logical domain.

This migration is initiated from the logical domain and it results in migration of single logical domain. When the Oracle VM Server is placed in a server pool and there are more than one logical domain in the host, you can use Migrate Logical Domains to migrate more than one logical domain from the host.

Migrating Multiple Logical Domains

1. Select an Oracle VM Server that is placed in a server pool.

2. Click **Migrate Logical Domains** in the Actions pane.

The Migrate Logical Domains Wizard is displayed.

3. The list of logical domains running in the Oracle VM Server are listed. Select one or more logical domains from the list.

Click **Next**.

4. The eligible Oracle VM Servers in the same server pool that have the required resources to host all the logical domains are listed in the decreasing order of preference. Select an Oracle VM Server from the list.

Click **Next**.

5. Review the summary and click **Finish** to migrate logical domains.

Automatic Recovery of Logical Domains

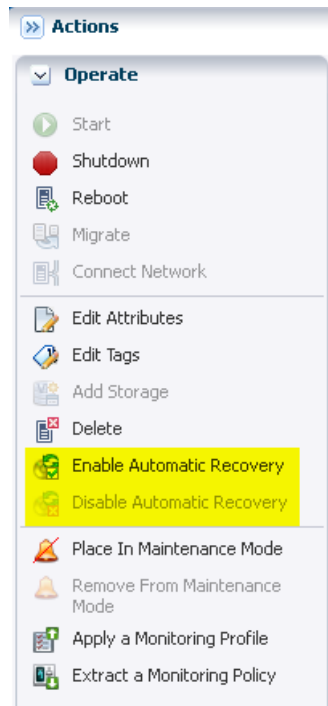
You can enable automatic recovery of logical domains. This ensures that when the Oracle VM Server for SPARC placed in a server pool fails, the logical domains are recovered and restarted on another server in the server pool.

You can set the priority of recovery value for the logical domain that decides the order of recovery of the logical domains in the Oracle VM Server for SPARC. The priority of recovery value can be set between 0 to 100. For example, if the priority of recovery value set for logical domain A is 12 and logical domain B is 15, then the logical domain B is recovered first.

Use the **Enable Automatic Recovery** and **Disable Automatic Recovery** actions, shown in [Figure 15-8](#), to set the automatic recovery for a logical domain. If you have disabled the automatic recovery of a logical domain, then it is not migrated to other servers in the pool when the underlying virtualization host fails. Instead, it is listed under

Shutdown Guests in the Server Pool. You can restart them on other servers as required.

Figure 15–8 Enable and Disable Automatic Recovery



When you edit the attributes of a logical domain, you can edit the priority of recovery value.

If there are no resources available to recover the logical domains in the server pool, Oracle Enterprise Manager Ops Center checks periodically for every minute for free resources to retry the automatic recovery mechanism.

You can also follow the procedure described in [Recovering Logical Domains](#) to recover the attached logical domains.

When a logical domain is configured for automatic recovery, the auto boot value for the logical domain is set to false. The auto boot value is controlled by Oracle Enterprise Manager Ops Center. This is because, when the failed server is repaired and restarted, the logical domains are not started automatically. Oracle Enterprise Manager Ops Center checks whether the logical domains are recovered and running on other servers. If the logical domains are not recovered, then the logical domains are started. If the logical domains are recovered on other servers, then Oracle Enterprise Manager Ops Center cleans up those logical domains on the server.

When you perform shut down and start operations using the UI for a logical domain that is configured to automatically recover, the operating system is also booted automatically by Oracle Enterprise Manager Ops Center.

When you perform the shut down and start operations using the CLI for a logical domain that is configured to automatically recover, you must boot the OS separately. This is not automatically done by Oracle Enterprise Manager Ops Center.

See [Chapter 17, "Server Pools"](#) and [Automatic Recovery](#) of Oracle VM Server for SPARC in a server pool for more information.

Recovering Logical Domains

The following procedure describes the actions that must be performed to recover the attached logical domains when an Oracle VM Server for SPARC server in a server pool has failed.

To Recover Logical Domains

1. Isolate the failed server.
Log in to the ALOM or ILOM of the physical server and shut down the server.
2. Power-off the failed server.
3. In the Oracle Enterprise Manager Ops Center UI, select the server pool in which the failed server is the member. Check whether the server is flagged as unavailable. This status is updated within 5 minutes approximately.
4. Select the Control Domain of the server in the UI and click **Delete Asset** in the Actions pane.
5. Select the server pool in which the failed server was a member and expand the **Shutdown Guests** list.
6. The attached logical domains of the failed server are listed. Select the logical domains and click **Start** to start the logical domains in the desired server in the server pool.
7. To bring back the failed server to the server pool, the server must be repaired and provisioned again with Oracle VM Server for SPARC using Oracle Enterprise Manager Ops Center.

Delete a Logical Domain

When you delete a logical domain, it is disconnected from the associated networks and is disassociated from the Oracle VM Server. All 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 plan remain unaffected.

When you delete a logical domain manually, that is using the native CLI, then the UI reflects the removal of the logical domain.

Connect to Logical Domain Console

You can attach to the logical domain console within the Oracle Enterprise Manager Ops Center UI. You enable the console connection and then you connect to the console.

Note: When the logical domain is not in view in the Assets tree, the console is logged out automatically but the connection exists until the connection time expires. You must log in again.

Manage Existing Oracle VM Server for SPARC Environments

You can manage existing Oracle VM Server for SPARC environments in Oracle Enterprise Manager Ops Center.

You can discover and manage Oracle VM Server for SPARC that has been provisioned manually. The underlying logical domains are automatically discovered and managed

in Oracle Enterprise Manager Ops Center. In the Assets section, the discovered logical domains are displayed in order and grouped under the Control Domain.

Logical domains that are created manually for an Oracle VM Server for SPARC provisioned through the UI are also automatically discovered and managed on the UI.

You can also provision OS on the manually created logical domains using the OS provisioning profile and plans in Oracle Enterprise Manager Ops Center.

You can perform the following functions on the UI for an existing logical domains:

- Start
- Shut down
- Reboot
- Modify the attributes of Logical Domain
- Delete Logical Domains

Limitations

The following are the limitations for managing such logical domains in Oracle Enterprise Manager Ops Center:

- Storage and network resources of the logical domains cannot be managed.
- Cannot be part of server pools.
- Cannot migrate the logical domains as it is not part of a server pool.

[Table 15–3](#) shows the options for managing existing logical domains as compared to managing new logical domains.

Table 15–3 Management Options for Oracle VM Server

Options	New Oracle VM Server for SPARC	New Logical Domain	Existing Oracle VM Server for SPARC	Existing Logical Domain
Modify attributes for Oracle VM Server	Yes	N/A	Yes	N/A
Reboot Oracle VM Server	Yes	N/A	Yes	N/A
Manage Networks	Yes (stand-alone)	Yes	No	No
Manage Storage	Yes (stand-alone)	Yes	No	No
Start Logical Domain	N/A	Yes	N/A	Yes
Shutdown Logical Domain	N/A	Yes	N/A	Yes
Reboot Logical Domain	N/A	Yes	N/A	Yes
Delete Logical Domain	N/A	Yes	N/A	Yes
Install Server Deployment Plan	Yes	Yes	No	Yes

Table 15–3 (Cont.) Management Options for Oracle VM Server

Options	New Oracle VM Server for SPARC	New Logical Domain	Existing Oracle VM Server for SPARC	Existing Logical Domain
Enable Console	N/A	Yes	N/A	Yes

Related Resources

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

- <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>
- [Chapter 2, "Asset Management"](#)
- [Chapter 17, "Server Pools"](#)
- [Chapter 7, "Networks"](#)
- [Chapter 6, "Storage Libraries"](#)
- *Oracle Enterprise Manager Ops Center Administration Guide*
- *Oracle Enterprise Manager Ops Center Deploying and Configuring Oracle VM Server for SPARC*

For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm

Oracle VM Server for x86

The following information includes:

- [Introduction to Oracle VM Server for x86](#)
- [Roles for Oracle VM Server for x86 Management](#)
- [Integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center](#)
- [Installing Oracle VM Manager](#)
- [Discovering Oracle VM Manager](#)
- [Discover Oracle VM Servers](#)
- [Installing Oracle VM Servers](#)
- [Administration of Oracle VM Manager](#)
- [Manage Oracle VM Servers](#)
- [Create Server Pool](#)
- [Create Virtual Machines](#)
- [Manage Virtual Machines](#)
- [Related Resources](#)

Introduction to Oracle VM Server for x86

Oracle VM Server for x86 is a platform that provides a fully equipped environment to leverage the benefits of x86 virtualization technology. Oracle VM Server enables you to deploy operating systems and application software within a supported virtualization environment.

Oracle VM Server for x86 is a Xen-based server virtualization technology which supports Linux, Oracle Solaris, and Windows guests, and provide features to manage guests lifecycle operations, allocate and monitor guest resource consumption. The components of Oracle VM are:

- **Oracle VM Manager:** Provides the user interface to manage Oracle VM Servers, virtual machines, and resources.
- **Oracle VM Server:** A self-contained environment that is designed to produce a lightweight, secure server-based platform for running virtual machines. The Oracle VM Server can perform one or more of the following functions:

- **Server Pool Master:** The master is the core of the server pool operations and it acts as the contact point for the server pool to Oracle VM Manager, and also as the dispatcher to other Oracle VM Servers in the server pool.
- **Utility Server:** Its function focuses on the creation and removal operations of virtual machines, Oracle VM Servers, and server pools.
- **Virtual Machine Server:** The primary function of virtual machine server is to run virtual machines.

Roles for Oracle VM Server for x86 Management

Table 16–1 lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 16–1 Oracle VM Server for x86 Roles and Permissions

Task	Role
Discover Oracle VM Manager and Oracle VM Server	Virtualization admin
Provision Oracle VM Manager	Ops Center admin
Manage Oracle VM Manager	Virtualization admin
Create Link Aggregation	Network admin
Create, manage, update, and delete virtual machines	Virtualization admin
Set monitor threshold	Asset admin
Manage storage resources	Storage admin
Create and manage profiles and plans	Profile and plan admin

Integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center

The integration of Oracle VM Server for x86 with Oracle Enterprise Manager Ops Center provides the platform to manage Oracle VM Manager, Oracle VM Servers, server pools, and the virtual machines through Oracle Enterprise Manager Ops Center UI.

The following operations can be done through Oracle Enterprise Manager Ops Center to manage Oracle VM Server for x86 deployments:

- Discover deployed Oracle VM Managers
- Provision Oracle VM Servers
- Discover existing Oracle VM Servers
- Launch Oracle VM Manager UI
- Create virtual machines
- Provision OS on virtual machines
- Create server pools
- Connect to Oracle VM Manager console

- Manage storage repositories of Oracle VM Server for x86
- Perform management operations on Oracle VM Servers and virtual machines

Installing Oracle VM Manager

Installation of Oracle VM Manager is outside the scope of this document. For detailed instructions about how to install the Oracle VM Manager, see *Oracle VM Installation and Upgrade Guide for x86*.

Discover the Oracle VM Manager in Oracle Enterprise Manager Ops Center and launch the console to access it.

Discovering Oracle VM Manager

When you discover the Oracle VM Manager in Oracle Enterprise Manager Ops Center, the following Oracle VM resources are automatically discovered and populated in the UI:


- Oracle VM Servers
- Server pools
- Virtual machines
- Storage servers
- Storage repositories
- Networks

To discover an existing Oracle VM Manager, you must have the following information:

- Discovery profile for Oracle VM Manager
- Credentials for Oracle VM Manager access

You must create a discovery profile for discovering the installed Oracle VM Manager. The discovery profile provides options to create the discovery profiles for all the components related to Oracle VM Server for x86.

[Figure 16–1](#) shows the different assets that are available under Oracle VM in the Discovery Profile Wizard.

Figure 16–1 Discovery Asset Type

To Create a Discovery Profile for Oracle VM Manager

The following procedure outlines the steps to create a discovery profile for Oracle VM Manager:

1. Select Discovery from the Plan Management section.
2. Click **Create Profile** in the Actions pane.
The Create Profile-Discovery Wizard is displayed.
3. Enter a name and description for the discovery profile.
4. Select the Asset Type as Oracle VM Manager from the list.
Click **Next** to define the tags.
5. (Optional) Define the tags for the asset to be discovered. Tags help to group the assets.
Click **Next** to enter the IP range.
6. (Optional) Enter the IP address of the Oracle VM Manager to be discovered. Enter the host name or IP Address when you execute the profile.
Click **Next** to select the credentials of the Oracle VM Manager.
7. When you have the credentials of the Oracle VM Manager, click **Select** and the list of saved credentials are displayed. Select the appropriate credential from the list.
Click **New** to display the Create Credentials window. Enter the following information:
 - Name and description for the credential.
 - Enter the username and password of the Oracle VM Manager. Re-enter the password to confirm.
 - Select the protocol and enter the port number.
 Click **Ok** to create the credential and click **Next** to review the summary of the Oracle VM Manager discovery parameters.
8. Click **Finish** to create the profile.

After creating the profile, execute the profile to discover the Oracle VM Manager. The following procedure describes how to execute the profile:

1. Select the discovery profile created for Oracle VM Manager.
2. Click the **Add Assets Using Profile** icon to launch the Add Assets Wizard.
3. When you have not provided the details of the IP address or host name in the profile, you must enter the details of the IP address or host name of the Oracle VM Manager.
4. Click **Add Now** to initiate the discovery job.

The Oracle VM Manager is considered an independent entity with which Oracle Enterprise Manager Ops Center integrates. The Oracle VM Manager appears in the Administration section of the UI, not as an asset in the Assets tree.

Figure 16–2 shows the Oracle VM Managers displayed in the Administration section of the UI.

Figure 16–2 Oracle VM Manager



When you discover an Oracle VM Manager, all of the Oracle VM Servers managed in it are also discovered and displayed in the Assets section. Existing server pools in the Oracle VM Manager are also displayed in the Server Pools view of the UI.

Discover Oracle VM Servers

When you discover and manage an Oracle VM Manager, all of the known Oracle VM Servers to the manager are also discovered and displayed in the software UI. When you have manually installed Oracle VM Servers and not discovered in the Oracle VM Manager, you can still discover the Oracle VM Server in the Oracle Enterprise Manager Ops Center UI.

When you create a discovery profile for Oracle VM Server, you must define the Oracle VM Manager to which it can be associated. This makes the Oracle VM Server known to the Oracle VM Manager.

You must use SSH discovery of the Oracle VM Server when you want access to the serial console and perform advanced configuration in the Control Domain. For example, in Oracle Exalogic systems, you must use SSH discovery of the Oracle VM Server to support the use of InfiniBand switches.

Discovery Profile for Oracle VM Server

The creation of discovery profiles is similar to Oracle VM Manager except for selecting the Discovery Type as Oracle VM Server. When you provide the Oracle VM Server host name or IP address, select the Oracle VM Manager with which you want to associate the Oracle VM Server.

Figure 16–3 shows the IP Ranges step of the Discovery Profile Wizard in which the user must select the Oracle VM Manager to discover the Oracle VM Server.

Figure 16–3 Select Oracle VM Manager

Execute the profile and discover the Oracle VM Server. This is a very basic discovery of the server. The storage and network configuration of the server are not populated in the UI and only basic information is displayed in the UI.

Discovering Oracle VM Servers Using SSH

Use this method when you want to access the serial console of the Oracle VM Server and perform advanced configuration.

To use SSH for discovery, define the credentials for the Oracle VM Server as explained in this procedure:

1. Select Credentials from the Plan Management section.
2. Click **Create Credentials** in the Actions pane.

The Create Credentials window displays.

3. Select the protocol SSH from the list.
4. Provide a name and description for the credential.
5. Select **Password** for Authentication Type.
6. Enter the root user and the password.
7. Click **Ok** to save the credential.

Create a discovery profile for the Oracle VM Server and select the credential created for ssh discovery. Discover the Oracle VM Server using the profile. Since the discovery of the Oracle VM Server is through the SSH discovery, the network and storage configuration of the server is populated in the UI.

Installing Oracle VM Servers

You can provision Oracle VM Servers using Oracle Enterprise Manager Ops Center. You can download an Oracle VM Server image into the storage libraries, create OS provisioning profiles and then deploy a plan with the profile on the selected target. When you import an image into Oracle Enterprise Manager Ops Center, a default profile and plan are created automatically. See [Uploading or Importing Images](#) to upload or import image into the software library.

Hardware and Software Requirements

Oracle VM Server is a managed virtualization environment based upon the Xen hypervisor technology and includes a small Linux-based management operating system. For hardware and software requirements for installing Oracle VM Server, see *Oracle VM Installation and Upgrade Guide for x86*.

Oracle VM Agent Installation

When you install Oracle VM Server through Oracle Enterprise Manager Ops Center, the Oracle VM Server is installed without Agent Controller. The Oracle VM Agent is installed automatically when you install Oracle VM Server. The Oracle VM Agent enables communication between the Oracle VM Manager and Oracle VM Server for all management tasks.

Oracle VM Server Password Management

The provisioning profile of installing Oracle VM Server is similar to installing a Linux OS profile with additional steps to provide the Oracle VM Server password and the management interface. The Oracle VM Server password is different from the root password. This password is used by the Oracle VM Manager to manage and monitor the Oracle VM Server and the virtual machines running in it. You must use this password while discovering the Oracle VM Server from Oracle VM Manager.

Creating OS Provisioning Profile for Oracle VM Server

You can create OS provisioning profile for provisioning Oracle VM Servers.

To Create a Provisioning Profile for Oracle VM Server

1. Select OS Provisioning profile in the Plan Management section.
2. Click **Create Profile** to create a new profile.
3. Provide a name and description for the profile.

4. Select **Oracle VM Server for x86** in the subtype.
Click **Next** to specify the OSP parameters.
5. Select the Oracle VM Server from the list.
Oracle VM Server does not have any specific software groups. Click **Next** to specify the OS setup.
(Optional) Select Include Custom Scripts to add some scripts to the installation. Select the scripts in the subsequent steps to execute.
6. Specify the time zone, language, terminal type and the root password for the OS. Also, specify the console baud rate and the serial port.
Click **Next** to specify the installation parameters.
7. Specify the following parameters for the Linux OS:
 - **Installation number:** Enter the installation number that is used to allow installation of all of the software that is included in your subscription.
 - **Partition action:** Select whether you want to change the disk partition of the system.
 - **Install protocol:** Specify HTTP or NFS as the install protocol.
 - **Kernel parameters:** If necessary, enter kernel parameters for the GRUB menu of the target system.
 - **Linux packages:** Specify the Linux packages to include or exclude during provisioning. To include a package, enter the package name in a line. To exclude a package, enter the package name preceded by a dash (-).

The options MD5 Checksum, Reboot action, Initialize Disk label, Use Shadow passwords, and Clear master boot record are enabled by default.

Click **Next** to specify the volume groups when you are using Logical Volume Manager.
8. Select **Use Logical Volume Manager** when you want to specify the volume groups in the Logical Volume Manager for the hard disk drives.
Click **Next** to specify the disk partitions and file systems.
9. Specify the disk partitions and file systems that you want to create on the target system. For the file system device, you can also select the logical device name.
Click **Next** to specify the naming service.
10. Specify the name service, domain name and the corresponding name server. Select one of the following name services:
 - **DNS:** Enter the domain name of the DNS server. Provide the IP address of the DNS server in the Name Server field. You can enter up to three IP addresses as the value for the Name Server. Provide the additional domains to search for name service information in the Domain Name Search List. You can specify up to six domain names to search. The total length of each search entry cannot exceed 250 characters.
 - **NIS or NIS+:** Enter the domain name of the NIS or NIS+ server. When you know the NIS server details, choose the option Specify an NIS Server and provide the NIS server host name and the IP address.
 - **LDAP:** Enter the domain name of the LDAP server. Specify the name of the LDAP Profile you want to use to configure the system. Enter the IP address of

the LDAP Profile Server. You can also optionally provide the Proxy Bind Distinguished Name and Password.

- **None:** Select None when there is no naming service configured.

Click **Next** to specify the network.

11. Select **None** to continue to specify the networks.

You can bond the interfaces after provisioning the Oracle VM Server. Select the Linux OS and aggregate the links.

12. Click the Add icon to add multiple networks.

All of the networks that are defined in Oracle Enterprise Manager Ops Center are displayed in the Network list. Enter the following information for each selected network:

- Select a NIC from the list of available logical interfaces for each network.
- Select the Address Allocation Method for the selected networks except the boot interface.

When you have selected Use Static IP for Address Allocation Method, then you must provide the IP address when you apply a plan with this profile. The specific IP address is assigned to the target system after provisioning.

Click **Next** to specify the Oracle VM Server parameters.

13. Enter the password to be used by the Oracle VM Agent. This password is used by the Oracle VM Manager to manage and monitor the Oracle VM Server. Re-enter the password to confirm it.

Figure 16–4 shows the step to enter the Oracle VM Server password.

Figure 16–4 Oracle VM Server Password

14. Specify the management interface, then click **Next**.

15. Review the information in the Summary, then click **Finish** to create the profile.

Applying Deployment Plan for Oracle VM Server

When you create a profile for provisioning Oracle VM Server, select the option to create a deployment plan with the profile. Either apply the created plan or create your own deployment plan.

To Apply a Deployment Plan for Oracle VM Server for x86

1. Select the plan created for provisioning Oracle VM Server in the Plan Management section.

2. Click **Apply Deployment Plan** in the Actions pane.
3. Specify the network resources defined in the profile. Provide the IP address for the boot interface.
4. (Optional) Provide a host name for the Oracle VM Server.
Click **Next** to select the Oracle VM Manager.
5. (Optional) Select an Oracle VM Manager from the list to discover and manage the Oracle VM Server. Select **None** when you do not want to associate with an Oracle VM Manager.

Place the Oracle VM Server in one of the server pools of the selected Oracle VM Manager. After provisioning, a separate job is initiated to add the Oracle VM Server to the server pool.

Click **Next** to schedule the job.
6. Schedule the provisioning job to run now or at a later time.
7. Review the summary and click **Apply** to launch the provisioning job.

Administration of Oracle VM Manager

The list of Oracle VM Managers discovered and managed in Oracle Enterprise Manager Ops Center are displayed in the Administration section.

Select an Oracle VM Manager and the following information is displayed as shown in Figure 16–5.

Figure 16–5 Oracle VM Manager View

The screenshot displays the 'Summary' tab of the Oracle VM Manager view. The main content area is divided into two columns. The left column lists various attributes: Oracle VM Manager Name, UUID, Description, Version, and counts for Server Pools, Servers, Virtual Machines, Storage Servers, Storage Repositories, and Networks. The right column shows 'Reachable: Yes' and YUM repository information. Below these is a 'Tags' section with a table for Tag Name and Value. At the bottom, a 'Server Pools (1)' table lists the details of the 'OvmPool'.

Server Pool Name	Description	Oracle VM Servers
OvmPool	0	

The Summary page displays information about the Oracle VM Manager, the list of server pools, the discovered Oracle VM Servers, and the virtual machines that are created in the corresponding Oracle VM Servers.

Manage Ownership of Oracle VM Servers

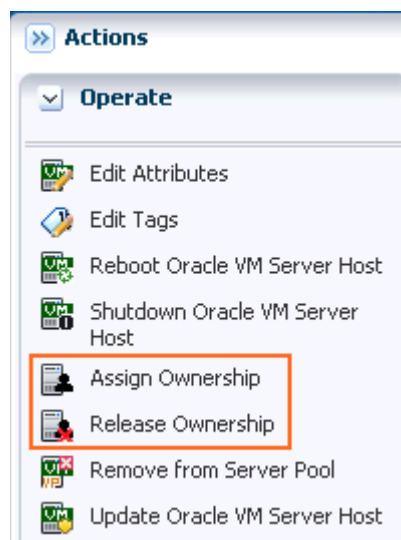
To use the Oracle VM Servers in the server pool and create virtual machines, it is required that you take ownership of the Oracle VM Server from the Oracle VM Manager.

Oracle VM Manager must take ownership of the Oracle VM Servers. Either assign ownership from the Oracle VM Server or take ownership through the Oracle VM Manager.

When you select an Oracle VM Server in the Assets section, the **Assign Ownership** and **Release Ownership** actions are available in the Actions pane.

Figure 16–6 shows the options Assign Ownership and Release Ownership displayed in the Actions pane.

Figure 16–6 Assign and Release Ownership From Oracle VM Server



When you select the Oracle VM Manager from the **Administration** section, the **Take Ownership** and **Release Ownership** options are available in the Actions pane.

When you take ownership of the Oracle VM Servers, add them to server pools and create virtual machines in them.

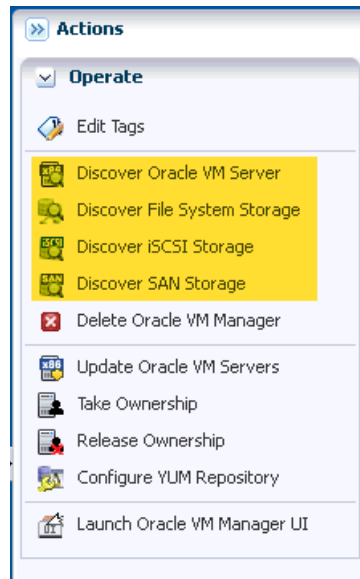
Discover Storage Resources

When you discover an Oracle VM Manager, Oracle Enterprise Manager Ops Center provides the option to discover the storage servers attached to it.

The following discovery options are available from the Oracle VM Manager:

- File system storage
- iSCSI storage
- SAN storage

Figure 16–7 shows the options to discover the storage resources from Oracle VM Manager.

Figure 16–7 Discover Storage Resources from Oracle VM Manager

Create discovery profiles for the storage servers and save them.

When you select to discover a storage server, existing discovery profiles are searched for the selected storage type. You can select from an available profile. When there are no profiles available, then the wizard to discover the corresponding storage server appears.

When you discover a storage resource, existing resources such as virtual disks, templates, ISO images, and virtual machine metadata are also discovered and displayed.

See [Chapter 6, "Storage Libraries"](#) for more information about the storage servers and the discovering them in Oracle Enterprise Manager Ops Center.

Manage Storage Resources

The following type of libraries setup are supported for storage of Oracle VM resources:

■ Filesystem Storage Libraries

Oracle Enterprise Manager Ops Center provides an option to create Oracle VM Storage Repositories. A storage repository is a logical disk space made available through a file system on top of physical storage hardware. The supported types of file system are NFS and OCFS. When the storage repository is created on an NFS file server, it is a NFS based storage repository. The NFS file server consists of NFS file systems. When the storage repository is created on a LUN, it is a LUN-based repository. The OCFS file system is created on the storage server. Create storage repositories on these file servers to be used by Oracle VM Servers to store resources. The resources include virtual machine metadata, templates, assemblies, ISO images and virtual disks.

To create Oracle VM storage repositories, you must have an Oracle VM Server discovered. Only NFS-based repositories can be shared by multiple server pools.

■ Static Block Storage Libraries

The LUNs from the storage arrays that are not managed by Oracle Enterprise Manager Ops Center form the static storage libraries. The LUNs are addressed by

iSCSI or Fibre Channel protocols. This forms the iSCSI and SAN static storage libraries. Also, add LUNs exported from managed storage arrays. The LUNs can belong to one library at a time.

- **Dynamic Block Storage Libraries**

The storage servers that are discovered and managed in Oracle Enterprise Manager Ops Center are presented under dynamic block storage libraries. The dynamic block storage libraries contains the exported LUNs from the storage array servers.

See [Chapter 6, "Storage Libraries"](#) for more information about discovering your storage servers, storage connect plug-ins, creating storage libraries and managing the libraries.

Manage Networks

During the installation of Oracle VM Server, the network interface used for the management is configured as a bonded interface. The bond is created with one interface and named as `bond0`. You can create additional bonds to add redundancy and load balancing of your network environment.

Attach networks to the Oracle VM Server or to the server pool that consists of a group of Oracle VM Servers on an Oracle VM Manager. Configure the network interfaces or the bonds to the network to be attached. You can assign different roles or functions to the networks attached to the Oracle VM Server.

The following are the network roles available for an Oracle VM Server:

- **Server Management:** Manages the Oracle VM Servers in a server pool. The Oracle VM Manager has one Server Management network.
- **Live Migrate:** Migrates the virtual machines from one Oracle VM Server to another in the server pool, without changing the state of the virtual machine.
- **Cluster Heartbeat:** Verifies that the Oracle VM Servers in the server pool are running.
- **Virtual Machine:** Monitors the network traffic between the virtual machines in a server pool.
- **Storage:** Transfers between virtual machines and virtual disks.

The management network created during the installation of Oracle VM Server has the following roles:

- Server Management
- Cluster Heartbeat
- Live Migrate

You can add and remove the roles of this management network, except for the Server Management role.

Depending on the available network interfaces on the Oracle VM Server, you can attach networks to Oracle VM Server and assign different roles to the networks. For example, you can attach the network in which your storage servers are placed and assign the Storage role to that network. You can assign a network with Live Migrate to be used only for migration.

For more information about setting different roles for network, refer to the *Oracle VM User's Guide*.

Attaching Networks

You can attach networks to the Oracle VM Server using the option **Attach Network**. When the Oracle VM Server is placed in a server pool and the server pool is associated with an user-defined network domain, you must add the network to the network domain.

1. Select the Oracle VM Server or the server pool.
2. Select **Attach Networks** from the Actions pane.

The Attach Network Wizard is displayed. The list of available networks in the network domain are displayed with the current connection to the Oracle VM Server.

3. Select one or more networks that you want to attach to the Oracle VM Server or the server pool.

Note: You cannot make multiple connections to the same network. You can connect to a network only once.

Click **Next** to configure the networks.

4. You cannot make more than one connection to the network. Do not increment the number of connections.

Click **Next** to configure the interfaces to the network.

5. For each selected network, select the NIC and the IP address.
 - **NIC:** Select an available NIC from the list. When the networks have different VLAN IDs, then you can assign the same NIC to another network. Otherwise, you cannot assign the same NIC to different network connection. When supported by the network, you can also select System Allocated for the system to take care of the NIC allocation.
 - **IP Address:** Select Use Static IP to provide the IP address manually or select Use System Allocated IP Address for the system to take care of the IP address allocation.

Click **Next** to view the summary of the selected information.

6. Review the information provided and click **Finish** to attach the networks to the Oracle VM Server or the server pool.

Manage Oracle VM Servers

When you install or discover the Oracle VM Server through the Oracle Enterprise Manager Ops Center UI, take ownership of the server from the Oracle VM Manager. This helps to place the Oracle VM Servers in the server pool. The following are some functions available for managing Oracle VM Servers:

- Edit Oracle VM Server information
- Reboot the server
- Shutdown the server
- Update the server
- Edit IPMI configuration

Edit Oracle VM Server Information

Select the Oracle VM Server and click **Edit Attributes** in the Actions pane. You can edit the following Oracle VM Server information:

- Name
- Description
- Tags

Editing IPMI Configuration

Intelligent Platform Management Interface (IPMI) allows you to remotely power off an Oracle VM Server, and to send a Wake-on-LAN message to power on an Oracle VM Server without having to physically press the power button. You can edit the IPMI configuration in the following method:

To Edit IPMI Configuration

1. Select the Oracle VM Server and click **Edit IPMI Configuration** in the Actions pane.

The Edit IPMI Configuration window is displayed as shown in [Figure 16–8](#).

Figure 16–8 IPMI Configuration

2. Enter the following information for IPMI configuration:
 - Username for the IPMI
 - Password for the IPMI
 - IP address of the IPMI
3. Click **Ok** to enter the IPMI configuration.

Placing Oracle VM Server in Maintenance Mode

To perform hardware or software maintenance, an Oracle VM Server can be placed in maintenance mode. When an Oracle VM Server is placed in maintenance mode, all of the virtual machines running on the Oracle VM Server are automatically migrated to other Oracle VM Servers in the server pool, if they are available, otherwise they are stopped. When the Oracle VM Server is the master Oracle VM Server in the server pool, this role is moved to another Oracle VM Server in the server pool, if available.

1. Select the Oracle VM Server and click **Place in Maintenance Mode** in the Actions pane.
2. Click **Place** to confirm the action.

The Oracle VM Server is placed in maintenance mode.

When you have finished with the maintenance on the Oracle VM Server and you are ready for it to rejoin the server pool, select the option **Remove From Maintenance Mode**.

Updating Oracle VM Server

Using a YUM repository, update or upgrade Oracle VM Servers. To access patch updates for Oracle VM, contact Oracle to purchase an Oracle VM Support contract and gain access to the Unbreakable Linux Network (ULN) which contains updates for Oracle VM. When you have access to ULN, you can use this to set up your own Yum repository to use when updating your Oracle VM Servers.

Configure YUM Repository

A YUM repository is required to update or upgrade Oracle VM Servers that are attached to an Oracle VM Manager. Set up a YUM repository and configure for automatic updates of the Oracle VM Servers managed by an Oracle VM Manager. Setting up a YUM repository is beyond the scope of this document.

After you set up a YUM repository, select the Oracle VM Manager and click Configure YUM Repository in the Actions pane. Enter the following details of the repository to configure it.

- Yum Repository Name: A name for the Yum repository.
- YUM Base URL: The URL to access the Yum repository, for example, the Oracle public YUM repository is at:
<http://public-yum.oracle.com/>
- Enable GPG Key: Select whether to use the GPG Key for the Yum repository. The GPG key or GnuPG key is the GNU project's implementation of the OpenPGP key management standard.
- Yum GPG Key: This field is enabled when you select Enable GPG Key. Enter the GPG key for the Yum repository, for example:

<http://public-yum.oracle.com/RPM-GPG-KEY-oracle-e15>

The GPG key must be available through HTTP, FTP, or HTTPS protocols. The GPG key for Oracle-signed updates from ULN is pre-installed on the Oracle VM Server at `/etc/pki/rpm-gpg/RPM-GPG-KEY-oracle`. When you want to use this GPG key, enter as:

`file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle`

To Update Oracle VM Server

1. Configure the Yum repository.
2. Select the Oracle VM Server and click **Update Oracle VM Server Host** in the Actions pane.

When you update the Oracle VM Server, the server is placed into maintenance mode and then the update is performed. When there are virtual machines in the Oracle VM Server, they are migrated to other Oracle VM Servers in the server pool, if they are

available, otherwise they are stopped. After the update, the Oracle VM Server remains in maintenance mode. Select the option **Remove From Maintenance Mode**.

Create Server Pool

You can group one or more Oracle VM Servers in an Oracle VM Manager and create server pools. You must create server pools before you start creating virtual machines. You can apply the virtual machines plans only on Oracle VM Servers that are in a server pool.

This section describes only how to create server pools using Oracle Enterprise Manager Ops Center. For more information about server pool policies, and managing server pools, see [Chapter 17, "Server Pools"](#).

When you create a server pool for Oracle VM Server for x86 systems, you must have the following information defined:

- Cluster file system
- Server pool master
- Virtual IP address for the server pool master

Cluster File System

Shared access to the server pool resources is a must for providing for high availability for the virtual machines running in the Oracle VM Servers of the server pool. This is achieved by cluster file system OCFS2 which allows multiple Oracle VM Servers to access the same disk at the same time. OCFS2 ensures that the Oracle VM Servers in a server pool can access and modify resources in the shared repositories in a controlled manner.

When you create a server pool, you must specify the server pool file system for the cluster heartbeat and other cluster information. The file system can be NFS shares or LUNs of iSCSI or SAN-based storage servers. Oracle VM formats the server pool file system as OCFS2 file system.

In Oracle Enterprise Manager Ops Center, the cluster is always enabled by default. The cluster configuration is pushed out to all the Oracle VM Servers in the server pool and the cluster heartbeat starts when the server pool is created. You can set a separate network for this cluster heartbeat. See [Manage Networks](#) for more information about setting up the networks and its role for an Oracle VM Server.

Server Pool Master

An Oracle VM Server is internally elected as Server Pool Master. You cannot set the role for an Oracle VM Server. When the elected Oracle VM Server fails, the role is set for another Oracle VM Server in the server pool. The virtual IP address provided while creating the server pool is assigned to the Oracle VM Server that has been elected as server pool master.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for x86 are as follows:

- Set the CPU utilization threshold. Places the virtual machines on the Oracle VM Server that has the lowest relative load.
- Set the CPU utilization threshold. Places the virtual machines on minimum number of Oracle VM servers to minimize the power consumption.

- You can select the automatic load balancing so that the virtual machines can be migrated automatically in the server pool whenever the thresholds are exceeded.

When you enable automatic load balancing for Oracle VM Server for x86 server pool, the server pool is checked continuously for the selected placement policy. When the threshold exceeds, the virtual machines are migrated from one Oracle VM Server to another.

When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM Server host servers are freed from the logical domains and powered off to minimize power consumption.

When the servers are overloaded and there are no other servers in the pool to host the virtual machines, then the policy decides to start a powered-off server using its Wake-on-LAN capability and live migrate the virtual machines to take up the load. The Wake-on-LAN capability must be enabled on the BIOS of the Oracle VM Server.

- When an Oracle VM Server fails, the virtual machines are migrated to another Oracle VM Server in the server pool, when you have selected Enable High Availability while creating the virtual machine.

Creating Oracle VM Server for x86 Server Pools

The following procedure describes the steps in the wizard to create an Oracle VM Server for x86 server pool.

1. Select Server Pool in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.
The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:
 - Name and description of the server pool.
 - Enter tags for categorizing your server pool.
 - Select Oracle VM Server for x86 from the Virtualization Technology.Click **Next** to define the server pool configuration.
4. Define the following configuration details:
 - Select the Oracle VM Manager in which the Oracle VM Servers are discovered and owned.
 - Select the network domain.
 - Select a network from the list of networks available in the network domain.
 - Enter the virtual IP address that is assigned to the server pool master. When the server pool master changes, the IP address is assigned to the new Oracle VM Server.
 - Select the server pool file system to store cluster heartbeat and other cluster information. The file system can be either NFS shares on a NFS file servers or LUNs of SAN and iSCSI based storages. The NFS file servers and the storage servers are known and reachable to the selected Oracle VM Manager.Click **Next** to select the members of the pool.
5. Select one or more Oracle VM Servers to add to the server pool.

The list of Oracle VM Servers displayed have the following characteristics:

- Owned by the Oracle VM Manager. Refer to [Manage Ownership of Oracle VM Servers](#) to own an Oracle VM Server.
- Not associated with server pool.
- In a healthy state.
- Not placed in maintenance mode.
- Do not have virtual machines in running, shutdown or suspended state.

Select Oracle VM Servers that are compatible for migration of virtual machines. To migrate virtual machines within a server pool, the Oracle VM Servers systems must be identical in model.

Click **Next** to associate the network domain.

6. Select the network domain to associate with the server pool.

- When you have selected default network domain, then go to Step 7.
- For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. You cannot bond the interfaces in Oracle VM Server for x86. Bonding can be done in only individual OS of the Oracle VM Server.

Click **Next** to select the networks and associate with the server pool.

7. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For an user-defined network domain, only the networks assigned to it are listed.

Select the networks that you want to associate with all the servers in the server pool. You cannot make multiple connections to a network. Limit the number of connections to 1.

Click **Next** to configure the interfaces of the servers.

8. For each network connection, specify the following details:

- Specify the NIC and IP address for each network connection. When a selected server is connected to the network, then no rows are displayed for that server.
- You can assign the same NIC to different network connection when the networks have different VLAN IDs.
- When supported by the network, you can select System Allocated for the NIC and IP address to be automatically allocated by the system.

Click **Next** to associate the storage libraries.

9. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store virtual machine metadata, ISO images, and for virtual disks of virtual machine.

You can select the following type of storage libraries:

- **Filesystem Storage Libraries:** This includes the Oracle VM storage repositories. You must have at least one Oracle VM storage repository associated with the server pool.
- **Block Storage Libraries:** This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of

storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

Click **Next** to select the server pool policies.

10. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
 - **Placement Policy:** This policy decides the preferred virtualization host in the server pool to place the logical domains.
 - **Auto Balancing Policy:** This policy performs load balancing of the server pool automatically at set intervals.
11. Select one of the following placement policies:
 - **Lowest relative load:** The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. Based on this, the server with the lowest relative load is considered to place a virtual machine. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and virtual machines are migrated to server with lowest relative load.
 - **Minimize power consumption:** The virtual machines running in a server pool are placed on the minimum number of Oracle VM Servers so that the unused Oracle VM Servers can be powered off. The threshold value set for the server over-utilization ensures that the servers are not over-loaded. Otherwise, the powered off Oracle VM Servers can be powered on to host the virtual machines.
12. Select the automatic load balancing policies:
 - Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the virtual machines automatically. Or send notifications about the approval.
 - You cannot manually balance the resources.

Click **Next** to view the summary of the server pool details.
13. Review the information to create a server pool for Oracle VM Server for x86. Click **Finish** to create the server pool.

Create Virtual Machines

A virtual machine comprises configurable set of resources and its own operating system. The resources include virtual CPU, memory, network, and virtual disk. You can start, stop, and restart each virtual machine independently. You can create virtual machines using one of the following installation sources:

- Templates
- Existing virtual machines
- ISO images in storage libraries
- Mounted ISO images on NFS, HTTP or FTP server
- Netboot

The operating system provisioned on the virtual machines can be hardware virtualized, paravirtualized and, hardware virtualized with paravirtualized drivers. Based on these virtualization types, you select the appropriate option to create the virtual machines.

Note: You can create virtual machines only in an Oracle VM Server for x86 server pools.

Virtualization Types

The following virtualization types are available for a virtual machine:

- **HVM**

Hardware virtualization or fully virtualized. For HVM, select an ISO image from the storage library and create a virtual machine. You must activate the hardware virtualization in the BIOS of the Oracle VM Server in which you want to create the virtual machine. The OS might be completely unmodified.

- **PVM**

Paravirtualized. You must use an ISO file mounted from an NFS share, HTTP, or FTP server. The OS is modified and recompiled to be made aware of the virtual environment. The paravirtualized guests run at near native speed, since most memory, disk and network accesses are optimized for maximum performance.

- **PVHVM or HVM with PV Drivers**

Hardware assisted virtualization with a paravirtualized driver. Install paravirtualized drivers on the hardware virtualized machines for optimized performance. This is mainly used for running Microsoft Windows guest operating systems.

Creating Virtual Machine Profile

Capture the virtual machine requirements in the form of profile in Oracle Enterprise Manager Ops Center. Create a deployment plan with this profile and apply it on the selected targets. Use these profiles and plans to create more than one virtual machine at a time.

To Create a Virtual Machine

1. In the Plan Management section, expand Profiles and Policies and select Virtual Machine from the list.
2. Click **Create Profile** in the Actions pane.
The Create Profile - Virtual Machine Wizard is displayed.
3. Enter the name and description of the profile.
Select Create a deployment plan for this profile to automatically create a plan using this profile. Click **Next** to specify the installation source of the virtual machine.
4. Select an installation source for creating the virtual machine. You have the following options to select from:
 - **Templates:** Templates are pre-configured, pre-installed virtual machine. You must upload the templates into Oracle VM Storage Repositories. In the profile,

the available storage repositories and the corresponding available templates are listed. Select a storage library and template.

- **Existing Virtual Machine:** When there are existing virtual machines, select the storage repository and the virtual machine image stored in it.
- **Install from ISO image (HVM):** You must have imported ISO images into the storage repository. Select the library and an ISO image from the list.
- **Install from mounted ISO image (PVM):** Mount the required ISO files to a NFS share, HTTP or FTP server and it must be accessible from the Enterprise Controller. Provide the network path of the mounted ISO file. For example, an ISO image accessed through HTTP server is represented like this.

`http://example.com/Enterprise-R5-U6-Server-x86_64-dvd.iso/`

- **Install from Network (Netboot):** Use this option only to create the virtual machine with Network as the boot order. You must provision the OS separately on it.

Select an option for the installation source and click **Next** to identify the virtual machines created.

5. Provide a name for the virtual machine and the number to be appended to the name. You can create one or more virtual machines simultaneously using this profile. To provide unique name for the virtual machines, enter the prefix name and the number to start the series.

For example, if the prefix name is VMachine, the number to start from is 10 and the number of virtual machines to be created is 3, then the virtual machines are created with the name VMachine10, VMachine11, and VMachine12.

6. Provide the description and tags for the virtual machine. All the virtual machines created with this profile carry the same description and tags. Use the tags to group the virtual machines based on the role.

Click **Next** to specify the configuration of the virtual machine.

7. When you have the installation source selected as templates or existing virtual machine, the values for all the parameters are set from the template. Otherwise, you must provide the following parameters:
 - **VM Type:** Based on the installation source, select the type as HVM, PVM or HVM with PV Drivers. For HVM, you must have the ISO image in the storage library. For PVM, you must use mounted ISO images.
 - **OS:** Select the OS from the list to be installed on the virtual machine.
 - **Enable High Availability:** High availability is to ensure the uninterrupted availability of a virtual machine. When an Oracle VM Server shuts down or fails in a server pool, all the virtual machines are migrated to other Oracle VM Servers in the server pool. High availability is implemented only in server pool.
 - **CPU Threads:** The number of CPU threads to be used by the virtual machine.
 - **CPU Priority:** The CPU priority of the virtual machine. You can select a high (100), intermediate (50), or low (1) priority for the virtual CPUs. The higher the priority more number of physical CPU cycles are allocated to the virtual machine.
 - **CPU Cap:** This parameter defines the maximum percentage to which the virtual CPUs can receive scheduled time. Set this option to keep the low priority virtual machines from using too many CPU cycles.

- **Memory:** The size of the memory the virtual machine is allocated to use when starting it.
- **Boot Property:** Select whether you want to start the virtual machine after creation.

Set these parameters and click **Next** to specify the boot configuration.

8. Select the boot media and order for the virtual machine. The boot options are set according to the installation source.

The default boot order for different installation source are:

- Templates – Disk
- Existing Virtual Machines – Disk
- HVM using ISO image – CDROM and Disk
- PVM using mounted ISO image – Network and Disk
- Netboot – Network

Select the correct order when you want to modify the boot order. Click **Next** to select the storage resources for virtual machine.

9. Select the storage resources to store the virtual machine configuration, also known as virtual machine metadata, and the virtual disks.

Use the following libraries for the virtual machine:

- **Oracle VM Storage Repositories:** The Oracle VM storage repositories are NFS file servers that are used to store virtual machine metadata, ISO images and virtual disk images. You can store the virtual machine metadata only in Oracle VM storage repositories. When you use the storage repository for virtual disk storage, you can either create a virtual disk or select existing virtual disks that are unshared and unused. You can also select ISO images to a virtual machine for HVM and PVHVM types.
- **Static Block Storage Libraries:** These are the storage array servers that are not managed by Oracle Enterprise Manager Ops Center but reachable to the Oracle VM Servers in the Oracle VM Manager. Select the LUNs that are exported as iSCSI and SAN storage types.
- **Dynamic Block Storage Libraries:** The exported LUNs that are available from the storage servers that are managed by Oracle Enterprise Manager Ops Center. Select the existing LUNs that are available from the storage servers. You can also create LUNs by specifying the volume group and the size of the LUN.

10. When you select to create virtual machines from existing virtual machines or templates, you are required to select the clone type for the virtual disks. You can select the following clone types:

- **Sparse type:** This clone type is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size.
- **Non-sparse type:** This clone type is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

Select the appropriate storage library, specify the virtual disks, disk size, and click **Next** to select the networks.

11. Select the network domain, and the networks that must be attached to the virtual machine.

You can connect to one or more networks. The number of connection to a network is limited to 1.

Click **Next** to view the summary.

12. Review the information provided to create a virtual machine profile. Click **Finish** to create the profile.

Create a deployment plan with this profile and apply the plan to create virtual machines. Except for the installation source Netboot, all the other source types results in creating a virtual machine with the OS installed. For virtual machine created with Netboot installation source, you must provision the OS on it separately using an OS provisioning plan.

The next section describes the different methods of deployment of virtual machine plans to create virtual machines.

Deploying Virtual Machine Plan

This section describes how to create virtual machines and provision OS in an Oracle VM Server using Oracle Enterprise Manager Ops Center.

The two new deployments plans for provisioning and creating multiple virtual machines are:

- Single step deployment plan to create only virtual machines
- Multi-step deployment plan to create virtual machine, provision OS, and install other applications.

In the single step deployment plan, the virtual machine can be created using Oracle VM templates or ISO images stored in Oracle VM repositories. Or only a virtual machine is created with ready to be booted using a separate OS provision plan.

In the complex deployment plan, you can configure and install the virtual machine, provision the OS, install software applications, apply monitoring profiles, and operational plans.

To Apply a Create Virtual Machine Plan

This is a single step virtual machine plan that creates a virtual machine and the OS is installed depending on the installation source selection. When you create HVM or PVHVM type virtual machines, select the ISO image to be used to create the virtual machine. You can select only one ISO image and multiple ISO files are not supported. While selecting the virtual disks, you have the option to select the ISO image. The following procedure describes how to apply the plan for creating a virtual machine:

1. Select the Virtual Machine plan in the Plan Management section.
2. Click **Apply Deployment Plan** in the Actions pane.
3. Select the targets on which you want to apply the plan. Click **Add Targets to List**.

You can apply the plan only on the Oracle VM Server for x86 server pools. You cannot create virtual machines on Oracle VM Servers that are not part of a server pool.

4. Select to apply the plan with minimal interaction when you do not want to modify the parameters in the profile.

Click **Next**.

5. You can modify the virtual machine identification, if required.

Click **Next** to specify the storage resource assignments.

6. The storage resources from the profile are populated. When the storage libraries are not available for a selected target, it is flagged in red. Modify the storage resources accordingly.
7. For the selected storage library, define the virtual disks.
 - For Oracle VM Storage repositories, you can create new virtual disks, select existing disks, or select ISO images. Select ISO image option is available only for HVM and PVHVM type virtual machines. Provide a disk name, select an existing disk or an ISO image.
 - For Static Block Storage, select an available LUN from the list.
 - For Dynamic Block Storage, either select an available LUN from the list or create a new LUN. When you create a new LUN, select the volume group and enter the size of the LUN.

Click **Next** to provide the network resources.

8. The network selected in the profile is populated. When the network is not available for a selected target, it is flagged in red. Modify the network resource.
9. Specify the network resource and the IP address for the virtual machine.

Click **Next** to view the summary.

10. Review the information provided and click **Apply** to apply the plan on the selected targets.

Except for the virtual machine created with Netboot as the installation source, all the other virtual machines are installed and provisioned with the OS.

Use the complex plan Configure and Install Virtual Machines to install the virtual machine and provision the OS on it. See [Chapter 8, "Plans and Profiles"](#) for more information about complex plans.

Provisioning OS on Virtual Machines

You can install an OS on virtual machines created on the Oracle VM Server. Use OS provisioning profiles to install the OS. See [Chapter 12, "Operating Systems"](#) for more information about creating provisioning profiles and plans to install an OS.

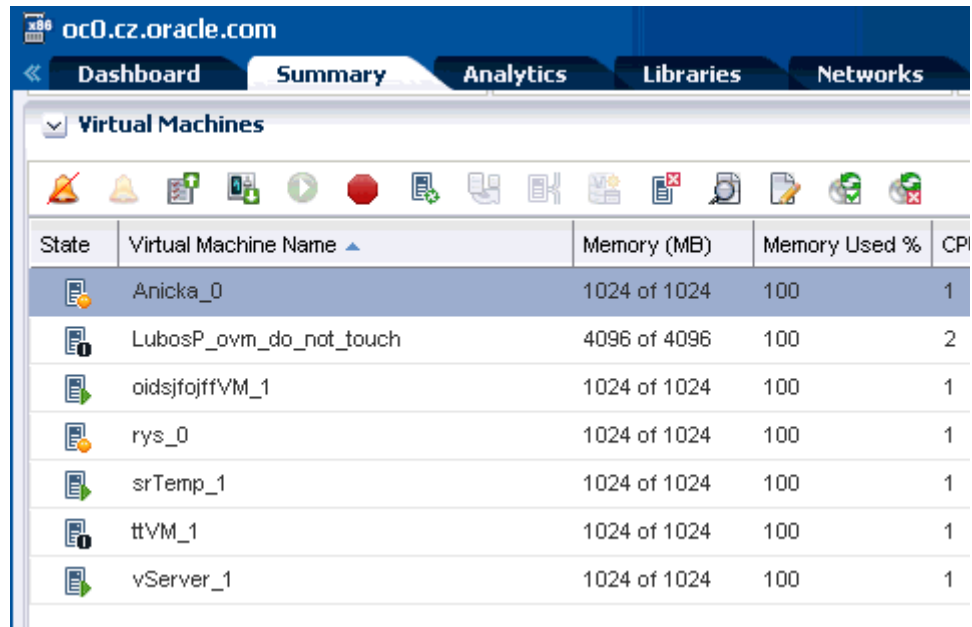
Manage Virtual Machines

Oracle Enterprise Manager Ops Center provides options to manage the lifecycle operations of virtual machines. You can start, suspend, resume, shut down, and delete virtual machines on the UI.

When you select the Oracle VM Server in the Assets section, the Summary tab in the center pane lists all the virtual machines that are currently running. All the options required to manage a virtual machine is available on the UI.

[Figure 16–9](#) shows the actions available for a virtual machine from Oracle VM Server.

Figure 16–9 Action Icons for Virtual Machines



State	Virtual Machine Name	Memory (MB)	Memory Used %	CPU
	Anicka_0	1024 of 1024	100	1
	LubosP_ovm_do_not_touch	4096 of 4096	100	2
	oidsjfojffVM_1	1024 of 1024	100	1
	rys_0	1024 of 1024	100	1
	srTemp_1	1024 of 1024	100	1
	ttVM_1	1024 of 1024	100	1
	vServer_1	1024 of 1024	100	1

Starting a Virtual Machine

Select the virtual machine and click the **Start Selected Guest** icon. The job is initiated to start the virtual machine.

Suspending a Virtual Machine

Select the virtual machine and click the **Suspend Selected Guest** icon. When you suspend a virtual machine, you cannot use the virtual machine. Click **Ok** to confirm the suspend action.

Rebooting a Virtual Machine

Select the virtual machine in the running state and click the **Reboot Selected Guest** icon. Click **Reboot Guest** to confirm the action.

Shutting Down a Virtual Machine

Select the virtual machine and click the **Shut Down Guest** icon. When you shut down a virtual machine, it is disconnected from its network. Click **Ok** to shut down the virtual machine.

Deleting a Virtual Machine

Select the virtual machine and click the **Delete Selected Guest** icon. Click **Ok** to confirm deleting the virtual machine. When you delete a virtual machine, all the references to the virtual machine such as the metadata and disk images are deleted from the system.

All these options are also available in the Actions pane when you select the virtual machine in the Assets section.

Edit Boot Order

The boot order that you have set during virtual machine installation is displayed. You can modify the boot order. You must specify the appropriate boot order depending on

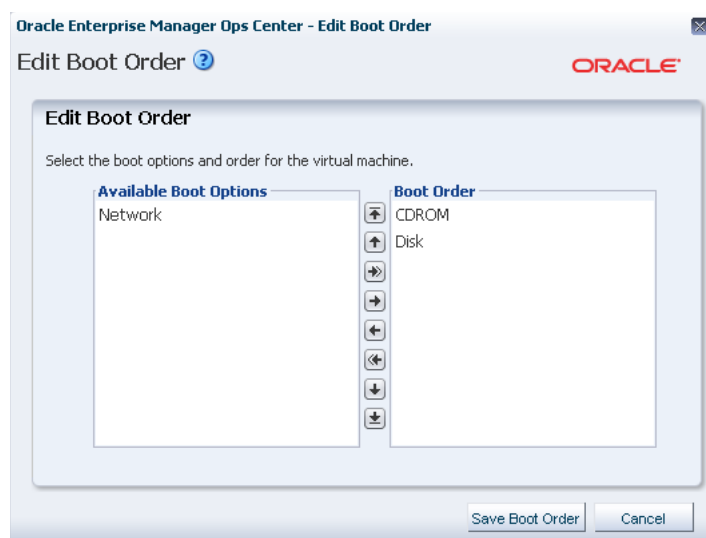
the installation source of the virtual machine. the following list shows the preferred boot options for different installation source:

- Hardware Virtualized Machine (HVM) – CDROM and Disk
- Paravirtualized Virtual Machine (PVM) – Network and Disk
- Templates and existing virtual machines – Disk
- Netboot – Network

If the virtual machine is in running state, then the modified boot order takes effect on the next reboot. Ensure to select the correct boot order for the virtual machine.

The **Edit Boot Order** option is available in the Actions pane of a selected virtual machine. [Figure 16–10](#) shows the Edit Boot Order window that is displayed to edit the boot order of the virtual machine.

Figure 16–10 Edit Boot Order Window



Modify the boot order and click Save Boot Order to save the changes.

Automatic Recovery

You can enable high availability for a virtual machine while creating a virtual machine profile. When the virtual machines are selected to be highly available and the Oracle VM Server shuts down or fails, the virtual machine is restarted on another available Oracle VM Server in the server pool.

You can also use the options **Enable Automatic Discovery** and **Disable Automatic Discovery** for a virtual machine to manage the high availability.

Related Resources

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

- [Chapter 2, "Asset Management"](#)
- [Chapter 17, "Server Pools"](#)
- [Chapter 7, "Networks"](#)

- [Chapter 6, "Storage Libraries"](#)
- *Oracle VM Installation and Upgrade Guide for x86*
- *Oracle VM User's Guide*
- *Oracle Enterprise Manager Ops Center Administration Guide*

For end-to-end examples, see the How To library at
http://docs.oracle.com/cd/E27363_01/nav/howto.htm

This chapter includes the following information:

- [Introduction to Server Pools](#)
- [Roles for Server Pool Management](#)
- [Server Pool Policies](#)
- [Automatic Recovery](#)
- [Server Pool Libraries](#)
- [Server Pool Networks](#)
- [Oracle VM Server for SPARC Server Pool](#)
- [Oracle Solaris Zones Server Pool](#)
- [Oracle VM Server for x86 Server Pool](#)
- [Manage Server Pools](#)
- [Related Resources](#)

Introduction to Server Pools

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

You can create server pools for the following types of virtualization servers:

- Oracle VM Server for SPARC
- Oracle VM Server for x86
- Oracle Solaris Zones

A virtualization host can refer to Oracle VM Server for SPARC, Oracle Solaris Zones, or Oracle VM Server for x86 that are managed through Oracle Enterprise Manager Ops Center. A virtual host or guest in a server pool refers to a non-global zone, logical domain or virtual machine running on the virtualization host.

Server pools are resource pools of virtualization 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 server pool manage many of the CPU utilization and resource balancing functions. Operations to the server pool are delegated to the individual virtualization hosts in the server pool.

To manage the guests within a server pool, you can perform warm and live guest migration and you can balance all of the guests' load among the members of the server pool. You can configure a policy to balance the load automatically, based on a schedule that you determine, or you can balance the load manually. When a virtualization host system shuts down, such as a hardware failure, you can start the guests on another host in the same pool.

Virtualization hosts in a server pool share network and storage libraries and several server pools can share the same networks and storage resources.

Server pools form a key concept for virtual datacenter management. Pooling your virtualization supported systems plays an important role and requires proper understanding of your requirements in your environment.

All the requirements, procedures, and different functions available for a server pool depend on the type of virtualization technology selected. See the appropriate sections to create a server pool for different virtualization technology.

- [Oracle VM Server for SPARC Server Pool](#)
- [Oracle Solaris Zones Server Pool](#)
- [Oracle VM Server for x86 Server Pool](#)

When you create a server pool, you must plan the storage and network resources such that they are accessible for all the members of the server pool.

You must set the policies to place the guests on the virtualization hosts and choose how you want to do the load balancing of the server pool. To understand more on this, refer to [Server Pool Policies](#).

Roles for Server Pool Management

The following table lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See [User and Role Management](#) in the *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 17–1 Server Pool Roles and Permissions

Task	Role
Create, manage, update, and delete server pools	Virtualization Admin
Provision and manage virtualization host	Virtualization Admin
Create, manage, update, and delete guests	Virtualization Admin

Server Pool Policies

When you create a server pool, you define guest placement, and automatic load balancing policies. When you create guests, you define the guest resource consumption, including physical and virtual CPUs. You can edit also the policies of a server pool.

Placement Policy

The placement policy is set to determine the preferred virtualization host for new guests within the server pool, automatic or manual load balancing of the server pool, and balance the server pool during any server failure or maintenance. The placement

policy is defined when a server pool is created. The following are the placement policy options:

- Place the guest on the virtualization host with the lowest relative load, based on the lowest CPU resource utilization. The calculation is based on a combination of the average load for the time period defined by the user. The default value is the last 10 minutes. The default threshold value for CPU utilization is 75%.
- Place the guest on the virtualization host with the lowest allocated CPU and memory, that is, the total static resource allocation across all guests on the host. The resource allocation is the sum of the number of vCPUs and virtual memory specified for each guest. This policy first verifies whether the resource allocation is possible and then ranks the server with available resources. The relative weight of CPU and memory resource is 1:1 while calculating the load on the host.
- Place the guest on minimum number of virtualization host and thereby consume less power.

You can modify the placement policy of a server pool. Following is an example of how the placement policy is worked upon a server pool.

Example 17-1 Example of Least Allocated Virtualization Host

A server pool has two virtualization hosts with different CPU and memory allocations. Host A has 8 available CPUs and 16 GB of memory. Host B has 4 available CPUs and 8 GB of memory.

Host A has three guests:

- Guest X has one vCPU and 1024 MB of memory.
- Guest Y has two vCPUs and 2048 MB of memory.
- Guest Z has one vCPU and 1024 MB of memory.

Host A's total static allocation is 4 vCPUs and 4 GB of memory.

Host B has one guest:

- Guest W has 3 vCPU and 4096 MB of memory.

Host B's total static allocation is 3 vCPUs and 4 GB of memory.

The allocation percentage for Host A is: CPU allocation is 4 vCPUs/8 physical CPUs, or 50%. Memory allocation is 4 GB/16 GB or 25%.

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

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

Minimize Power Consumption Policy

The policy to minimize power consumption places the guests on minimum number of virtualization hosts and powers off the idle servers. The idle virtualization hosts are powered off or set to low-power mode on explicit approval.

Automatic Load Balancing Policy

Scheduling the automatic balancing is applicable only for Oracle VM Server for SPARC and Oracle Solaris Zones server pool.

Use the Automatic Load Balancing Policy to schedule load balancing within a server pool. You can schedule the automatic balancing to occur weekly, daily, or hourly on a

specific day and time of the week. The default is to balance the load on the servers in the server pool every Saturday at midnight according to the defined placement policy.

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

When you do not want to balance the server pool's load automatically, schedule a reminder to balance the server pool's load manually.

Maintenance Mode

Automatic load balancing policy does not consider the virtualization hosts that are placed in maintenance mode as a target in the server pool for migrating the guests. When you place a virtualization host in maintenance mode, then all the migratable guests in it are migrated to other servers in the server pool. This action is triggered automatically.

Resource Balancing Enhancements

You can define the resource usage threshold for the servers in the server pool. The servers are considered to be loaded heavily depending on the threshold set for the resources. Depending on the placement policy selected, the CPU or total CPU and memory resource usage is calculated.

The virtualization hosts are checked whether the utilization exceeds the threshold. If so, the servers with enough resources are identified and then the guests that can be migrated from the server are identified to balance the load.

When the virtualization hosts are well within the threshold and power minimization policy is selected, the automatic balancing policy attempts to free a virtualization host from its guest so that the administrator can power off the server.

Automatic Recovery

The automatic recovery option is available for all types of supported virtualization technology. You can enable or disable the automatic recovery of zones, logical domains, and the virtual machines.

This feature provides the option to recover the guests that are attached to a failed server. The recovered guests are restarted on other servers in the server pool.

The automatic recovery for virtual machines is managed by Oracle VM Manager. For logical domains and zones, the automatic recovery is managed by Oracle Enterprise Manager Ops Center.

Recovery Process

Oracle Enterprise Manager Ops Center monitors the managed assets by checking the connection between the agent deployed on the assets and the Proxy Controller. To detect the unreachability of the assets, it takes about 5 minutes approximately. The unreachability detection steps are as follows:

1. A heartbeat time-out of 3 minutes allows the Proxy Controller to detect that the connection is down.
2. A re-connection attempt is made for every 30 seconds for four times.

When the connection cannot be re-established, the asset is seen as unreachable from Oracle Enterprise Manager Ops Center.

An automatic recovery job is initiated when the global zone or Oracle VM Server for SPARC is:

- Identified as unreachable by Oracle Enterprise Manager Ops Center, and its attached zones or logical domains are also seen as unreachable.
- Placed in a server pool.
- Has at least a zone or logical domain that is configured to recover automatically.

If an automatic recovery failed or cannot be performed because of non availability of free resources in the server Pool, Oracle Enterprise Manager Ops Center periodically checks for every minute if there are enough free resources and retries the automatic recovery.

You can also manually recover the zones and logical domains. See [Recovering Zones](#) for more information about manual recovery.

See the sections [Oracle Solaris Zones Server Pool](#), [Oracle VM Server for SPARC Server Pool](#), and [Oracle VM Server for x86 Server Pool](#) for more information about automatic recovery options.

Server Pool Libraries

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

The virtualization hosts in a server pool share the libraries associated with the server pool. The type of library you can associate with the server pool depends on type of virtualization hosts in the server pool and on whether all members of the server pool can get access to the storage resource.

See [Chapter 6, "Storage Libraries"](#) and [Chapter 5, "Software Libraries"](#) for more information about setting up these libraries.

Server Pool Networks

Oracle Enterprise Manager Ops Center provides groups related networking components such as fabrics and networks as network domains. All the managed and declared networks in Oracle Enterprise Manager Ops Center are available contained in the default network domain. Associating a virtualization host or a server pool with the default network domain does not require fabric connection. All the networks in the default network domain are available to be attached to the server pool.

You can also create network domains. Associating a server pool with a user-defined network domain requires connection between the physical interface of each server in the pool and each fabric in the network domain. Only networks assigned to the new network domain are available to be attached to the server pool.

Plan your networks and network domains that you must attach to the server pool. See [Chapter 7, "Networks"](#) for more information about creating network domains and managed networks.

Oracle VM Server for SPARC Server Pool

You can pool the Oracle VM Server for SPARC resources and manage your logical domains. Oracle VM Servers for SPARC must meet the following requirements to be added to a server pool:

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

Overview of the type of Oracle VM Servers present. You must group your servers that have the same CPU architecture and frequency. This enables migration of logical domains within the server pool. For example, you cannot migrate a logical domain of Oracle VM Server for SPARC 2.1 version to a server lower than this version. So, plan your servers that must be pooled.

- Must be running at least Oracle VM Server for SPARC 1.2 version.
- Must be running on at least the Solaris 10 10/09 OS and must meet specific patch and firmware requirements.
- Plan the network that must be attached to the server pool. Either place those networks in the user-defined network domain or use the default network domain. You can connect to a network multiple times. For each network connection, a virtual switch is created. For each network connection, you must define the NIC and IP address in each of the Oracle VM Servers in the server pool. Select the virtual switches when you start a logical domain. The multiple network connection allows you to create IPMP or aggregate the links in the logical domain.
- The storage libraries must be associated with the server pool to store logical domain metadata, ISO images, and for virtual disk storage of logical domains. You can associate file system and block storage libraries with the server pool. Filesystem storage includes the NAS storage libraries. Block storage includes the SAN, iSCSI, and Dynamic Storage libraries. Ensure that you have at least one NAS storage library assigned to the server pool. For using migration capabilities, the server pool must be on shared storage facility.
- Logical domain in the virtualization host must be created using Oracle Enterprise Manager Ops Center.
- Do not have logical domains in running, shutdown, or suspended state.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for SPARC are as follows:

- **Placement Policy**

Set the CPU utilization threshold. Place the logical domains on the Oracle VM Server for SPARC that has the lowest relative load.

Set the CPU and memory allocation threshold. Place the logical domains on the Oracle VM Server for SPARC allocated with lowest CPU and memory resources.

Set the CPU utilization threshold. Place the logical domains on the minimum number of Oracle VM Servers to minimize the power consumption.

- **Automatic Balancing**

Set the automatic load balancing. Logical domains are migrated automatically in the server pool whenever the thresholds are exceeded. Schedule the automatic load balancing to take place at definite time interval.

When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM

Server host servers are freed up from the logical domains and powered off to minimize power consumption.

Set whether administrator approval is required for migrating the logical domains.

Automatic Recovery

You can enable automatic recovery for a logical domain in the following scenarios:

- During logical domain creation, you can enable automatic recovery option and set the value for priority of recovery. See [Creating a Logical Domain Profile](#) for more information about setting the priority of recovery while creating a logical domain profile.
- For an existing logical domain, use the options Enable Automatic Recovery and Disable Automatic Recovery to set the automatic recovery. Also, you can edit the attributes of a logical domain to set the value for priority of recovery. See [Automatic Recovery of Logical Domains](#) for more information.

In an Oracle VM Server for SPARC server pool, the option to power-off the failed server and recover the logical domains is enabled by default. When a server fails, Oracle Enterprise Manager Ops Center tries to power-off the failed server, ensure that it is stopped and then initiates the automatic recovery of the logical domains. The logical domains that are configured for automatic recovery are recovered and restarted on other servers in the server pool.

When the option to power-off the failed server is enabled and Oracle Enterprise Manager Ops Center cannot power-off the failed server, that is the Service Processor is also unreachable, the automatic recovery of the logical domains does not take place. This is to avoid any potential data corruption during the recovery process of the logical domains.

When Oracle Enterprise Manager Ops Center detects that the failed server is unreachable, the logical domains that were running in the server disappear from the Navigation pane in the UI until they are recovered to other servers in the server pool.

The metadata of disappeared logical domains is still available in the storage libraries. The logical domains that are not configured for automatic recovery or the logical domains that cannot be recovered re-appear in the UI in the following scenario:

- When the failed server is restarted, the logical domains are booted and attached to the server again.
- When you have followed the manual recovery procedure and deleted the failed server from Oracle Enterprise Manager Ops Center, the logical domains re-appear in the list of Shutdown Guests in the Server Pool. You can start them on other server in the server pool. See [Recovering Logical Domains](#) for the logical domains recovery.

[Table 17-2](#) provides a quick view of different conditions that exist for recovering a logical domain.

Table 17–2 Recovery of Logical Domains

	Conditions	Result
Logical domains	<ul style="list-style-type: none"> Failed Oracle VM Server for SPARC The server pool option to power-off the failed server on automatic recovery is enabled Oracle Enterprise Manager Ops Center can power-off the failed server Automatic recovery is enabled for the logical domain 	<p>The logical domain is restarted and placed in another server in the server pool.</p> <p>If there are no server available in the server pool, then you must manually recover the logical domains.</p> <p>See Recovering Logical Domains for more information.</p>
Logical domains	<ul style="list-style-type: none"> Failed Oracle VM Server for SPARC The server pool option to power-off the failed server on automatic recovery is enabled Oracle Enterprise Manager Ops Center cannot power-off the failed server Automatic recovery is enabled for the logical domain 	<p>The automatic recovery of the logical domain does not take place.</p> <p>This is to avoid any data corruption during the recovery procedure.</p> <p>You must manually recover the logical domains.</p>
Logical domains	<ul style="list-style-type: none"> Failed Oracle VM Server for SPARC The server pool option to power-off the failed server on automatic recovery is not enabled 	<p>Oracle Enterprise Manager Ops Center proceeds to recover the logical domains configured to automatically recover without trying to power-off the failed server from its processor.</p> <p>This can lead to data corruption if the failed server is just isolated from the network point of view, and has still logical domains writing to disks.</p>
Logical domains	<ul style="list-style-type: none"> Failed Oracle VM Server for SPARC Automatic recovery is not enabled for the logical domain 	<p>Follow the procedure in Recovering Logical Domains to recover the logical domains.</p> <p>Only after this the logical domains appear in the Shutdown Guests list in the UI. You can restart them on the available server in the server pool.</p>

Creating an Oracle VM Server for SPARC Server Pool

Oracle Enterprise Manager Ops Center launches a wizard that collects information about the servers, storage and network resources, and placement policies to create a server pool.

To Create a Server Pool for Oracle VM Server for SPARC

1. Select Server Pools in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.

The Create Server Pool Wizard is displayed.

3. In the Identify Server Pool step, enter the following information:

- Name and description of the server pool.
- Enter tags for categorizing your server pool.
- Select Oracle VM Server for SPARC from the Virtualization Technology.

Click **Next** to select the members of the server pool.

4. Select one or more Oracle VM Servers to add to the server pool.

The list of Oracle VM Servers displayed have the following characteristics:

- Not associated with a server pool
- In a healthy state
- Not placed in maintenance mode
- Have no logical domains in the running, shutdown or suspended state

The list displays the details of CPU frequency, type, and architecture of Oracle VM Servers. Select Oracle VM Servers that are compatible for migration of logical domains.

Click **Next** to associate the network domain.

5. Select the network domain to associate with the server pool.

- When you have selected default network domain, then go to Step 8.
- For a user-defined network domain, select the physical interfaces of each server to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 8.

6. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

Click **Next** to configure the bonding parameters.

7. In the Configure Bonding step, specify the following parameters for Link Aggregation:

- Load balancing policy.
- LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
- LACP timer.
- MAC address policy and provide the MAC address.

Click **Next** to select the networks and associate with the server pool.

8. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For an user-defined network domain, only the networks assigned to it are listed.

Select the networks that you want to associate with all the servers in the server pool. You can make multiple connections to a network. Enter the total number of connections. For each network connection, a virtual switch is created. Choose the networks that you want to use for migration of logical domains.

Click **Next** to configure the interfaces of the servers.

9. For each network connection, provide the connection details:

- Specify the NIC and IP address for each network connection. When a selected server is connected to the network, then no rows are displayed for that server.

You can assign the same NIC to different network connection when the networks have different VLAN IDs.

- You can select System Allocated for the NIC and IP address to be automatically allocated by the system.
- You can select Do Not Allocate IP to allocate no IP address for the network connection.

Click **Next** to associate the storage libraries.

10. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store logical domain metadata, ISO images, and for virtual disks of logical domain.

You can select the following types of storage libraries:

- **Filesystem Storage Libraries:** This includes the NAS storage libraries. Associate at least one NAS storage library with the server pool to store the logical domain metadata.
- **Block Storage Libraries:** This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

Click **Next** to select the server pool policies.

11. Select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
 - **Placement Policy:** This policy decides the preferred virtualization host in the server pool to place the logical domains.
 - **Auto Balancing Policy:** This policy performs load balancing of the server pool automatically at set intervals.
12. Select one of the following placement policies:
 - **Lowest relative load:** The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. The server with the lowest relative load is considered to place a logical domain. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and logical domains are migrated to server with lowest relative load.
 - **Lowest allocated CPU and memory resources:** The total number of virtual CPU and memory resources allocated for all the logical domains in an Oracle VM Server are calculated. The server with the lowest allocated resources is considered to place a logical domain in the pool. Provide the threshold values for CPU and memory allocation. When the allocation is exceeded, the server is considered to be over-allocated and logical domains are migrated to server with lowest allocated CPU and memory resources.
 - **Minimize power consumption:** The logical domains running in a server pool are placed on the minimum number of Oracle VM Servers and the unused

Oracle VM Servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded.

13. Select the automatic load balancing policies:

- Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the logical domains automatically. Or send notifications about the approval.
- You can also select to manually balance the resources. See [Balancing Resources](#) for more information.

14. For Automatic Recovery, you can specify whether an attempt to power-off the failed server must be performed before initiating the automatic recovery of the logical domains.

When Oracle Enterprise Manager Ops Center has the capability to power-off the Service Processor of a failed server and you have selected the automatic recovery option, the failed server is powered off and then the logical domain recovery is started.

If Oracle Enterprise Manager Ops Center cannot power-off the failed server, the recovery of the attached logical domains does not take place. This is to avoid any data corruption during the recovery process.

Click **Next** to view the summary of the server pool details.

15. Review the information to create a server pool for Oracle VM Server for SPARC. Click **Finish to create the server pool.**

Oracle Solaris Zones Server Pool

You can now pool Oracle Solaris Zones resources and manage your zones. Oracle Solaris Zones must meet the following requirements to be added to a server pool:

- All the assets in the pool must have compatible architecture for supporting guest migration. For a global zone to be added to the server pool, it must have the same release or must be at least Oracle Solaris 10 10/08 version. This ensures that the zones can be migrated to a global zone in the server pool.
- The global zones must be in healthy state.
- The global zones must not be associated with a server pool.
- The global zones must not be placed in maintenance mode.

IP Stack in Server Pool

The network deployment for zones server pool vary depending on the Oracle Solaris OS version.

For a stand-alone Oracle Solaris 10 OS, the network can be attached to in either shared IP or exclusive IP mode. Whereas, in a server pool, you can attach a network only in shared IP mode for Oracle Solaris 10 OS. You cannot make multiple connections to a network.

For Oracle Solaris 11 OS, the network is always attached to in exclusive IP mode. In a server pool also, the network for Oracle Solaris 11 OS is always deployed in exclusive IP mode. You can make multiple connections to a network.

When you want to have a server pool with mixture of Oracle Solaris 10 and 11 OS, then you cannot make multiple connections to a network. So, plan your server pools with compatible Oracle Solaris OS versions.

Server Pool Policies

The server pool policies that are applicable for Oracle Solaris Zones are as follows:

■ Placement Policy

Place the zones on the global zone that has the lowest relative load. Set the CPU utilization threshold.

Place the zones on the global zone allocated with lowest CPU and memory resources. Set the CPU and memory allocation threshold.

Place zones on minimum number of global zones to minimize the power consumption. Set the CPU utilization threshold.

■ Automatic Balancing

You can select the automatic load balancing so that the zones can be migrated automatically in the server pool whenever the thresholds are exceeded. You can schedule the automatic load balancing to take place at definite time interval.

When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some servers are freed from the zones and powered off to minimize power consumption.

You can set whether the approval is required from the administrator for migrating the zones.

■ Automatic Recovery

You can set the automatic recovery for the zones in the following scenarios:

- During zone creation, you can enable automatic recovery of the zones. You can set the value for priority of recovery. Zone with a highest priority is migrated first.
- For existing zones, use the option Enable Automatic Recovery and Disable Automatic Recovery to set the automatic recovery. You can edit the zone configuration to modify the value for priority of recovery.

In an Oracle Solaris Zones server pool, you can specify that the automatic recovery of a failed server must start first with an attempt to power off the failed server, ensure that it is stopped and then initiate the automatic recovery of the zones.

Creating a Zones Server Pool

Oracle Enterprise Manager Ops Center takes you through a series of steps to collect information for creating a zones server pool.

To Create a Zones Server Pool

1. Select Server Pool in the Systems Group list on the Navigation pane.
2. Click **Create Server Pool** in the Actions pane.
The Create Server Pool Wizard is displayed.
3. In the Identify Server Pool step, enter the following information:
 - Name and description of the server pool.

- Enter tags for categorizing your server pool.
 - Select Oracle Solaris Zones – SPARC or x86 from the Virtualization Technology.
- You cannot have a zones server pool with a mixture of SPARC or x86 architecture.

Click **Next** to select the members of the server pool.

4. Select one or more global zones to add to the server pool.

The list of global zones displayed have the following characteristics:

- The same architecture – SPARC or x86
- Oracle Solaris 10 10/08 OS or higher version
- Not associated with a server pool
- In a healthy state
- Not placed in maintenance mode

The list displays the details of global zones such as CPU architecture and OS version. Select global zones that are compatible for migration of zones within the pool.

Click **Next** to associate the network domain.

5. Select the network domain to associate with the server pool.

- When you have selected default network domain, go to Step 8.
- For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 8.

6. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

Click **Next** to configure the bonding parameters.

7. In the Configure Bonding step, specify the following parameters for Link Aggregation:

- Load balancing policy.
- LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
- LACP timer.
- MAC address policy and provide the MAC address.

Click **Next** to select the networks and associate with the server pool.

8. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For a user-defined network domain, only the networks assigned to it are listed.

Select the networks that you want to associate with all the global zones in the server pool. Depending on the Oracle Solaris OS versions selected for the server pool, you can make multiple connections to a network. See [IP Stack in Server Pool](#) for more information about IP stack mode for networks in the zones server pool.

Enter the total number of connections and click **Next** to configure the interfaces of the servers.

9. For each network connection, provide the connection details:
 - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.

You can assign the same NIC to different network connection when the networks have different VLAN IDs.
 - When supported by the network, you can select System Allocated for the NIC and IP address to be automatically allocated by the system.
 - You can select Do Not Allocate IP to allocate no IP address for the network connection.

Click **Next** to associate the storage libraries.

10. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store zone metadata, ISO images, and for virtual disks of zones.

You can select the following type of storage libraries:

- **Filesystem Storage Libraries:** This includes the NAS storage libraries. Associate at least one NAS storage library with the server pool to store the zone metadata.
- **Block Storage Libraries:** This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

Note: When you have unmanaged storage attached to the non-global zones, then upload the script to be placed in all the global zones in the server pool. See [Script to Migrate a Zone With Dependencies](#) for more information about migrating zone with unmanaged storage.

Click **Next** to select the server pool policies.

11. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
 - **Placement Policy:** This policy decides the preferred virtualization host in the server pool to place the zones.
 - **Auto Balancing Policy:** This policy performs load balancing of the server pool automatically at set intervals.
12. Select one of the following placement policies:
 - **Lowest relative load:** The recent lowest memory and CPU utilization for the global zones in the server pool are calculated. The server with the lowest relative load is considered to place the zone. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and zones are migrated to server with lowest relative load.

- **Lowest allocated CPU and memory resources:** The total number of virtual CPU and memory resources allocated for all the zones in an Oracle VM Server are calculated. The server with the lowest allocated resources is considered to place a zone in the pool. Provide the threshold values for CPU and memory allocation. When the allocation is exceeded, the server is considered to be over-allocated and zones are migrated to server with lowest allocated CPU and memory resources.
 - **Minimize power consumption:** The zones running in the server pool are placed on the minimum number of global zone and the unused servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded. Otherwise, the powered off servers can be powered on to host the zones.
13. Select the automatic load balancing policies:
- Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the zones automatically. Or send notifications for the approval.
 - You can also select to manually balance the resources. See [Balancing Resources](#) for more information.
14. For Automatic Recovery, you can specify whether an attempt to power-off the failed server must be performed before initiating the automatic recovery of the zones.
- When Oracle Enterprise Manager Ops Center has the capability to power-off the Service Processor of a failed server and you have selected the automatic recovery option, the failed server is powered off and then the zone recovery is started.
- Click **Next** to view the summary of the server pool details.
15. Review the information to create a server pool for Oracle Solaris Zones. Click **Finish** to create the server pool.

Oracle VM Server for x86 Server Pool

You can group one or more Oracle VM Servers in an Oracle VM Manager and create server pools. When you create a server pool for Oracle VM Server for x86 systems, you must have the following information defined:

- Cluster file system
- Server pool master
- Virtual IP address for the server pool master

Cluster File System

Shared access to the server pool resources is a must for providing for high availability for the virtual machines running in the Oracle VM Servers of the server pool. This is achieved by cluster file system OCFS2 which allows multiple Oracle VM Servers to access the same disk at the same time. OCFS2 ensures that the Oracle VM Servers in a server pool can access and modify resources in the shared repositories in a controlled manner.

When you create a server pool, you must specify the server pool file system for the cluster heartbeat and other cluster information. The file system can be NFS shares or LUNs of iSCSI or SAN based storage servers. Oracle VM formats the server pool file system as OCFS2 file system.

In Oracle Enterprise Manager Ops Center, the cluster is always enabled by default. The cluster configuration is pushed out to all the Oracle VM Servers in the server pool and the cluster heartbeat starts when the server pool is created. You can set a separate network for this cluster heartbeat. See [Manage Networks](#) for more information about setting up the networks and their role for an Oracle VM Server.

Server Pool Master

An Oracle VM Server is internally elected as server pool master. You cannot set the role for an Oracle VM Server. When the elected Oracle VM Server fails, the role is set for another Oracle VM Server in the server pool. The virtual IP address provided while creating the server pool is assigned to the Oracle VM Server that has been elected as server pool master.

Server Pool Policies

The server pool policies that are applicable for Oracle VM Server for x86 are as follows:

■ **Placement Policy**

Places the virtual machines on the Oracle VM Server that has the lowest relative load. Set the CPU utilization threshold.

Places the virtual machines on minimum number of Oracle VM servers to minimize the power consumption. Set the CPU utilization threshold.

■ **Automatic Balancing**

You can select the automatic load balancing so that the virtual machines can be migrated automatically in the server pool whenever the thresholds are exceeded.

When you enable automatic load balancing for Oracle VM Server for x86 server pool, the server pool is checked continuously for the selected placement policy. When the threshold exceeds, the virtual machines are migrated from one Oracle VM Server to another.

When you have selected power minimization and automatic load balancing policy, and the servers in the server pool are not overloaded, some Oracle VM Server host servers are freed from the logical domains and powered off to minimize power consumption.

When the servers are overloaded and there are no other servers in the pool to host the virtual machines, then the policy starts a powered-off server using its Wake-on-LAN capability and live migrate the virtual machines to take up the load. The Wake-on-LAN capability must be enabled on the BIOS of the Oracle VM Server.

■ **Automatic Recovery**

When an Oracle VM Server fails, the virtual machines are migrated to another Oracle VM Server in the server pool. The automatic recovery for a virtual machine is set in the following scenarios:

- During virtual machine creation, select the option Enable High Availability. This ensures that the virtual machines are migrated when an Oracle VM Server fails.
- For an existing virtual machine, use the option Enable or Disable Automatic Recovery to set the automatic recovery.

Creating an Oracle VM Server for x86 Server Pools

The following procedure describes the steps in the wizard that is launched to create an Oracle VM Server for x86 server pool.

To Create an Oracle VM Server for x86 Server Pool

1. Select Server Pool in the Systems Group list on the Navigation pane.

2. Click **Create Server Pool** in the Actions pane.

The Create Server Pool Wizard is displayed.

3. In the Identify Server Pool step, enter the following information:

- Name and description of the server pool.
- Enter tags for categorizing your server pool.
- Select Oracle VM Server for x86 from the Virtualization Technology.

Click **Next** to define the server pool configuration.

4. Define the following configuration details:

- Select the Oracle VM Manager in which the Oracle VM Servers are discovered and owned.
- Select the network domain.
- Select a network from the list of networks available in the network domain.
- Enter the virtual IP address that is assigned to the server pool master. When the server pool master changes, the IP address is assigned to the new Oracle VM Server.
- Select the server pool file system to store cluster heartbeat and other cluster information. The file system can be either NFS shares on a NFS file servers or LUNs of SAN and iSCSI based storages. The NFS file servers and the storage servers are known and reachable to the selected Oracle VM Manager.

Click **Next** to select the members of the pool.

5. Select one or more Oracle VM Servers to add to the server pool.

The list of Oracle VM Servers displayed have the following characteristics:

- Owned by the Oracle VM Manager. Refer to [Manage Ownership of Oracle VM Servers](#) to own an Oracle VM Server.
- Not associated with a server pool
- In a healthy state
- Not placed in maintenance mode
- Have no virtual machines in the running, shutdown or suspended state

Select Oracle VM Servers that are compatible for migration of virtual machines. To migrate virtual machines within a server pool, the Oracle VM Servers systems must be identical in model.

Click **Next** to associate the network domain.

6. Select the network domain to associate with the server pool.

- When you have selected default network domain, go to Step 7.

- For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. You cannot bond the interfaces in Oracle VM Server for x86. Bonding can be done in only individual OS of the Oracle VM Server.

Click **Next** to select the networks and associate with the server pool.

7. For default network domain, all the networks that are declared and managed in Oracle Enterprise Manager Ops Center are listed. For a user-defined network domain, only the networks assigned to it are listed.

Select the networks that you want to associate with all the servers in the server pool. You cannot make multiple connections to a network. Limit the number of connections to 1.

Click **Next** to configure the interfaces of the servers.

8. For each network connection, specify the following details:
 - Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.
 - You can assign the same NIC to different network connection when the networks have different VLAN IDs.
 - You can select System Allocated for the NIC and IP address to be automatically allocated by the system.

Click **Next** to associate the storage libraries.

9. The storage libraries that are reachable from the selected members of the pool are displayed. The storage libraries are required to store virtual machine metadata, ISO images, and for virtual disks of virtual machine.

You can select the following type of storage libraries:

- **Filesystem Storage Libraries:** This includes the Oracle VM storage repositories. At least one Oracle VM storage repository must be associated with the server pool.
- **Block Storage Libraries:** This includes the Static Block Storage and Dynamic Block Storage. Static block storage libraries comprises the exported LUNs of storage arrays that are not managed by Oracle Enterprise Manager Ops Center. Dynamic block storage libraries comprises the exported LUNs of storage array servers that are discovered and managed in Oracle Enterprise Manager Ops Center.

Select the storage libraries from the list that you want to associate with all the virtualization hosts in the server pool.

Click **Next** to select the server pool policies.

10. You must select the following policies in the server pool to manage the under utilized and overutilized servers in the pool:
 - **Placement Policy:** This policy decides the preferred virtualization host in the server pool to place the logical domains.
 - **Auto Balancing Policy:** This policy performs load balancing of the server pool automatically at set intervals.
11. Select one of the following placement policies:
 - **Lowest relative load:** The recent lowest memory and CPU utilization for the Oracle VM Servers in the server pool are calculated. The server with the

lowest relative load is considered to place a virtual machine. Provide the threshold for CPU utilization above which the server is considered to be over-utilized and virtual machines are migrated to server with lowest relative load.

- **Minimize power consumption:** The virtual machines running in a server pool are placed on the minimum number of Oracle VM Servers and the unused Oracle VM Servers are powered off. The threshold value set for the server over-utilization ensures that the servers are not overloaded. Otherwise, the powered off Oracle VM Servers can be powered on to host the virtual machines.

12. Select the automatic load balancing policies:

- Select automatic load balancing and set the interval in which the server pool must be checked for resource balancing. You can set the approval to migrate the virtual machines automatically. Or send notifications about the approval.
- You can also select to manually balance the resources. See [Balancing Resources](#) for more information.

Click **Next** to view the summary of the server pool details.

13. Review the information to create a server pool for Oracle VM Server for x86. Click **Finish to create the server pool.**

Manage Server Pools

You must monitor the requirements of the server pool so that there are resources to host the guests and run them efficiently. Also, the server pools gives the infrastructure support for the virtual datacenters. You might need to add storage and network resources to run the guests. Oracle Enterprise Manager Ops Center provides the following management functions to manage the resources in the server pool:

- [Editing Server Pool Parameters](#)
- [Adding Virtualization Hosts](#)
- [Associating Network Domains](#)
- [Attaching Networks](#)
- [Associating Libraries](#)
- [Creating Guests](#)
- [Migrating Multiple Guests](#)
- [Migrating Zones](#)
- [Migrating Logical Domains](#)
- [Migrating Virtual Machines](#)
- [Balancing Resources](#)
- [Deleting Server Pool](#)

Editing Server Pool Parameters

You can modify the parameters set during the creation of a server pool. Click **Edit Attributes** in the Actions pane and you can edit the following parameters of a server pool:

- Name
- Description
- Change the placement policy
- Change the auto balancing policy

To add, remove or modify the tags added to the server pool, use the option **Edit Tags** in the Actions pane.

For Oracle Solaris Zones and Oracle VM Server for SPARC server pool, you can edit the following additional parameters:

- Enable or disable automatic recovery option.
- Upload scripts to manage the file systems on unmanaged storage. See Step 10 in [Oracle Solaris Zones Server Pool](#) for more information.

Adding Virtualization Hosts

You can add more virtualization hosts to a server pool when you require more CPU and memory resources to take up the load in the server pool. Depending on the virtualization type, you can add global zones, and Oracle VM Servers of x86 and SPARC architecture.

Oracle Solaris Zones

The list of global zones that are displayed to be added to the server pool has the following characteristics:

- Same release as the existing global zones in the pool or at least Oracle Solaris 10 10/08 OS or higher version
- Not associated with a server pool
- In a healthy state and not placed in maintenance mode

When the server pool has Oracle Solaris 11 OS and multiple network connections, then the list of available global zones is limited to only Oracle Solaris 11 OS. You cannot add Oracle Solaris 10 OS as you cannot make multiple network connections to it.

Load balancing in the server pool requires the zones to be migrated within the pool. Therefore, ensure compatibility for migration within the pool when you select the members.

Oracle VM Server for SPARC

The list of Oracle VM Servers displayed to be added to the server pool have the following characteristics:

- Compatible architecture with the other members of the server pool
- Not associated with a server pool
- In a healthy state and not placed in maintenance mode
- Have no logical domains in the running, shutdown or suspended states

Oracle VM Server for x86

The list of Oracle VM Servers displayed to be added to the server pool have the following characteristics:

- Owned by the Oracle VM Manager

- Not associated with a server pool
- In a healthy state and not placed in maintenance mode
- Have no virtual machines in the running, shutdown or suspended states

To Add Virtualization Hosts to the Server Pool

1. Select the server pool and click **Add Global Zones** or **Add Oracle VM Servers** in the Actions pane.

The Add Virtualization Host to Server Pool Wizard is displayed.

2. Select one or more virtualization hosts to be added to the server pool.

When the server pool is associated with the default network domain, you must proceed to configure the interfaces. Otherwise, you must associate the network domain with the selected assets.

3. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 6. For Oracle VM Server for x86, you cannot bond the interfaces and go to Step 6.

4. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

Click **Next** to configure the bonding parameters.

5. In the Configure Bonding step, specify the following parameters for Link Aggregation:

- Load balancing policy.
- LACP mode. If the Ethernet switch to which the physical interface connects supports aggregation, then specify the LACP mode.
- LACP timer.
- MAC address policy and provide the MAC address.

Click **Next** to define the network connection.

6. For each network connection, specify the following details:

- Specify the NIC and IP address for each network connection. When a selected server is connected to the network, no rows are displayed for that server.

You can assign the same NIC to different network connection when the networks have different VLAN IDs.

- You can select System Allocated for the NIC and IP address to be automatically allocated by the system.

Click **Next** to view the summary.

7. Review the information provided and click **Finish** to add the selected assets to the server pool.

Associating Network Domains

All networks in a network domain that is associated with a server pool are available to the members of the server pool. To associate a user-defined network domain with a server pool, there must be a connection between the virtualization servers and the

fabrics in the network domain. The fabric connection is not required for the default network domain.

The default network domain includes all the networks in Oracle Enterprise Manager Ops Center. When you associate a server pool with the default network domain, all the networks in Oracle Enterprise Manager Ops Center are available for attaching with the server pool.

For a user-defined network domain, you must define the physical interfaces of each virtualization server in the server pool to connect to each fabric in the network domain. You can also bond the interfaces of the virtualization servers.

Note: You cannot bond the interfaces of Oracle VM Server for x86 servers.

To Associate the Network Domain to a Server Pool

1. Select the server pool.
2. Click **Associate Network Domain** in the Actions pane.
The Associate Network Domain Wizard is displayed.
3. Select the network domain from the list.
4. For a user-defined network domain, select the physical interfaces of each selected servers to connect to each fabric in the network domain. When you do not want to bond the interfaces, go to Step 7. For Oracle VM Server for x86, you cannot bond the interfaces and go to Step 7.

For default network domain, you are taken to the Summary step.

5. You can bond the interfaces of the servers into a single logical link. The aggregation is done according to the standard IEEE 802.3ad Link Aggregation. Select Bond Interface in the Physical Interface column. A Bond ID is provided. Select the physical interfaces for the aggregation.

Click **Next** to configure the bonding.

6. In the Configure Bonding step, specify the following parameters for Link Aggregation:
 - Load balancing policy.
 - LACP mode. If the Ethernet switch to which the physical interface connects to supports aggregation, then specify the LACP mode.
 - LACP timer.
 - MAC address policy and provide the MAC address.

Click **Next** to view the summary.

7. Review the information and click **Finish** to associate the network domain with the server pool.

When you associate with a user-defined network domain, only the assigned networks in the network domain are available for attaching with the server pool. Therefore, ensure that you assign the required networks to the domain to attach to the server pool.

Attaching Networks

After the network domain is associated with a server pool, you can attach the networks in the network domain. Depending on the type of virtualization, the multiple connection to a network is allowed.

For Oracle VM Server for SPARC, you can make multiple connections to a network. Refer to the [Chapter 16, "Oracle VM Server for x86"](#) for more information.

For Oracle Solaris Zones, you can make multiple connections to a network when the server pool contains only Oracle Solaris 11 OS. For Oracle Solaris 10 OS, you can make only one connection. Also, when the server pool has mixture of Oracle Solaris 10 and 11 OS, then also you can make only single connection.

To Attach Networks to Server Pool

1. Select the server pool.

2. Click **Attach Networks** in the Action pane.

The Attach Network Wizard is displayed. All the networks assigned to the network domains are listed. For default network domain, all the networks are listed. The table also displays existing number of connections to the server pool.

3. Select one or more networks to attach to the server pool.

Click **Next** to specify the number of connections.

4. Enter the total number of connections required as per the virtualization type:

- For Oracle VM Server for SPARC, you can enter multiple connections. For each network connection, a virtual switch is created. You must physical interface for each virtual switch. You assign these virtual switches to the logical domains.
- For Oracle VM Server for x86, you can enter only one connection per network. You cannot make multiple connections.
- For Oracle Solaris Zones, the Oracle Solaris OS version plays an important role. For Oracle Solaris 10 OS, the network is always deployed in shared IP mode. You cannot make multiple connections. For Oracle Solaris 11 OS, the network is always deployed in exclusive IP mode. You can make multiple connections to the network. For each network connection, a VNIC is created when you boot the zone. For a server pool of mixture of Oracle Solaris 10 and 11 OS, the network connection is limited to single connection.

Note: The number of connections do not limit the number of guests to be connected to be network.

Click **Next** to configure the interfaces.

5. When one of the virtualization hosts are connected to the network, the corresponding rows for network configuration for that hosts are not displayed. If not, specify the information for each connection:

- Specify the NIC and IP address for each connection.

You can assign the same NIC to different network connection when the VLAN IDS are different for the networks.

- Use System Allocated for the system to take care of the NIC and IP address allocation.

- Use Do Not Allocate IP which is available for zones and Oracle VM Server for SPARC. You can connect to the network without specifying the IP address.

Click **Next** to view the summary of the information.

6. Review the information provided and click **Finish** to attach the networks to the server pool.

Ensure that you have provided correct information. Otherwise, the network attach job fails.

Associating Libraries

[Table 17-3](#) lists the types of storage libraries that you can associate with the server pool.

Table 17-3 Supported Libraries

	Oracle VM Storage Repositories	NAS Libraries	SAN Fibre Channel Libraries	SAN iSCSI Libraries
Oracle VM Server for SPARC	No	Yes	Yes	Yes
Oracle VM Server for x86	Yes	Yes	Yes	Yes
Oracle Solaris Zones	No	Yes	Yes	Yes

When you associate libraries with the server pool, only the libraries that are reachable from the virtualization hosts in the server pool are listed. Depending on the type of library, you either provide virtual disks or LUNs to the guests' storage.

To Associate a Storage Library

1. Select the server pool.
2. Click **Associate Libraries** from the Actions pane.
The Associate Library window is displayed.
3. The storage libraries that are accessible from the virtualization hosts in the server pool are listed. Select the libraries that you want to associate with the server pool. You can select one or more storage libraries from the list.
4. Click **Associate Libraries** to associate the selected libraries with the server pool.

Creating Guests

The guests refer to the logical domains, zones or the virtual machines that can be created in the virtualization hosts. According to the virtualization type of the server pool, you have the following options:

- Oracle VM Server for SPARC – Create Logical Domains
- Oracle Solaris Zones – Create Zones
- Oracle VM Server for x86 – Create Virtual Machines

These options trigger the deployment plans for the corresponding guests. Refer to the following chapters to refer to the profile and plan creation for the guests:

- [Chapter 14, "Oracle Solaris Zones"](#)
- [Chapter 15, "Oracle VM Server for SPARC"](#)
- [Chapter 16, "Oracle VM Server for x86"](#)

Migrating Multiple Guests

When you place a virtualization host in a server pool, the option to migrate one or more guests is enabled. The following options become available for the virtualization hosts:

- Oracle Solaris Zones – [Migrating Zones](#)
- Oracle VM Server for SPARC - [Migrating Logical Domains](#)
- Oracle VM Server for x86 – [Migrating Virtual Machines](#)

Migrating Zones

You can migrate multiple zones from a global zone which is in a server pool. Migrate the zones to an individual global zone or other zones server pool. When you migrate zones, the source and the target global zone must be compatible. The target global zones must have the following characteristics:

- Running at least Oracle Solaris 10 8/07 OS.
- Can access all the storage libraries associated with the zone.

Apart from compatibility, the target and the source global zone might have differences in the patches and packages installed on them. Choose to update the zone to match the patches and packages of the target global zone.

Note: You cannot downgrade the patches and packages of the zone. The migration fails in such scenario.

You can also force the migration of the zone without updating the patches and packages. Select the update options while migrating the zone.

To Migrate Multiple Zones

1. Select the global zone from which you want to migrate the zones.

2. Click **Migrate Zones** in the Actions pane.

The Migrate Zones Wizard is displayed.

3. The list includes the zones running in the global zone. Select one or more zones from the list.

Click **Next**.

4. Select an individual global zone or server pool to which you can migrate the zones.

The table displays the list of eligible global zones and server pool to which you can migrate the zones. The target global zone in the server pool depends on the server pool placement policy.

5. Select an update option to continue with migration.

The source and the target global zones might not be in the same patch level. Either you can select to update the patches and packages of zone to match the target global zone or continue migration without updating the zone.

6. Review the details and click **Finish** to migrate the zones.

Migrating Logical Domains

Migrate multiple logical domains from an Oracle VM Server to another Oracle VM Server in the same server pool. You cannot do cross server pool migration.

Only from Oracle VM Server for SPARC 2.1 version, live migration of logical domains is supported. For earlier releases, it is warm migration. The logical domains are shut down and then restarted on the target server.

To Migrate a Logical Domain

1. Select the Oracle VM Server in a server pool.
2. Click **Migrate Logical Domains** in the Actions pane.
3. The list of logical domains running in the Oracle VM Server are listed. Select one or more logical domains from the list.
4. Select an Oracle VM Server from the list.

The table lists the Oracle VM Servers that have enough resources to host the logical domains in the same server pool. The Oracle VM Server are listed in the decreasing order of preference.

5. Review the information and click **Finish** to migrate the logical domains.

Migrating Virtual Machines

For Oracle VM Server for x86 systems, you can migrate virtual machines only within a server pool. Live migration of virtual machines is supported. The eligible Oracle VM Servers have the following characteristics:

- Owned by the Oracle VM Manager.
- Placed in the same server pool as the source Oracle VM Server.
- Have required resources to host the virtual machines.
- Identical in machine make and model of the source Oracle VM Server.

To Migrate a Virtual Machine

1. Select the Oracle VM Server in the server pool.
2. Click **Migrate Virtual Machines** in the Actions pane.
3. Select the virtual machines that you want to migrate.
4. Select the Oracle VM Server to which you want to migrate the virtual machines.

The table lists the Oracle VM Servers that are eligible to host the virtual machines.

5. Review the summary and click **Finish** to migrate the virtual machines.

Balancing Resources

When you have selected to balance server pool resource manually, use the Balance Resources action to check for the under-utilization or over-utilization of resources. Select a server pool and click **Balance Resources** in the Actions pane to display the

current utilization of the virtualization hosts in the server pool. When balancing is required according to the placement policy, Oracle Enterprise Manager Ops Center displays a list of target server that can accept migrated guests. Click **Balance Resources** to start the migration job.

Deleting Server Pool

Use the option **Delete Server Pool** to release the virtualization hosts back to stand-alone state. The type of server pool imposes some limitations to delete a server pool.

Oracle VM Server for SPARC

When you want to delete an Oracle VM Server for SPARC server pool, shut down the logical domains running in the server pool. Only when the logical domains are detached, you can delete the server pool.

Oracle Solaris Zones

For a zones server pool, you can delete the server pool and the zones continue to be attached to the global zone. Oracle Enterprise Manager Ops Center does not provide any restrictions to shut down the zones before deleting a zones server pool.

Oracle VM Server for x86

Shut down all the virtual machines and remove all Oracle VM Servers except for the Master Oracle VM Server before deleting the server pool. The virtual machine metadata and virtual disks are stored in the storage library. The virtual machine details are not lost and you can start the virtual machine.

Related Resources

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

- [Chapter 16, "Oracle VM Server for x86"](#)
- [Chapter 15, "Oracle VM Server for SPARC"](#)
- [Chapter 14, "Oracle Solaris Zones"](#)
- [Chapter 7, "Networks"](#)
- [Chapter 2, "Asset Management"](#)
- [Chapter 6, "Storage Libraries"](#)
- [Chapter 18, "Virtual Datacenters"](#)
- *Oracle Enterprise Manager Ops Center Administration Guide*

For end-to-end examples, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

Virtual Datacenters

This chapter includes the following information:

- [Introduction to Cloud Management](#)
- [Oracle Engineered Systems](#)
- [Overview of Virtual Datacenter](#)
- [Roles for Managing Virtual Datacenter](#)
- [Actions Available in vDC Management](#)
- [Virtual Datacenter Creation](#)
- [vDC Management](#)
- [Create and Manage Cloud Users](#)
- [Account Creation](#)
- [Creating an Account](#)
- [Account Management](#)
- [Create and Manage vServer Types](#)
- [Overview of Cloud Users](#)
- [Roles for Cloud User Tasks](#)
- [Actions Available for a Cloud User](#)
- [View Account Quotas and Virtual Resources](#)
- [Creating and Managing vServers](#)
- [Creating vNets](#)
- [Managing vNets](#)
- [Managing Server Templates](#)
- [Volume Creation](#)
- [Creating Snapshots](#)
- [Creating Distribution Groups](#)
- [Related Resources](#)

Introduction to Cloud Management

Oracle Enterprise Manager Ops Center provides the platform to consolidate the physical resources in your data center, and to build and operate your cloud services. The physical resources such as virtualization servers, storage, and network are pooled which are accessed by users to build their applications.

Oracle Enterprise Manager Ops Center provides comprehensive management solution to deploy, configure, and manage the virtualization servers, storage resources, and network fabrics. Also, pool the virtualization resources that share the storage and network resources. All these features are leveraged and can be deployed as Infrastructure-as-a-Service (IaaS) cloud platform service model in Oracle Enterprise Manager Ops Center.

In Oracle Enterprise Manager Ops Center, the provision to setup the IaaS cloud platform service is available as vDC Management. Virtual Datacenter or vDC is a collection of the server pools that share the common storage and network resources. A cloud administrator sets up the infrastructure and provide access to cloud users. The cloud users use the allocated resources to create guest and install an OS of different flavours, deploy applications, monitor, and manage the applications. The cloud user is provided the access to Oracle Enterprise Manager Ops Center UI to view and manage their applications.

Oracle Enterprise Manager exposes APIs and command-line interface (CLI) to enable the access to a subset of the Virtual Datacenter functionality. The cloud user can also manage the allocated resources programmatically by calling the IaaS web services directly.

The features accessible through the Cloud Infrastructure API and CLI are available for a Cloud User. The features that are restricted to the Cloud Administrator, such as infrastructure configuration and setup are not available through the Cloud Infrastructure API and CLI except for listing the accessible accounts and to create key based access to cloud users who have access to the accounts already. All the functions that are available for a Cloud User are also available for the Cloud Administrator.

When you log in to Oracle Enterprise Manager Ops Center, the views and access for different sections of the UI are different for a cloud administrator and a cloud user.

Cloud User

If you are a cloud user looking for more information about creating and managing virtual servers, then go to the section [Overview of Cloud Users](#).

See *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide* for more information about using Cloud Web Service and Cloud Infrastructure API and CLI.

Cloud Administrator

If you are a cloud administrator, continue with the following sections that describe how to create and manage the virtual datacenters in Oracle Enterprise Manager Ops Center.

Oracle Engineered Systems

If you want to manage virtual datacenters on engineered systems such as Oracle Exalogic Elastic Cloud, then refer to [Chapter 20, "Oracle Engineered Systems"](#) for more information.

Overview of Virtual Datacenter

In Oracle Enterprise Manager Ops Center, consolidate the virtualization servers, storage, and network resources, and enable those resources to be utilized optimally and securely for mixed and dynamic workloads. This optimization of your resources is done by creating a Virtual Datacenter (vDC) in Oracle Enterprise Manager Ops Center.

The physical resources allocated of a vDC are entitled to accounts as virtual resources under quotas set by the cloud administrator. Accounts provide the required capabilities to manage the allocated resources. An account entitles designated cloud users the right to use its virtual computing, network, and storage resources.

A cloud user with access to different accounts can create virtual servers, known as vServers, and host or access applications.

A cloud administrator is involved in setting up the infrastructure for the vDC, creating and managing vDCs, creating and managing accounts, and managing access to cloud users.

Roles for Managing Virtual Datacenter

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

Table 18–1 Virtual Datacenter Roles and Permissions

Task	Role
Create and Manage Virtual Datacenters	Cloud Admin
Create and Manage Accounts	Cloud Admin
Create and Manage vServer Types	Cloud Admin
Add and Manage Cloud Users	Cloud Admin

Required User Role

A cloud administrator requires the Cloud Admin role to create and manage the virtual datacenters, accounts, and cloud users. The Cloud Admin role has the necessary permissions of the Asset Admin and Network Admin roles to setup and manage the virtual datacenter infrastructure.

Specify vDC Privileges

By default, the addition of the Cloud admin role to a user doesn't give full vDC privileges. To access an existing vDC as a cloud administrator, you must have the proper privileges set to view and manage it.

Contact your Ops Center administrator to manage your user's role to specify proper vDC privileges. See *Oracle Enterprise Manager Ops Center Administration Guide* for more information about how to manage user's roles and privileges.

Actions Available in vDC Management

The virtual datacenter configuration involves the setup of the infrastructure, managing the resource workloads, and understanding cloud user requirements. The user with the cloud administrator role has the following actions:

- Create and manage virtual datacenters
- Create and manage accounts
- Manage access to accounts
- Manage physical resources of virtual datacenter
- Create and manage vServer types

Virtual Datacenter Creation

Virtual datacenter (vDC) is a consolidation of your physical resources that share the network and storage resources. The physical resources that form a virtual datacenter are:

- **Homogenous server pools**

The server pools are groups of virtualization supported servers that share compatible chip architecture. The supported virtualization types are:

- Oracle Solaris Zones
- Oracle VM Server for x86
- Oracle VM Server for SPARC

Server pools are of single virtualization technology. For example, you can have a server pool of Oracle VM Server for SPARC servers only and not a mixture of Oracle VM Server for x86 and Oracle VM Server for SPARC servers. You cannot have server pools that have mixture of different virtualization technology. A vDC is based on any one of the supported virtualization technology. This infrastructure difference brings in some major differences in using the vDC. The cloud user might not be aware of the background infrastructure but the cloud administrator ensures that the vDC resources are always highly available to a cloud user. See [Server Pool Setup](#) for more information about the server pool setup required for virtual datacenter.

- **Storage**

The storage is inherited from the server pools. The storage is used for storing rootdisk of the virtual server created in the accounts, virtual server configuration data, templates, ISO images, FLAR images, and also used as volumes. While creating a vDC, you can allocate some storage resources that are used only for volumes. Volumes are used to attach to the vServers. See [Storage Setup](#) for more information about the storage resources required for virtual datacenter.

- **Networks**

Plan the network requirements and attach the network to the server pools. vDC inherits the network resources from the server pools. See [Network Setup](#) for more information about setting up your network resources for virtual datacenter.

Server Pool Setup

You can create vDC with server pools based on the following types of virtualization technology:

- Oracle VM Server for x86
- Oracle Solaris Zones
- Oracle VM Server for SPARC

See [Chapter 17, "Server Pools"](#) for more information about how to create server pools.

Ensure all the required networks are assigned to the user-defined network domain that is associated with the server pool. The vDC creation wizard displays the list of available server pools in Oracle Enterprise Manager Ops Center. The list includes the server pools that conform to the following characteristics:

- The server pool is not empty and there are virtualization servers placed in the server pool.
- The server pool is associated with a user-defined network domain.
- Server pools of supported virtualization technology.
- vServers have compatible storage resources to be used as root disk and volumes.
- Compatible storage resources for saving vServer metadata.

For Oracle Solaris Zones server pool, the pool members must not be a mixture of Oracle Solaris 10 and 11 operating system.

Storage Setup

The vDC inherits the storage resources allocated for the server pool. The virtualization type of the server pool defines all the possible types of storage resources available for the vDC. Oracle VM Server for x86 based server pools can have the following type of storage resources:

- Oracle VM Storage Repositories
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

For Oracle VM Server for x86 server pools, the Oracle VM Storage repositories must be associated with the server pool.

The Oracle Solaris Zones server pools can have the following types of storage resources:

- NAS libraries
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

The Oracle VM Server for SPARC server pools can have the following types of storage:

- NAS libraries
- LUNs allocated from Static Block Storage libraries
- LUNs from Dynamic Block Storage libraries

You require the Oracle Solaris Zones and Oracle VM Server for SPARC server pool to be associated with NAS libraries if they are associated only with Block Storage libraries. NAS libraries are required for storing guest metadata details.

Volume Storage

When you create a vDC, you allocate the storage resources to be used for all the accounts created in the vDC. The allocated storage resources are used as root disks of the virtual servers (vServers) and volumes. Volumes are additional storages that you can attach to the vServers. For volumes, the storage is allocated from the storage libraries that are associated with the server pools. A cloud user can also import external volumes into the account.

Network Setup

Oracle Enterprise Manager Ops Center provides comprehensive network management options that enables secure management of the virtual datacenter.

The server pools must be associated with user-defined network domain to be added to the vDC.

Understanding Network Domain Implementation

The network domain is a logical grouping of related networking components in Oracle Enterprise Manager Ops Center. The networks are built on the following type of fabrics:

- **Fully-managed switched fabrics**

The switches and the fabrics they support are discovered and managed in Oracle Enterprise Manager Ops Center and you can create VLANs or InfiniBand partitions. The network domain created from fully-managed fabrics provides the option to create dynamic private networks.

- **Host-managed fabrics**

Set VLAN IDs to unmanaged Ethernet fabrics and make it host-managed fabrics. Though the switches are not managed in Oracle Enterprise Manager Ops Center, you can assign VLAN IDs to the Ethernet fabrics. You must enable the VLAN IDs manually on the switch ports connected to the hosts in that fabric. Host-managed fabrics allow you to create dynamic private networks.

Oracle Enterprise Manager Ops Center uses the VLAN IDs set for the creation of dynamic private networks.

- **Unmanaged fabrics**

Neither the switches are managed nor the VLAN IDs are available in Oracle Enterprise Manager Ops Center to create the network. The networks are declared or discovered while discovering an asset in Oracle Enterprise Manager Ops Center. The network domain created out of these networks does not provide the option to create dynamic private networks. You can set aside the networks created on the unmanaged fabrics as private while creating a network domain.

You must ensure that these networks assigned as private are not in use and there is no routing between the selected networks and other networks. These private networks are defined as static private networks. For vDC, a network domain with unmanaged fabrics can be assigned to a server pool only if there are static private networks available in it.

Dynamic Private Networks

Each dynamic private network created is allocated a subnet address according to the standards specified in IETF RFC1918. While creating network domains, you can select the option to select the fabric on which you want to create the network, limit the number of networks to be created on the fabric, and the network addresses to be excluded from use.

The number of private networks that you can create on a fully-managed fabric depends on the type of physical fabric. Each port on an Ethernet switch can support 128 logical fabrics. Each partition on an InfiniBand switch can support 32000 logical fabrics. The number of logical fabrics determines how many private networks you can create. For an InfiniBand fabric, each P-key can support a private network. For an Ethernet fabric, each VLAN ID can support a private network.

When you create a network domain, you can set the network creation limit for each selected fabric. This determines how many private vNets that can be created in an account. You can spread the number of private networks across all of the accounts in the vDC.

Plan Your Network Domains

The network setup must be implemented in such a way that the cloud user has the required network resources available to allocate to the virtual servers created in the accounts. The cloud user must be able to create private networks. The networks must be created upon either fully-managed switched fabrics, host-managed fabrics or unmanaged fabrics.

Create a network domain, assign the managed networks, and associate it with the server pool which must be added to the vDC. The vDC inherits the network domain from the server pool. From the available networks in the network domain, attach the required networks to the server pool. These managed networks form the public external networks for the vDC. You can assign these networks to the accounts created in vDC.

The vDC enables cloud users to create virtual machines and run applications. The cloud user provides networking to its virtual machines from the public networks that are allocated to each account or by creating private networks. The cloud user can use the public networks, or create a private vNet for use in the account. When the cloud users create a private vNet, either a dynamic private network is created or the static private network is available for use in that account.

You must allocate the IP addresses in the public networks and static private networks. This enables the cloud users to use the allocated IP addresses.

Note: All the server pools in the vDC must be associated with the same user-defined network domain.

When the cloud user creates a private vNet, it is listed under the network domain of the vDC.

Avoid to delete any EoIB network resource when a cloud user has just deleted a vServer. When you delete a EoIB network resource, wait for approximately 5 minutes for the VLAN maps to be refreshed automatically.

See [Chapter 7, "Networks"](#) for more information about managing fabrics, creating network domains, and private networks.

CPU Oversubscription

You can oversubscribe the CPU resources allocated to a vDC for an increased utilization of the resource. You define two parameters that define the CPU oversubscription:

- The ratio of number of virtual CPUs (vCPUs) to physical CPU Threads.
For example, if the vCPU to physical CPU threads is set to 2, each virtual CPU receives at least 50% of the cycles of a physical CPU thread.

Note: As the CPU oversubscription ration increases, the performance might be affected, but the utilization of the CPU resources improve. The CPU oversubscription ratio that you might want to use is at most 3:1. At extremely high ratios, the risk of instability of the system increases.

- CPU cap that defines the maximum share of the physical CPU thread's cycle that can be allocated to a vCPU.

For example, if the CPU Cap is set to 50, a vCPU can use up to 50% of the cycles of a physical CPU thread.

Provide appropriate CPU cap for the ratio of vCPU to physical CPU thread to get accountable vServer performance and balanced system.

This is applicable only for vDCs based on Oracle Solaris Zones and Oracle VM Server for x86 virtualization technologies. For Oracle VM Server for SPARC, this is not applicable as there is a one to one relationship between vCPU and physical CPU thread.

Creating a Virtual Datacenter

Before you create the vDC, you must have your server pools, network, and storage set up. When they are set up properly, you can select the required resources for the vDC.

To Create a Virtual Datacenter

1. Select vDC Management in the Navigation pane.
2. Click **Create Virtual Datacenter** in the Actions pane.
The Create Virtual Datacenter Wizard is displayed.
3. The first step provides a introduction to the vDC and the prerequisites for creating a vDC.

You can select to skip this step in the future when you create a vDC again. Click **Next** to specify identification details for the vDC.

4. In the Specify Virtual Datacenter Details, enter the following information:
 - Provide a name and description for the vDC.
 - Enter tags for better identification and classification of the vDC in Oracle Enterprise Manager Ops Center.

Click **Next** to select the server pools.

5. Select a server pool from the list.
6. Select whether you want to add more compatible server pools as that of the selected server pool.
7. Select the compatible server pools that you want to assign to the vDC.

Click **Next** to configure the volume storage.

8. Designate some storage resources to be used for volumes. Designate the following storage resources as volume:
 - Oracle VM Server for x86 based vDC
 - Oracle VM Storage Repositories

- Dynamic Block Storage Libraries
- Oracle Solaris Zone based vDC
 - NAS Libraries
 - Block Storage Libraries that include both Dynamic and Static
- Oracle VM Server for SPARC based vDC
 - NAS Libraries
 - Block Storage Libraries that include both Dynamic and Static

Click **Next** to specify the vCPU sizing.

9. Define the following values for vCPU sizing:

- Enter the ratio of number of vCPUs to a physical CPU thread. The value must be greater or equal to 1.0. Decimal values are also supported. The total number of vCPUs can be higher than the number of physical CPU threads. You can run more vCPUs than the existing physical CPU threads through timesharing of the CPU cycles.
- Specify the CPU cap which is the maximum share of the physical CPU thread's cycle that can be allocated to a vCPU.

The following information is displayed in the vCPU sizing screen:

- **Total number of vCPUs:** The total number of vCPU in the vDC for the updated vCPU to physical CPU Threads ratio.
- **Avg memory per vCPU:** The average memory per vCPU in GB. The total memory available for the vDC by the total number of vCPUs. When you want to use the CPU and memory resources to the full extent, then the vServers must be using this amount of resources.

When vServers are created in a zone based vDC, the Fair Share Scheduler is the default scheduler. The physical CPU allocation for each vServer is provided by the Fair Share Scheduler.

This is not applicable for vDCs based on Oracle VM Server for SPARC virtualization technology.

Click **Next** to view the summary.

10. The summary lists the total resources that are allocated for the vDC. The summary lists the following:

- **vDC Name:** The name given to the vDC.
- **Virtualization type:** The virtualization technology of the server pool.
- **Total Physical CPU Threads:** The physical CPU threads available from all the virtualization servers in the selected server pools for the vDC.
- **Total memory:** The total RAM allocated to the management domains of the selected server pool. For example, Oracle Solaris Zones based vDC, it is the total memory of the global zones in the server pool. For Oracle VM Server for x86 based server pools, it is the total memory of the Oracle VM Servers in which the virtual machines run.
- **Total disk space:** The sum of all the disk spaces allocated from the storage libraries associated with the selected server pools.
- **vCPU to Physical CPU Threads ratio:** The number of vCPUs to physical CPU Threads.

- **vCPU cap:** The maximum share of the physical CPU Thread's cycle that can be allocated to a vCPU.

Confirm the vDC configuration and click **Finish** to create the vDC.

vDC Management

Perform the following actions for managing the vDC:

- [Update vDC Configuration](#)
- [Resource Management](#)
- [Delete vDC](#)

Update vDC Configuration

Update the following details of vDC:

- vDC identification and tags.
- Assign more server pools.
- Assign more storage for volume. You cannot remove the storage assigned for volume.
- Modify the vCPU sizing.

Select a vDC and click the **Update Virtual Datacenter** option in the Actions pane. Modify the required attributes and complete the wizard.

Resource Management

Oracle Enterprise Manager provides a complete view of the resources in the vDC and monitors the utilization of the virtualization servers, network, and storage. The cloud administrator gets a good picture of the utilization of the resources which helps to manage the under-utilized and over-utilized resources.

The Dashboard shows the total resources available for a vDC and the amount allocated for different accounts. The committed resources indicate the total resources entitled to all the accounts in the vDC.

For example, if you have 4 physical threads in the vDC and set the vCPU to physical CPU threads ratio to 2, then you have 8 virtual CPUs. You can oversubscribe the vCPUs and allocate 5 vCPUs to one account and another 5 vCPUs to another account. The total committed resources to all the accounts in the vDC is 10 vCPUs.

The Network tab displays detailed network usage in the vDC. The page lists the fabric and public networks of the vDC, and all of the private networks created by the cloud user. The fabric on which the private vNets are created are also displayed in this page.

Figure 18–1 Network Usage in a vDC

vdc1 - Virtual Datacenter

Dashboard | **Networks** | Storage | Server Pools | Accounts | vServer Types | Incidents | Jobs

Network Configuration | **Network Allocation**

Fabrics:

Fabric Name	Fabric Media Type	Fully Managed
pfab at eth0 at 10.169.79.77		No

Public External Networks

Search

Network Name	Fabric Name	Network Address	Logical Fabric ID	Role
10.169.69.0/24		10.169.69.0		PUBLIC EXTERNAL

Application Internal Networks

Search

Network Name	Fabric Name	Network Address	Logical Fabric ID	Used by Account
vnetClt	lfab at null at null; fabTag:1000	192.168.0.0		account1

The Server Pools tab shows the usage of the server pool resources. This lists the server pools that are in the vDC. The total and committed resources of vCPU, memory and storage are displayed. This helps you to plan whether you must add more resources to the vDC.

Figure 18–2 vDC Server Pool Usage

myCloud - Virtual Datacenter

Dashboard | Networks | Storage | **Server Pools** | Accounts | vServer Types | Incidents | Monitoring

Server Pool Summary

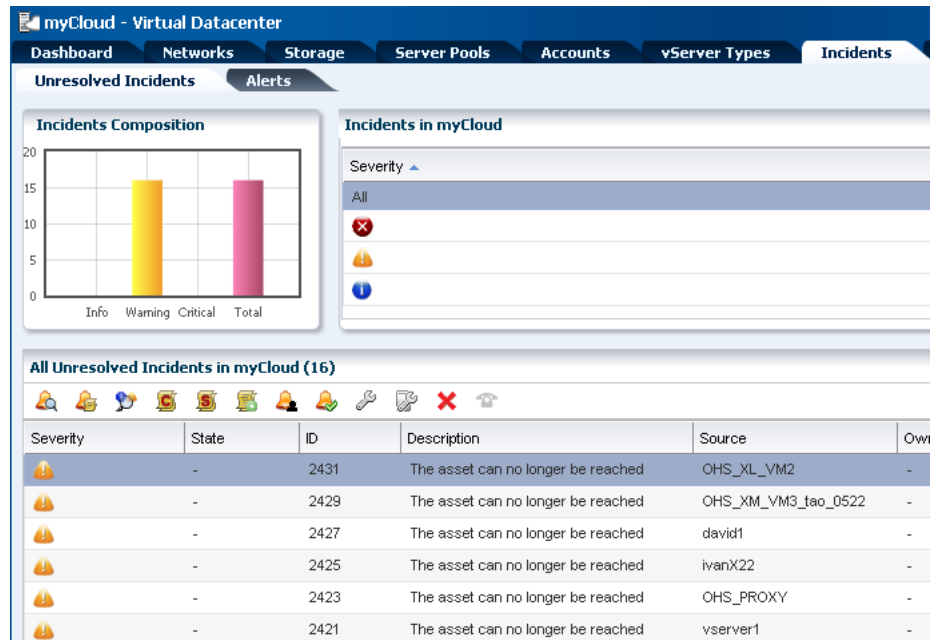
Pool Type: Oracle VM Server for x86
 Allocated vCPUs: 38

Total Physical CPU Threads: 24
 vCPU to CPU Thread ratio: 1 to 1

Server Pools

Server Pool	Servers	Total vCPUs	vServers	Average vCPU Utilization (%)
spool1	1	24	16	158

The Incident tab lists all the incidents reported from all the infrastructure components in the vDC. For each incident, view the source of the incident which helps you to identify the component and the source of the problem. The incidents also covers the warnings issued when the vDC resources are over-utilized.

Figure 18–3 vDC Incidents Tab

Delete vDC

You can delete a virtual datacenter and release all of the resources attached to it. To delete a virtual datacenter, ensure the following conditions:

- No vServers are in the running state in the virtual datacenter. Shut down all the vServers in the virtual datacenter.
- Delete all the accounts in the vDC.

When you delete a virtual datacenter, all of the suspended or shut down vServers are deleted and the resources are released.

Create and Manage Cloud Users

You can create cloud users and provide access to the accounts so that they can utilize the computing resources allocated to them. The cloud user is associated with an existing user on the OS under the Enterprise Controller. The ways to create this operating system user varies by the type of OS and type of name services, for example, file based, NIS, or LDAP.

Use the following Oracle Solaris OS command for a simple local file based user account.

```
useradd <cloud user name>
passwd <cloud user name>
```

Adding and Removing Cloud Users to an Account

You can add cloud users to the account and give them the right to use the resources allocated to an account. You can add the cloud users when you create an account or using the option **Add Users**.

To Add a Cloud User

1. Select the account and click **Add Users** in the Actions pane.

The Add Users Wizard is displayed.

2. The lists of available cloud users and cloud users assigned to the account are displayed. Select users you want to assign to the account and move them to the Assigned Users.

Click **Next** to view the summary.

3. View the summary and click **Finish** to add the user to the account.

To Remove a Cloud User

1. Select the account.
2. Select the **Users** tab in the center pane.

The list of users assigned to the account are displayed.

3. Select the cloud user whom you want to remove.
4. Click the **Delete** icon.
5. Click **Remove** to confirm the action.

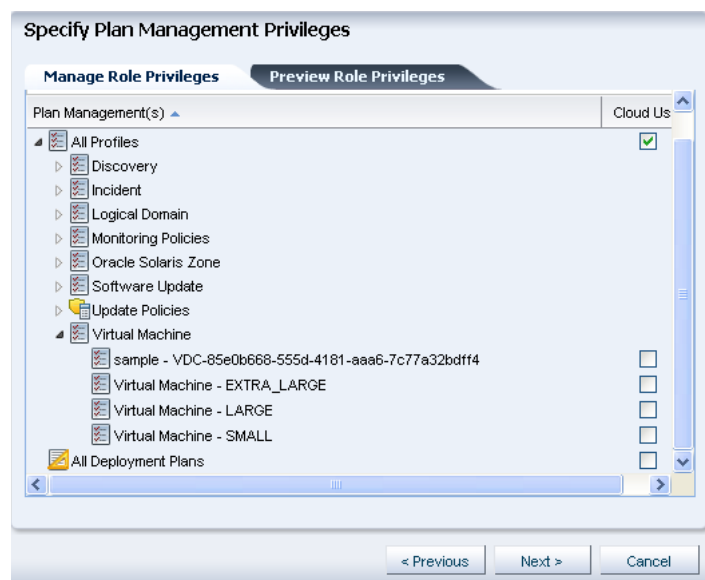
Manage Cloud User Roles

By default, the cloud user is provided with default privileges to use all the profiles, plans, and vServer Types available in the selected vDC. Ensure that the cloud user has appropriate privileges to the vServer types. Otherwise, the users cannot create vServers in their environment.

Refer to the *Oracle Enterprise Manager Ops Center Administration Guide* for more information about managing user roles.

Figure 18–4 shows the step to set the privileges for a cloud user.

Figure 18–4 Specify Privileges



Account Creation

An account entitles designated cloud users the right to use computing, network, and storage resources of vDC. The account provides the required capabilities to manage these resources.

The prerequisites for creating an account are:

- Estimate the resource quotas to be allocated for the account
- Identify the cloud users to be assigned to the account

The quota for vCPU, memory and storage resources are defined during account creation. The Resource Quota Information display in the account wizard creation indicates how much of the corresponding vDC resources are subscribed. The resource usage indicates whether the vDC resources are oversubscribed or undersubscribed.

You can create a maximum of 4096 private vNets in an account. You can set the limit of number of private vNets that can be created in an account.

During account creation, the public networks that are available in the vDC are listed. You can set the number of public IP addresses allocated to the account from the resource. A cloud user can assign the public IP address to a vServer, as needed.

Assign the cloud users to the account during account creation or separately. Cloud users have access to only specific accounts. As a cloud administrator, you can manage the access of the cloud users to all the accounts.

Creating an Account

You provide an entitlement to the virtual resources for an account. You allocate the resources from the vDC to an account. The resource allocation for all the accounts in a vDC can be more than the actual resources in a vDC. This oversubscription of the resources must be identified and planned for a vDC. You must configure the virtual resources for an account properly and update the resource configuration when the requirement increases.

As a cloud administrator watch the resource usage and properly configure the resources for an account.

To Create an Account

1. Select the vDC in which you want to create the account.
2. Click **Create Accounts** in the Actions pane.

The Create Account Wizard displays.

3. The first step provides a introduction to accounts and the prerequisites for creating an account.

You can skip this step in the future when you create an account again. Click **Next** to specify identification details for the account.

4. Enter the name and description for the account.

Add tags for categorizing and identifying the account. Click **Next**.

5. Specify the resource quotas for vCPU, memory and storage to the account.

The available resources from vDC are allocated to the vDC. You can oversubscribe the resources. The Resource Quota Information displays whether the vDC resources are undersubscribed or oversubscribed. This gives a complete picture of the vDC resource usage.

6. Select the number of private vNets a cloud user can create in an account.
The maximum number of private vNets that you can create in an account is 4096.
7. Select the public networks from the vDC that you can allocate to the account.
For each network, the available number of IP addresses is displayed. Enter the number of IP addresses to be allocated to the account. Click **Next**.
8. Select the cloud users who can access the account. You must assign a cloud user while creating the account. Thereafter, use **Add Users** option to add more users.
9. Review the account information and click **Finish** to create the account.

Account Management

You have the following options to manage the account configuration in a vDC:

- Update account resource configuration
- Assign cloud users to accounts
- Delete an account

Update Account Configuration

As a cloud administrator, you can modify the identification, and resources allocated to an account. When you want to allocate more resources to an account, you can modify by selecting an account and click **Update Account** in the Actions pane. Modify the resources and complete the wizard.

You can remove the resources allocated to an account if they are not used by the vServers. Similarly, you can reduce the limit of public IP addresses if they are unused by vServers.

Account Resource Management

When an account does not have enough resources, the cloud user receives notifications that they cannot create vServers due to unavailability of resources. Also, when there are no physical resources available from the corresponding vDC, the cloud user cannot create vServers even if the account resources show that they are not fully used up.

For example, assuming that you have 100 GB physical storage allocated to a vDC and an account with 125 GB storage. The vDC storage is oversubscribed. The cloud users create vServers in the account and use the storage. When the storage usage exceeds 100 GB, the cloud user cannot create vServer as the vServer job fails with the message that enough space is not available.

Therefore, the cloud administrator must watch the resource usage and add more resources to the vDC.

Maintaining OS Images

For vDCs based on Oracle VM Server for SPARC virtualization technology, you must manage OS provisioning profiles that the cloud user can use for creating vServers in their environment. The Oracle Solaris 10 ISO or FLAR images must be imported in the NAS libraries associated with the server pools placed in the vDCs.

Configure your Oracle Solaris 11 Software Update Library in the Enterprise Controller to synchronize with the Oracle Solaris 11 Package Repository. The cloud user can use all the profiles for Oracle Solaris 11 OS.

Ensure to configure DHCP in your environment so that the cloud user can install Oracle Solaris 10 ISO image.

OS provisioning profiles that are based on Oracle Solaris 10 FLAR images and Oracle Solaris 11 OS, use WAN boot and do not require DHCP configuration.

See [Chapter 5, "Software Libraries"](#) for more information.

Delete an Account

Ensure that you do not have running virtual servers in the account. Use the option Delete Account to remove the account.

When there are vServers in running state, the following message is displayed:

Figure 18–5 Delete Account



Create and Manage vServer Types

vServer Type are profiles of virtual machines that defines the computing resources such as virtual CPU, memory, and storage size. A cloud user can use this to implement and create vServers.

Note: For vDCs based on Oracle VM Server for x86, the storage size is not defined in the vServer type. The storage size is defined from the server template of the vServer.

A vServer type created is available for all the accounts in a virtual datacenter.

As a cloud administrator, you can capture the user requirements and create different vServer Types. Cloud users can use this to create vServers. By default, there are three system-defined vServer types that define the computing resources.

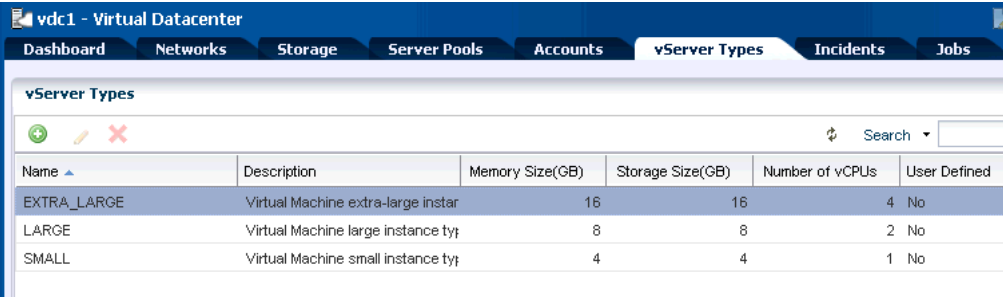
System-Defined vServer Types

The three system-defined vServer types provided for all vDCs are:

- Extra Large
- Large
- Small

The vServer type configuration varies depending on the virtualization technology on which the vDC is constructed.

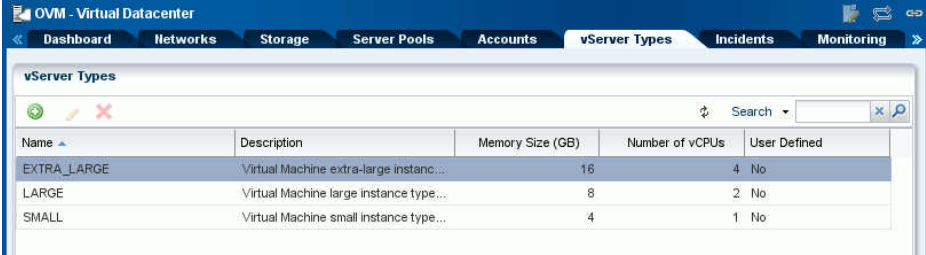
Figure 18–6 vServer Types



Name	Description	Memory Size(GB)	Storage Size(GB)	Number of vCPUs	User Defined
EXTRA_LARGE	Virtual Machine extra-large instar	16	16	4	No
LARGE	Virtual Machine large instance typ	8	8	2	No
SMALL	Virtual Machine small instance typ	4	4	1	No

For Oracle VM Server for x86 based vDCs, the vServer Type is displayed as shown in [Figure 18–7](#).

Figure 18–7 vServer Type for Oracle VM Server for x86 based vDCs



Name	Description	Memory Size (GB)	Number of vCPUs	User Defined
EXTRA_LARGE	Virtual Machine extra-large instanc...	16	4	No
LARGE	Virtual Machine large instance type...	8	2	No
SMALL	Virtual Machine small instance type...	4	1	No

Creating a vServer Type

When you create a vServer type, the VM hosting details display the following information in the wizard based on the resources defined:

- The number of virtualization servers in the vDC that have sufficient physical resources to host a vServer with the selected resources.
- An estimation of number of vServers that can be hosted with the total number of physical resources of the vDC.
- A warning when the current value of the memory size exceeds the selected storage size.

To Create a vServer Type

1. Select the vDC and click **Create vServer Type** in the Actions pane.

The Create vServer Type Wizard is displayed.

2. Enter the name and description for the vServer type.

Add tags for identification and classification of the vServer type. Click **Next**.

3. Specify the vCPU, memory and storage resources for the vServer type.

As you enter the values, the VM Hosting displays the following information:

- The number of virtualization servers that have sufficient resources in the vDC to host a vServer of this type.

- The number of vServers that the vDC can host with the total number of physical resources of the vDC.

A warning message is displayed when the current memory size is more than the storage size. The existing vServer types are displayed with the configuration details. This helps to avoid creating duplicate vServer types.

For vDCs based on Oracle VM Server for x86, the storage size is not defined in the vServer Type.

Provide the values for the resources and click **Next**.

4. Review the information provided and click **Finish** to create the vServer type.

Update vServer Types

Modify the following details of a vServer type:

- Name and description
- Modify, or add new tags
- Modify the resource configuration

Use the option **Update vServer Type** to launch the Update vServer Type Wizard. Modify the configuration and complete the wizard.

Delete vServer Type

You cannot delete the system-defined vServer types. You can delete only the vServer types that you have created.

Overview of Cloud Users

A cloud user with access to an account is entitled to manage and use computing, network, and storage resources allocated in a vDC within the limits of the account quotas.

Cloud users can create and manage the life cycle of vServers for their applications. Creation and management of vServers involve the setup of virtual resources, the management of virtual resource workloads, and understanding application requirements.

Cloud Users can manage the following virtual resources:

- **Virtual Networks (vNets):** Used to connect and restrict network access of vServers.
- **Server Templates:** Designates the operating system and how it is installed while creating a vServer.
- **Virtual Storage:** Includes volumes that you can attach to vServers, and snapshots to capture the current state of a volume for different purposes.
- **vServers:** An entity that provides the outward interface of a stand-alone operating system. A vServer has its own identity, local storage, interfaces, and configuration that exist for the full lifetime of the vServer.

Availability and management of some virtual resources vary depending on the virtualization technology of the vDC or resource type supported by the vDC.

Oracle Enterprise Manager Ops Center offers to cloud users the option to perform their tasks using its browser interface or through:

- **APIs:** A Web service API and a Java API to programmatically manage the allocated resources in a vDC account for automation or integration purposes.
- **CLI:** To manage allocated resources in a vDC account from a text-based console that can be also used for automation or integration purposes.

The following sections in this guide describe the tasks a cloud user can perform in Oracle Enterprise Manager Ops Center using the UI. For more information about the use of APIs and CLI, refer to *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide*.

Roles for Cloud User Tasks

[Table 18–2](#) lists the tasks and the role required to complete the task. Contact your administrator if you do not have the necessary role or privilege to complete a task. See *Oracle Enterprise Manager Ops Center Administration Guide* for information about the different roles and the permissions they grant.

Table 18–2 Cloud User Tasks and Roles

Task	Role
View Account Quotas and Virtual Resources	Cloud User
Create and Manage vServers	Cloud User
Create and Manage Virtual Networks	Cloud User
Create and Manage Volumes	Cloud User
Create and Manage Snapshots	Cloud User
Create and Manage Server Templates	Cloud User

Actions Available for a Cloud User

A cloud user can perform different actions to use and manage the allocated resources in a vDC account. This section lists the actions a cloud user can perform for each type of resource and also provides information about the actions that are available for the different vDC account types.

- View account quotas and virtual resources
- Create and manage vServers
- Create and manage virtual networks
- Create and manage volumes
- Create and manage snapshots
- Create and manage server templates

Based on the virtualization technology on which the vDC and its accounts are created, there might be differences in the availability of the options in the cloud user view. The following tables provide a detailed list of options that are available for a cloud user when the vDC is based on different virtualization technology.

Account

Table 18–3 Account Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
View account quotas and resources	Yes	Yes	Yes
View account Incidents	Yes	Yes	Yes
View Account Jobs	Yes	Yes	Yes

vServers

Table 18–4 vServers Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create vServer	Yes	Yes	Yes
Stop and start vServer	Yes	Yes	Yes
Pause and Resume vServer	Yes	No	No
Update vServer	Yes	Yes	Yes
Launch Virtual Console	Yes	No	No
Shutdown All vServers	Yes	Yes	Yes
Attach and Detach Volume	Yes	Yes	Yes
Enable and Disable HA	Yes	No	No
Delete vServer	Yes	Yes	Yes

Networks

Table 18–5 Networks Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create private vNet	Yes	Yes	Yes
Update private vNet	Yes	Yes	Yes
Delete private vNet	Yes	Yes	Yes
Allocate vIP	Yes	Yes	Yes
Deallocate vIP	Yes	Yes	Yes
View vIPs	Yes	Yes	Yes

Volumes and Snapshots

Table 18–6 Volumes and Snapshot Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create Volume	Yes	Yes	Yes
Create Volume from Snapshot	Yes	No	No
Import Volume	Yes	No	No
Update Volume	Yes	Yes	Yes
Delete Volume	Yes	Yes	Yes
Create Snapshot	Yes	No	No

Server Templates

Table 18–7 Server Templates Option Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Upload Server Template	Yes	No	No
Create Server Template	No	Yes	No
Update Server Template	Yes	Yes	No
Delete Server Template	Yes	Yes	No
Register and Unregister Server Template	Yes	Yes	No
Save vServer as Template	Yes	No	No

Distribution Groups

Table 18–8 Distribution Groups Options Availability

Action	Oracle VM Server for x86 Based vDC	Oracle VM Server for SPARC Based vDC	Oracle Solaris Zones Based vDC
Create Distribution Group	Yes	Yes	No
Update Distribution Group	Yes	Yes	No
Delete Distribution Group	Yes	Yes	No

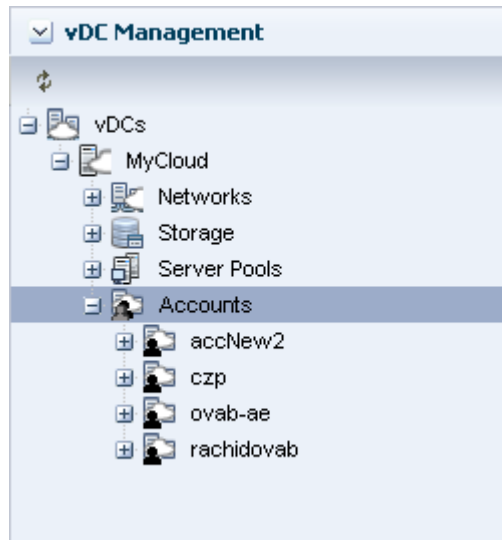
View Account Quotas and Virtual Resources

Account quotas limit the creation of new virtual resources in the account. A cloud user must be aware of the account quotas and manage the workloads of the virtual resources.

To see the quotas and virtual resources of an account, expand vDC Management in the Navigation pane.

This lists all the accounts to which a cloud user has access under vDC Accounts.

Figure 18–8 vDC Management



You can select a listed account to display general account information and quotas in the Dashboard tab of the center pane. The Dashboard tab displays general information about the selected account and a summary of the virtual resources status and usage.

Virtual resources details appear in the other tabs across the center pane:

- **Networks:** Shows the vNets quotas. Lists and displays information for each private vNet and public network available for the account. Also displays the actions bar for the actions that you can perform for a vNet.
- **Storage:** Shows the storage quotas. Lists and displays information for each vServer root disk, volume, and snapshot available for the account. Also displays the actions bar for the actions that you can perform for a storage resource.
- **vServers:** Shows the CPU quotas. Lists and displays information for each vServer available for the account. Also displays the actions bar for the actions that you can perform for a vServer.
- **Server Templates:** Lists and displays information for each server template available for the account. Also displays the actions bar for the actions that you can perform for a server template.

Creating and Managing vServers

A vServer is an entity that provides the outward interface of a stand-alone operating system that consumes CPU, storage, and memory resources. A vServer has its own identity, local storage, interfaces, and configuration that exist for the full lifetime of the vServer.

You determine the creation of new vServers according to the account quota limits and applications requirements.

Before You Begin

You must complete the following tasks as prerequisites for creating vServers:

- **Define the memory, storage, and number of vCPUs required for the vServers**

After memory, storage, and number of vCPUs are defined, you can select the best suitable vServer type from those available for the account. vServer types are visible to cloud users during the vServer creation process. Contact your cloud administrator when you require a vServer type.

- **Define the operating system and configuration for the vServers**

You must define the OS and configuration for the vServers. Based on that, you can choose the correct server template. The proper server template must exist before creating a vServer. You can either select a server template from those that exist for the account or create a new server template. Actions to use or create server templates vary based on the virtualization technology used in the vDC. For more information about server templates, see [Managing Server Templates](#)

- **Define the virtual networks to attach to vServers for connectivity**

Plan the network connectivity of the vServers. vServers are only assigned to virtual networks at vServer creation time. You must choose one or more vNets from the available vNets or create new ones before creating a vServer. For more information about vNets, see [Creating vNets](#)

- **Define the IP address assignment method**

Select from two different methods for IP address assignment to vServers:

- **Static method:** Used to assign a specific IP address to a vServer. When using this method you must allocate in advance an IP address from the selected virtual networks. You must consider using this method when creating a single vServer at a time.
- **Automatic method:** This method dynamically assign an IP address from each selected virtual network. When creating a multiple vServer at a time, only the use of the automatic IP address assignment method is allowed.

- **Define the number of vServers to create**

You can create single or multiple vServers at a time. When creating multiple vServers at a time, vServers are created with the same configuration and a suffix is added to each vServer's name. You must also consider the IP address assignment method when defining the number of vServers.

- **Define the use of the high availability support option for vServers**

Choose to use the high availability support option for vServers. When a high availability is enabled, the vServer is available uninterruptedly. The vServer is not shutdown when the backend infrastructure fails and it is migrated to another server and started. When high availability is not enabled, the vServer starts only after the backend infrastructure is available. Cloud administrator takes care of server pool settings to provide high availability feature in a vDC.

- **Determine the use of the vServers access control mechanism**

You must determine the access mechanism to log in to your systems. An option available as an alternative secure method is to use a public key to authenticate SSH root access to the vServer without the use of a password. This option is available to all cloud users, other options to create local user account or provide password for remote SSH access are also available.

- **Determine the use of distribution groups for vServers**

You can determine if a vServer needs to be added to a distribution group. This feature is available for vDCs based on Oracle VM Server for SPARC virtualization technology and Oracle Exalogic systems. Distribution group enforces that two vServers are not running in the same Oracle VM Server. If you choose to assign a vServer to a distribution group, the distribution group must exist before creating the vServer. For more information about distribution group, see [Creating Distribution Groups](#).

Some additional tasks that you can perform before or after creating a vServer:

- **Define the use of additional storage for vServers**

When you require additional storage for vServers, attach volumes to vServers. You might attach volumes to vServers at vServers creation time, in that case, volumes must exist before creating the vServers. For more information about creating volume, see [Volume Creation](#).

You can skip this option at vServer creation time, you can also attach volumes at a later time.

Creating vServers

Oracle Enterprise Manager Ops Center provides a wizard that walks you through a series of steps to collect information and resources to create vServers. Before running this wizard, it is important to have your planning tasks completed.

To Create vServers

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create vServer** in the Actions pane.
The Create vServer Wizard is displayed.
4. Enter the following information in the **vServer Details** step:
 - Name and description for the vServer.
 - Tags for better identification and classification of the vServer.
 - Number of vServers to create.
 - High Availability Support option is selected by default.Click **Next** to select a server template.
5. Select a server template from the list.
Click **Next** to select a vServer type.
6. Select a vServer type from the list and click **Next**.
7. Select one or more volumes from the Available Volumes list. Use the arrow keys to move the selected volumes to the Attached Volumes list.
If not required, skip this step as the volumes can also be attached after creating a vServer. Click **Next** to continue.
8. Select one or more vNets from the list.
Click **Next** to select the IP address assignment method.
9. Select the IP address assignment method for each vNet and click **Next**.

When the static method is selected, choose an IP address from the list.

10. When the Distribution Group Selection step is present, select a distribution group from the list if required.

If not required, skip this step. Click **Next** to provide a public key.

11. In the vServer Access step, choose from one of the available options to get remote SSH access to the vServer:
 - Provide a public key to authenticate the SSH root access to the vServer. Use either of the following options to register the public key:
 - Check the **Upload Local Key** option to upload the public key from a local file, then click **Next** to view the Summary.
 - Paste the public key directly into the Public Key text area, then click **Next** to view the Summary.
 - Provide login credentials to authenticate the SSH access to the vServer. Select the option to supply the password for the remote SSH root access, and click **Next**. You are directed to the Login Credentials step. Enter the password and confirm it, then click **Next** to view the Summary.
 - For cloud users using Zones or Oracle OVM for SPARC as virtualization technology in the vDC, you can select the option to create a local user account and password for remote SSH access, and click **Next**. You are directed to the Login Credentials step. Enter the local user name and the password, then click **Next** to view the Summary.

You can also skip this step and click **Next** to view the summary.

12. Confirm the vServer information provided in the Summary and click **Finish** to launch the job to create the vServer.

After the job completes, the vServer is created and listed in the Navigation pane. By default, the DNS information is taken from the selected vNet and added in the `/etc/resolv.conf` file of the vServer.

Managing vServers

After creating a vServer, you can manage the vServer life cycle by executing the available actions for vServer management.

To locate all the actions available for managing a vServer:

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Select the vServer tab in the Center Pane.
4. You can then select a vServer from the vServers list displayed in the Center Pane and then choose one of the actions displayed.

Figure 18–9 vServer Management Options

Name	Description	OS	Memory Size (GB)	Storage Size (GB)	V...	Created by	Creation Date	St...
diskTestOC	Type: wls CatalogId: ...		4.0	5.3	1	rachidovab	Mon Oct 08 20...	

Cloud users can perform the following actions to manage a vServer life cycle. The availability of some actions depends on the cloud infrastructure. See [Actions Available for a Cloud User](#) for available actions under different cloud infrastructure:

- **Update a vServer:** To modify the current name or description of a vServer. This action also allows the creation or deletion of vServer tags.
- **Stop a vServer:** To stop a running vServer. When a vServer is stopped, the guest operating system is shutdown. A stopped vServer can be started later.
- **Start a vServer:** To start a stopped or shutdown vServer. The restarted vServer might not get the same IP address as the original vServer.
- **Pause a vServer:** To suspend a running vServer.
- **Resume vServer:** To start a suspended vServer.
- **Shutdown all vServers:** To shutdown all vServers listed for an account. This action is only available in the Actions pane.
- **Attach vServer volumes:** To attach one or more new volumes to a vServer. Do not stop the vServer to attach volumes. You might require to stop the vServer for vDCs based on zones.
- **Detach vServer volumes:** To detach one or more volumes from a vServer. Do not stop the vServer to detach volumes. You might require to stop the vServer for vDCs based on zones
- **Enable HA:** To enable high availability of the vServer. This ensures that the vServer is migrated and restarted on another Oracle VM Server when the current Oracle VM Server fails.
- **Disable HA:** To disable high availability of the vServer. If the Oracle VM Server fails, the vServer is shut down and restarted when the Oracle VM Server becomes available.
- **Save vServer as Template:** To save the vServer as a server template and use for new vServer creation.
- **Delete a vServer:** To delete a vServer from a vDC account. Deleting a vServer results in shutting down the vServer followed by the deletion of the vServer.

Creating vNets

vServers are bound to one or more vNets to restrict network connectivity. The different types of vNets that a cloud user can use are:

- **Public networks:** Created by cloud administrators. Cloud users cannot create, update, or delete this type of vNet. Cloud administrators can also share this type of vNet among a number of accounts in a vDC. vServers that are members of public vNets have also external communication beyond vDCs. You can also use the vServer to host public services.
- **Private vNets:** Created by cloud users according to their requirements and within the limits of the account quota. A private vNet is created based on the private network from the network domain of the vDC. Private vNets are only accessible within an account. All vServers that have membership to a private vNet in common can communicate freely through that subnet.

Cloud users define which vNets are associated with a vServer. You can specify the membership of a vServer to one or more vNets when you create a vServer. Once a vNet is associated with a vServer, the association persists until the vServer is deleted.

A cloud user can release a reserved IP address that is not allocated to a vServer. Allocated IP addresses are only released when the vServer is deleted.

Before You Begin

Cloud users can create private vNets. To create a private vNet, you must:

- Plan vServers connectivity.
- Define the number of vServer that can be part of a private vNet.

Creating a Private vNet

Private vNet is a private virtual network set up exclusively for an account. The vServers associated with this vNet have private virtual IP address for internal communication.

To Create a Private vNet

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Private vNet** in the Actions pane.
The Create Private vNet Wizard is displayed.
4. Enter the following information in the Private vNet Details step:
 - Name and description for the private vNet.
 - Tags for better identification and classification of the private vNet.
 Click **Next** to configure the private vNet.
5. Select the number of elements for the private vNet.
This is the maximum number of vServers that can be part of this vNet. Use the slide bar to set the value. The values entered are rounded to the next value of 1, 5, 13, 29, 61, 125, 253, 509, 1021, 2045, 4093, and 8189.
Click **Next** to view the summary.
6. Confirm the private vNet information and click **Finish** to create the private vNet.

Managing vNets

Virtual network management involves the necessary actions to connect and restrict network access to vServers. After a vNet is created, cloud users can perform the following actions for virtual networks management:

- **Allocate vIP:** To allocate one or more IP addresses from a private or public vNet. IP addresses are dynamically allocated from those available IP addresses that have not been assigned to a vServer or allocated previously. You can use allocated vIP addresses for static assignment to vServers.
- **Deallocate vIP:** To release an IP address that was previously allocated from a public or private vNet. IP addresses assigned to a vServer are not listed and they cannot be deallocated. Once an IP address is deallocated, the IP address is available to the account.
- **Update a Private vNet:** To modify the name or description of a private vNet. This action also allows the creation or deletion of vNet tags.

- **Delete Private vNets:** To delete a private vNet. You cannot delete vNets associated with a vServer.
- **View Reserved IP Addresses:** To see a complete list of the reserved IP addresses for a public or private vNet. Lists all IP addresses of the vNet, the list includes allocated and unallocated IP addresses.
- **View Allocated IP Addresses:** To see a complete list of the allocated IP addresses from a public or private vNet. This option lists all allocated IP addresses, regardless of whether or not they are assigned to a vServer.
- **View Used IP Addresses:** To see a complete list of the used IP addresses from a public or private vNet. This list includes all IP addresses assigned to a vServer.

Managing Server Templates

Server templates designate the operating system and how it is installed while creating a vServer. Server templates are specific to processor architecture of the server pool and virtualization type and can be pre-built images or identifies the OS distribution.

Server templates are loaded into the storage libraries associated with the vDC and cannot be changed later. By default, a server template is bound to a specific account.

The proper server template must exist before creating vServers. Depending on the virtualization type of the vDC account, cloud users can have the following options:

- Use a default server template provided by the system for creating vServers. This is available only for Zones based cloud infrastructure.
- Upload a new server template to be used for creating vServers. This is available only for Oracle VM Server for x86 based cloud infrastructure.
- Create a new server template for installing vServers. This is available only for Oracle VM Server for SPARC based cloud infrastructure.
- Save a server template from an existing vServer to be used for creating vServers, see [Managing vServers](#).

Uploading Server Templates

To upload a server template, you can choose from two different server templates sub types:

- **Template:** Single virtual machine template that is ready to be deployed into virtualized platforms. Templates can be of format `.tar` or other file types. This option allows you to upload a multi-file template that is stored as a single server template for the account.
- **Assembly:** Collection of interrelated software appliances that can include a configuration of multiple virtual machines with their virtual disks and their inter connectivity. An assembly is contained in a single `.ova` (Open Virtualization Format Archive) file. When uploading an assembly, you can also create snapshots when registering an `.ova` file with shared virtual disks. Those snapshots are tagged with the template ID and the assembly ID.

Either select the server templates stored in the local host machines or from other locations that you can access. Use FTP, HTTP, or HTTPS protocols to upload the server templates from other locations.

Before uploading a server template, you must:

- Decide on a suitable server template sub type.

- Make sure the file for the server template is of the correct format.
- Check whether the file is accessible when uploading the server template from a file in other location.

To Upload a Server Template

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Upload Server Template** in the Actions pane.
The Upload Server Template Wizard is displayed.
4. Enter the following information in the Identify Server Template step:
 - Name and description for the server template.
 - Tags for better identification and classification of the server template.
 Click **Next** to specify server template details.
5. Choose a server template sub type.
6. Choose an option to upload the server template file.
When the URL option is selected, enter the complete URL. When using the sub type template, you can also specify multiple URLs when uploading a multi-file template.
Click **Next** to view the summary.
7. Confirm the server template information and click **Finish** to create the server template.

Creating Server Templates

The option to create a server template is available only when the vDC is built on Oracle VM Server for SPARC virtualization technology. You can create server templates from existing OS provisioning profiles. The list of profiles is filtered to display only the distribution which are available to vDC. If a required profile is not available, contact your cloud administrator for required OS provisioning profiles to install your vServers. The OS profiles that have been created for a selected ISO image is made available to create the server template.

To Create a Server Template

1. Select the account in the vDC Management section.
2. Click **Create Server Template** in the Actions pane.
The Create Server Template Wizard is displayed.
3. Skip the introduction and click **Next**.
4. Enter the name and description for the server template.
If required, enter tags for easy classification and categorization.
Click **Next**.
5. Select a profile that provisions Oracle Solaris OS.
You can install Oracle Solaris 10 or Oracle Solaris 11 OS on the vServer. Click **Next**.
6. Review the information and click **Finish** to create the server template.

The new template is created and available to be used to create a vServer.

Server Template Management

Perform the following actions for server templates management:

- **Update Server Template:** To modify the name or description of a server template. This action also allows the creation or deletion of server template tags.
- **Register Server Template:** To register a server template for public use. Registering a server template make the server template available to other accounts in the vDC.
- **Unregister Server Template:** To unregister a server template is accessible for public use. Unregistering a server template make the server template available to only to the vDC account from which the server template was created.
- **Delete Server Template:** To delete a registered or unregistered server template. You can only delete a registered server template from the account where the server template was originally created. Deletion of a server template does not influence a vServer that was created based on the server template. Deleting a server template that was created based on an assembly causes the deletion of all snapshots associated with that template.

These options are disabled when the vDC is built on Zones virtualization technology.

Volume Creation

A volume is a virtual block storage device that you can attach or detach from vServers. Volumes are bound to an account. Storage space for volumes is limited by the account's quota.

You can attach one or more volumes to a vServer at vServer creation time or at a later time.

To create a volume, you can:

- Create a new empty volume specifying only the size.
- Create a volume from a snapshot. A volume created from a snapshot can be empty or not. The size of the volume is defined by the snapshot. Availability of this action depends of the virtualization type of the vDC account. The action might be disabled for some cloud users.
- Import volume from another location. Volumes must be contained in a file of the format `.img`. Volume files must be accessible by the Enterprise Controller, using HTTP, HTTPS or FTP protocols. The action is available only when the vDC is built on Oracle VM Server for x86 based virtualization technology.

Volumes can be shared at volume's creation time. When a volume is shared, the volume is available for all the cloud users of the account.

Before You Begin

Before creating a volume a cloud user must:

- Verify the storage space available for the account and plan accordingly.
- Decide for a suitable option to create the volume.
When importing a volume, check for file accessibility and format.
- Define whether the volume must be shared.

Creating Volumes

You can create or import volumes. The procedures to create and import volumes are described in this section.

To Create an Empty Volume

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Volume** in the Actions pane.
The Create Volume Wizard is displayed.
4. Enter the following information in the Volume Details step:
 - Name and description for the volume.
 - Tags for better identification and classification of the volume.Click **Next** to configure the volume.
5. Check the **Shared** option if required.
6. Enter the size of the volume.
Click **Next** to view the summary.
7. Confirm the volume information and click **Finish** to create the volume.

To Create a Volume from a Snapshot

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Select the **Storage** tab in the center pane.
4. Select the **Snapshot** sub tab in the center pane.
5. Select a snapshot for the Snapshots list.
6. Click the **Create Volume from Snapshot** action.
The Create Volume From Snapshot Wizard is displayed.
7. Enter the following information in the Volume Details step:
 - Name and description for the volume.
 - Tags for better identification and classification of the volume.
8. Check the **Shared** option if required.
9. Click **Create** to create the volume.

To Import a Volume

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Import Volume** in the Action action.
The Import Volume Wizard is displayed.
4. Enter the following information in the Volume Details step:
 - Name and description for the volume.
 - Tags for better identification and classification of the volume.

5. Enter the URL in which the external volume resides.
6. Check the **Shared** option if required.
7. Click **Import** to import the volume.

Volume Management

After a volume is created, perform the following actions for volume management.

- **Update Volume:** To modify the name or description of a volume. The action also allows the creation or deletion of volume tags.
- **Delete Volume:** To delete a volume from a vDC account. Only volumes that are not currently attached to a vServer can be deleted. Deleting a volume results in releasing storage space. The process of deletion of a volume does not influence a snapshot that has been created previously based on that volume.

Creating Snapshots

A snapshot is an image of a volume at a given time. A snapshot captures the current state of the volume and is immutable. You can create snapshots for the following purposes:

- Backup of data stored on a volume
- Creation of new volumes based on a snapshot.

Snapshots availability and management are possible only when supported by the storage of the vDC. Contact a cloud administrator for verification.

You can create a snapshot from:

- An existing volume.
- Uploading an assembly. See [Managing Server Templates](#).

In an assembly's template, there are four types of exposed disks. Three of them are presented to cloud users as snapshots, this allows cloud users to optionally create a shared volume from the snapshot. The types of disk created as snapshots are:

- **Public Populated:** Extra disk with present data
- **Private Raw:** Dynamically created as empty disk, no disk data, just a size
- **Shared Raw:** Dynamically created as empty disk, no disk data, just a size

You can create a volume from a snapshot and attach those volumes to vServers.

Before You Begin

Before creating a volume, you must:

- Verify the storage space available for the account and plan accordingly.
- Decide for a suitable option to create the snapshot: from a volume or from an assembly.

To Create a Snapshot from Volume

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Snapshot** in the Actions pane.

The Create Snapshot Wizard is displayed.

4. Enter the following information in the Snapshot Details step:
 - Name and description for the snapshot.
 - Tags for better identification and classification of the snapshot.

Click **Next** to select a volume.
5. Select a volume from the list.

Click **Next** to view the summary.
6. Confirm the snapshot information and click **Finish**.

Snapshot Management

Perform the following actions for storage management

- **Update Snapshot:** To modify the name or description of a snapshot. The action also allows the creation or deletion of snapshot tags.
- **Create Volume from a Snapshot:** To create a volume base on a snapshot, see [Volume Creation](#)
- **Delete Snapshot:** To delete a snapshot from a vDC account. Deleting a snapshot does not affect volumes, snapshots exist independently of the volume.

Creating Distribution Groups

Distribution groups are necessary for properties similar to anti-affinity scaling. Distribution groups are available for vDCs based on Oracle VM Server for SPARC virtualization technology and Oracle Exalogic systems. The action to create a distribution group might be disabled to some cloud users.

A distribution group is bound to a specific account. You can assign a vServer to a distribution group only when you create the vServer.

You can create a distribution group and define its size. The size of the distribution group represents the minimum amount of anti-affinity desired for running vServers in the distribution group.

Distribution group rules are applied every time a vServer assigned to a distribution group is created or started. Distribution group rules enforce the following behavior for vServers assigned to a specific distribution group:

- When the number of running vServers in the distribution group doesn't exceed the size of the distribution group, then all running vServers are distributed in a separate server of the vDC. If a spare server is not available in the vDC, then the create or start task fails.
- When the number of running vServers in the distribution group exceeds the size of the distribution group, then running vServers are distributed in the servers of the vDC according to the normal account resources capacities.

At all times, distribution rules satisfy one of the following conditions for vServers assigned to a specific distribution group:

- If the number of running vServers in the distribution group exceeds the size of the distribution group, then the total number of servers in the vDC hosting running vServers is greater or equal to the distribution group size.

- If the number of running vServers in the distribution group does not exceed the size of the distribution group, then the number of servers in the vDC hosting running vServers is greater or equal to the total of vServers of the distribution group in running status.

If the distribution group rules are broken, then an alert is generated and it is displayed on the account and vDC level incidents report. Events that can break the distribution group rules are stopping, deleting, or migrating a vServer.

Before You Begin

Before creating a distribution group, you must:

- Be aware of the size limit for creating a distribution groups. Contact a cloud administrator to verify it.
- Plan the creation and vServers distribution for the distribution groups.

To Create a Distribution Group

1. Expand vDC Management in the Navigation pane.
2. Select the account from the vDC Accounts list.
3. Click **Create Distribution Group** in the Actions pane.
The Create Distribution Group Wizard is displayed.
4. Enter the following information in the Distribution Group Details step:
 - Name and description for the distribution group.
 - Tags for better identification and classification of the distribution group.Click **Next** to configure the distribution group.
5. Define the size of the distribution group.
Click **Next** to view the summary.
6. Confirm the distribution group information and click **Finish**.

Distribution Group Management

Perform the following actions for distribution groups management:

- **Update a distribution group:** To modify the name or description of a distribution group. The action also allows the creation or deletion of distribution group tags.
- Delete a distribution group.

If the distribution group rule is broken, then an alert is generated and displayed on the account and vDC level incidents report. The vServer that has violated the distribution group rule must be restarted. When a vServer is started, the accounts resources and distribution group rules are checked to place the vServer.

Related Resources

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

- [Chapter 17, "Server Pools"](#)
- [Chapter 16, "Oracle VM Server for x86"](#)
- [Chapter 14, "Oracle Solaris Zones"](#)

- [Chapter 7, "Networks"](#)
- [Chapter 6, "Storage Libraries"](#)
- *Oracle Enterprise Manager Ops Center Cloud Infrastructure API and CLI Reference Guide*
- *Oracle Enterprise Manager Ops Center Administration Guide*

For end-to-end examples, see the How To library at
http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

Part V

Command Line Interface

Part V contains the following chapters:

- [Chapter 19, "Command Line Interface"](#)

Command Line Interface

The Oracle Enterprise Manager Ops Center Command Line Interface (CLI) is an alternative to the browser user interface for Oracle Enterprise Manager Ops Center.

This chapter explains how to use the Oracle Enterprise Ops Center CLI and all the commands included within this tool.

The following information is included:

- [CLI Features](#)
- [Starting the CLI](#)
- [General Commands](#)
- [Universal Output Filters](#)
- [Command Scripts](#)
- [Mode Commands](#)

CLI Features

The Oracle Enterprise Manager Ops Center CLI can perform many, but not all, functions of Oracle Enterprise Manager Ops Center. The following charts show the available features.

Asset Discovery, Management, and Grouping

[Table 19–1](#) shows the functions that the CLI can perform for asset discovery, management, and grouping.

Table 19–1 *Asset Discovery, Management, and Grouping*

Function	Possible with CLI	Mode
Find asset	Yes	Discover
Execute discovery profiles	Yes	Discover
Add asset	No	n/a
Declare asset	No	n/a
Manage asset	Yes	Gear
Unmanage asset	Yes	Gear
Register asset	Yes	Gear
View asset data	Yes	Gear

Table 19–1 (Cont.) Asset Discovery, Management, and Grouping

Function	Possible with CLI	Mode
Update asset data	Yes	Gear
Reset a server	Yes	Gear
Refresh a server	Yes	Gear
Manage locator lights	Yes	Gear
Manage power settings	Yes	Gear
Reboot an OS	Yes	Gear
Reinstall an asset on an OS	Yes	Gear
Create a group	Yes	Groups
Add asset to a group	Yes	Groups
Move assets between groups	Yes	Group
Manage power settings of an asset in a group	Yes	Group
Delete a group	Yes	Group
Delete an asset	Yes	Gear

Operating System Provisioning and Patching

Table 19–2 shows the functions that the CLI can perform for operating system provisioning and patching.

Table 19–2 Operating System Provisioning and Patching

Function	Possible with CLI	Mode
Import an OS image	Yes	OSImage
Delete an OS image	Yes	OSImage
Create an OS profile	No	n/a
Import an OS profile	Yes	OSProfile
Export an OS profile	Yes	OSProfile
Clone an OS profile	Yes	OSProfile
Delete an OS profile	Yes	OSProfile
Import a JET template	Yes	OSProfile
Provision an OS	Yes	Gear
Check inventory	Yes	Update
Upload a knowledge base bundle	Yes	Update
Update operating systems	Yes	Update
Create OS update profiles	Yes	Update
Apply OS update profiles	Yes	Update
Add a configuration file	Yes	Update
Add a file to a distribution	Yes	Update
Run OS update profiles	No	n/a

Firmware Provisioning

Table 19–3 shows the functions that the CLI can perform for firmware provisioning.

Table 19–3 Firmware Provisioning

Function	Possible with CLI	Mode
Create a firmware image	Yes	FWImage
Update a firmware image	Yes	FWImage
Delete a firmware image	Yes	FWImage
Create a firmware profile	Yes	FWProfile
Update a firmware profile	Yes	FWProfile
Delete a firmware profile	Yes	FWProfile
Update firmware	Yes	Gear
Check firmware compliance	Yes	Gear
Run firmware reports	No	n/a

Administration

Table 19–4 shows the functions that the CLI can perform for administration.

Table 19–4 Administration

Function	Possible with CLI	Mode
Change connection mode	Yes	Update
Deploy a Proxy Controller	No	n/a
Manage users and roles	Yes	User
View roles	Yes	Jobs
Manage jobs	Yes	Jobs
View notifications	Yes	Notifications
Delete notifications	Yes	Notifications
Manage incidents	Yes	Incidents

Starting the CLI

The CLI must share information with the Enterprise Controller, therefore you must connect the CLI to the Enterprise Controller.

When the CLI runs on a system other than the Enterprise Controller, the system's credentials are necessary to connect to the Enterprise Controller. When the CLI runs on the same system as the Enterprise Controller, the CLI does not need credentials.

Invoking the Command Line Interface

You get access to the command line interface from the following locations:

- Oracle Solaris: /opt/SUNWoccli/bin
- Linux: /opt/sun/occli/bin

Connecting From the Enterprise Controller

1. Enter the `oc` command.

```
# ./oc
xvmSh>
```

You can also add some options. For more information about the options for the `oc` command, see [Using the CLI](#).

2. Enter the `connect` command.

```
xvmSh > connect
localhost >
```

The CLI is connected to the Enterprise Controller.

Connecting From a Remote System

1. Enter the `oc` command.

```
# ./oc
xvmSh>
```

You can also add some options. For more information about the options for the `oc` command, see [Using the CLI](#).

2. Enter the `connect` command with the `-h <enterprise controller>`, `-u <user name>`, and `-p <password>` | `-p@<filename for password>` options.

Examples:

```
xvmSh > connect -h EnterpriseController -u root -p rootpass
EnterpriseController >
```

```
xvmSh > connect -h EnterpriseController -u root -p@/tmp/pwd
EnterpriseController >
```

The CLI is connected to the Enterprise Controller.

Note: If you do not use the `-h`, `-u`, and `-p` options, you are prompted for the Enterprise Controller name, user name, and password.

Using the CLI

The CLI is accessed with the `oc` command. The `oc` command can be run with the following options:

- Use `-V` | `--version` to show the CLI version and exit.
- Use `-h` | `--help` to show the CLI help and exit.
- Use `-q` | `--quiet` to be as quiet as possible.
- Use `-d` | `--debug <debug level>` to specify a debug logging level. Debug logging levels are `DEBUG`, `FINEDEBUG`, `CMDSTAT`, `INFO`, `WARNING`, `ERROR`, `CRITICAL`, `OFF`, and `ALL`.
- Use `-c` | `--cmdfile <file>` to execute the `<file>` command file.

- Use `-e | --eval <command>` for one-line execution. The `<command>` command is the only command to be executed. This option can be run with one-liners. For more information about one-liners see [Using the CLI One-Liner](#).
- Use `-o | --outfile <file>` to write the output to the `<file>` file.

Examples:

```
#./oc -V
OpsCenter Command Line Interface 12.1.0

#./oc -e history
495 connect
496 version
```

Defining User Variables

Variables can be defined and then used within any command.

To define a user variable, enter `set <variable name>=<value>` in the command prompt.

Example:

```
#./oc
xvmSh> set JOBID=foo.3
xvmSh> connect
localhost> jobs.list -x $JOBID
```

Using the CLI One-Liner

The CLI accepts a list of commands as an argument.

To use the CLI one-liner, enter `./oc -e 'list commands separated by a semicolon'` in the command prompt.

Example:

```
#./oc -e 'connect ; jobs.list | grep FAILED'
```

General Commands

The following are general commands used either for CLI execution or as universal commands.

Note: You can use the Tab key to see a list of the available commands or as a completion key when typing a command.

Checking Connectivity

The `ping` command checks connectivity to the Enterprise Controller.

To check the connectivity to the Enterprise Controller, enter `ping` in the command prompt.

Example:

```
localhost > ping
OEMEC None is alive, version : 12.1.0.1848
```

Displaying User Configuration

The `printenv` command displays the current user configuration. The variables included in the user configuration are:

- `Output_format`: Format for the command output.
- `Parsable_text_separator`: Text separator for parsable command output.
- `Log_level`: Logging level.
- `History_length`: Number of commands saved in the history.
- `Debug`: Debug logging.
- `Pager`: Pager for long output.
- `Continue_on_error`: Continue when a command fails.
- `Text_wrap_width`: Maximum width, in characters, to be displayed in a line.

To display the current user configuration, enter `printenv` in the command prompt.

Example:

```
localhost > printenv
output_format=text
parsable_text_separator=
log_level=OFF
history_length=500
debug=false
pager=false
continue_on_error=false
text_wrap_width=-1
```

The user configuration can be changed using the `setenv` command.

Setting the User Environment

The `setenv` command sets the user environment in the following ways:

```
output_format=text|parsable
parsable_text_separator=text
log_level=OFF|ON
history_length=integer
debug=true|false
pager=true|false
continue_on_error=true|false
text_wrap_width=integer
```

To set the user environment, enter `setenv <variable>=<value>` in the command prompt.

Example:

```
localhost > setenv pager=true
```

Note: You must set the pager variable to `true` before running some commands, such as the `man` command.

Recording Console Input and Output

The `record` command copies all console inputs and outputs to a text file. You can save the resulting file in the directory you prefer. Issue the command without a file name to stop the recording.

To record the console inputs and outputs into a text file, enter `record <filename>` in the command prompt.

Example:

```
localhost > record foo.file
```

Note: Use the Tab key for filename completion.

Using a Source File

The `source` command uses command script files to source commands as if they were entered at the prompt. For more information about command scripts, see [Command Scripts](#). This command is also used to specify a file as the source of an operation when executing other commands. This file can be used for information that is shared by more than one script, or to keep secure information separate.

To use a source file, enter `source <filename>` in the command prompt.

Example:

```
EnterpriseController > source /tmp/cmds
```

Viewing the Version

The `version` command displays the version of the CLI, the Oracle Enterprise Manager Ops Center software, the Java SE Runtime Environment, and the Java HotSpot Server VM software.

To view the version of the CLI, enter `version` in the command prompt.

Example:

```
localhost > version
CLI version: 12.1.0
OpsCenter version: 12.1.0.1848
Platform: Java-1.6.0_21-Java_HotSpot-TM-_Server_VM,_17.0-b16,_Sun_Microsystems_
Inc.-on-SunOS-5.10-x86
Jython Version: 2.5.1
```

Waiting for a Job to Complete

You can set the CLI to wait until either the most recent job or a specific job finishes before launching a new job, using the `wait` command. This can be useful if one job, such as a discovery, is a requirement for another job, such as grouping or managing discovered assets. You can identify a job by its jobID or by its runID.

The `wait` command is used in command scripts where many commands are run almost at the same time, and a command must wait to be executed until the job launched by the previous command is completed. For more information about command scripts, see [Command Scripts](#). This command can be used from within any mode at any time or from the main command line.

To wait for a job to complete, enter `wait` in the command prompt.

Use the `-i | --jobid <job ID>` option to specify a job ID. The CLI waits for that job to complete.

Use the `-r | --runid <run ID>` option to specify a job runID. The CLI waits for that job to complete. Default is to wait for the last job.

Use the `-a | --all` option to wait for all jobs.

Use the `-n | --maxwait <maximum wait time>` to specify a timeout, a maximum number of seconds to wait for a job to complete.

Note: To interrupt a wait operation during interactive session, press Ctrl+C.

Viewing Help

You can use the `help` command to show the available commands. This command can be used from within any mode at any time or from the main command line.

To show the available commands, enter `help` in the command prompt.

Use the `-l` option to include detailed usage information.

Reviewing Commands

The `history` command shows recent commands. This command can be used from within any mode at any time or from the main command line.

To show the recent commands, enter `history` in the command prompt.

Use the `-n <number of commands>` option to limit the history list.

Example:

```
localhost > history -n 3
```

The history includes a command number for the sequence. To repeat a command in the history enter the `!` character with the command number and press the Tab key, as shown in the following example:

```
localhost > history -n 3
495 connect
496 version
498 history -n 3
#!496 <TAB>
#version
```

Displaying a Man Page

You can display the man page for a specific mode, using the `man` command. The syntax usage and description for all mode subcommands are included in the man page. This command can be used from within any mode at any time or from the main command line.

To display the man page for a specific mode, enter `man <mode name>` in the command prompt.

Example:

```
localhost > man gear
```

If no mode name is given, the CLI man page is displayed.

To see a list of man pages, enter the command and press the Tab key.

Note: The `man` command requires the `pager` environment variable to be set to `true` for multipage output before being run. For more information about how to set the `pager` variable to `true`, see [Setting the User Environment](#).

Ending a Mode

You can exit a mode and return to the initial CLI prompt, using the `end` command. This command can be used from within any mode at any time or from the main command line.

To exit a mode, enter `end` in the command prompt.

Example:

```
localhost/gear > end
localhost >
```

Disconnecting from Oracle Enterprise Manager Ops Center

The `disconnect` command closes the connection between the Oracle Enterprise Manager Ops Center Command Line Interface and the Enterprise Controller.

To disconnect from the Enterprise Controller, enter `disconnect` in the command prompt.

Example:

```
localhost > disconnect
xvmSh >
```

Exiting the CLI

You can exit the CLI from any mode, using the `exit` command. This command can be used from within any mode at any time or from the main command line.

To exit the CLI, enter `exit` in the command prompt.

Example:

```
localhost > exit
#
```

Universal Output Filters

For any main or mode command, you can constrain the output by appending one of the following filters to the command:

- `|grep regex` to filter by the regular expression
- `|count integer` to display only the number of output lines
- `|tee filename` to write the output to a file

Example:

```
localhost/jobs > list | tee /var/tmp/todays_jobs.txt
```

Note: For better results with `grep` and `count` filters, activate parsable output, using the `setenv` command.

Command Scripts

The CLI can also be used in a non-interactive way, using command scripts. The following example shows a script named `update.xvm` that writes the inventory of a server to a file:

```
connect finance_svr -u zz3282134 -p b@seball
update
csv_inventory -h finance_svr -f ~/inventory_latest.csv
wait
end
```

To create the file at any time, issue the following command:

```
oc --cmdfile update.xvm
```

To run a command script use the `source` command.

Example:

```
xvmSh > source update.xvm
```

Mode Commands

The commands in this section start a mode that includes a set of subcommands. Use the `man` command to display the man page for each mode.

- [Controllers](#) – Use controllers mode to migrate assets from one Proxy Controller to another Proxy Controller.
- [Credentials](#) – Use credentials mode to view and manage driver credentials.
- [Deploy-Setup](#) – Use deploy-setup mode to set up the plan and target to deploy an asset.
- [Deploy](#) – Use deploy mode to deploy an asset.
- [Discover](#) – Use discovery mode to discover gear (assets).
- [FWImage](#) – Use fwimage mode to create, view, and manage firmware images.
- [FWProfile](#) – Use fwprofile mode to create, view, and manage firmware profiles.
- [Gear](#) – Use gear mode to view and manage gear (assets) and provision operating systems and firmware.
- [Groups](#) – Use groups mode to view, create, and manage groups.
- [Jobs](#) – Use jobs mode to view and manage jobs.
- [Incidents](#) – Use incidents mode to view and manage incidents.
- [Notifications](#) – Use notifications mode to view notifications and change notification settings.
- [OSImage](#) – Use osimage mode to view, manage, and install OS images.
- [OSProfile](#) – Use opsprofile mode to view, manage, and import OS profiles.
- [ServerPools](#) – Use serverpools mode to view, and manage server pools.

- **Stats** – Use stats mode to view statistics about the Enterprise Controller.
- **Update** – Use update mode to view and manage OS update profiles, policies, snapshots, and distributions, and view and change the connection mode.
- **User** – Use user mode to create, view and manage users and roles.
- **Virtualization** – Use virtualization mode to view, and manage virtual hosts and virtual machines.

Unlike the `ecadm` and `proxyadm` commands, which are used for administering the Enterprise Controller and its Proxy Controllers, the CLI operates on the assets of the data center.

To execute a subcommand within a mode:

- Enter the mode, then enter the subcommand.
- Enter `<mode name>.<subcommand name>`.

Example:

```
localhost > jobs.list
```

For information about the mode commands that have a corresponding action in the browser user interface, see [Finding the Equivalent Browser User Interface Actions For the Mode Commands](#).

Controllers

Controller mode allows you to evacuate the assets managed by a failed Proxy Controller to other available Proxy Controllers. This mode also allows you to view the available Proxy Controllers that you can use to evacuate the assets.

Viewing Controllers

You can view a list of the available Proxy Controllers, Enterprise Controllers, and Virtualization Controllers.

- Enter `list`.

```
localhost/controllers > list
```

Controller list:

Name	UUID	Status	Version
10.131.110.95-PC	h4fm76fd-28be-4892-8uae-4i8b20b780e1	UNKNOWN	12.1.4.2311
10.131.105.97-PC	6485cdab-7654-ac6b-a781-4fg796791054	OK	12.1.4.2311
10.131.105.97-EC	abcd2f30-ef67-4564-ab12-bc989b6ab671	OK	12.1.4.2311
x44-sal-orac1-VC	987me955-1edt-4678-adi5-192ggg1e74e6	OK	12.1.4.2311

Evacuating Assets From a Failed Proxy Controller

You can evacuate the assets managed by a failed Proxy Controller to other available Proxy Controllers. The Proxy Controller to which the assets can be migrated is automatically determined based on the networks managed by the remaining available Proxy Controllers and their load.

- Enter `migrate_assets -U| --UUID <proxy_uuid>`.

```
localhost/controllers > migrate_assets -U h4fm76fd-28be-4892-8uae-4i8b20b780e1
```

WARNING: The failed Proxy Controller will be removed, and its MBeans will be unregistered after the job launched by this command is finished.

Credentials

Credentials mode allows you to view and manage the driver credentials necessary to discover and manage assets. The credentials also help to establish trust between internal components.

Viewing Credentials

You can view a list of existing driver credentials

- Enter `list`.

```
localhost/credentials > list
Driver Credentials:
  ID |      Name      | Type |      Description      |
-----|-----|-----|-----|
 6850 |      admin      |  SSH | root/admin SSH creds |
 6840 |   occosIPMI     |  IPMI |          IPMI         |
```

Use the `-i | --ids <credential IDs>` option to give a specific credential ID or a list of credential IDs. IDs must be separated by a comma.

Use the `-l | --detail` option to include detailed credential information.

Modifying Credential Details

You can modify the values of a driver credential. The `modify` command starts a submode with its own set of subcommands.

1. Enter `modify -i | --id <credential ID>`.

```
localhost/credentials > modify -i 6840
localhost/credentials/modify >
```

2. (Optional) Use the `show` command to display the current properties of the credential.

```
localhost/credentials/modify > show
description=IPMI
*sharedSecret=*****
*login=IPMI_user
*name=occosIPMI
```

Note: Properties marked with an asterisk (*) at the beginning must be set.

3. Enter `set <property>=<value>` to modify a property value of the credential.

```
localhost/credentials/modify > set name=IPMI_cred
```

4. If necessary, use the `unset <property>` command to unset a property or the `reset` command to reset all properties.
5. Use the `commit` command to apply the change.

Deploy-Setup

Deploy-setup mode allows you to specify a deployment plan and its target. The deployment plan must exist already and the asset must be already a managed asset.

Starting Deploy-Setup Mode

You must switch to deploy-setup mode to run commands for deploying assets. To start the deploy-setup mode, enter the deploy mode, which redirects you to the deploy-setup mode.

- Enter deploy in the CLI prompt.

```
localhost > deploy
localhost/deploy-setup >
```

Selecting a Deployment Plan

You must select an existing deployment plan.

- Enter plan <plan name>

Use the `-v` | `--version` option to specify a version of a plan.

```
localhost/deploy-setup > plan plan-fwp
Plan "plan-fwp" selected: version 1, description "None"
```

Selecting a Target Asset

You must select an asset as the target of the deployment plan. You must always indicate an asset name or public UUID. In case of ambiguity, where multiple assets have the same user friendly name, the user is informed to use UUID instead.

- Enter target <asset name or public UUID>

You can also add some other options:

Use the `-a` | `--add` option to add assets to current set of targets.

Use the `-d` | `--delete` option to delete assets from the current set of targets.

Use the `-g` | `--gear` <asset name> option to give an asset name or a list of asset names, separated by comma.

Use the `-U` | `--UUID` <public UUID> option to give an asset UUID or a list of asset UUIDs, separated by comma.

Example using an asset's friendly user name:

```
localhost/deploy-setup > target foo.us.oracle.com
```

Example using an asset public UUID:

```
localhost/deploy-setup > target 53f6e87f-3582-4cda-9c63-9121a3a8beb3
```

After you specify the plan and target, the mode changes from the deploy-setup mode to the deploy mode. This is reflected at the command prompt as follows:

```
localhost/deploy>
```

Deploy

Deploy mode allows you to deploy an asset.

Viewing the Plan

You can view the attributes of the deployment plan or the deployment plan structure.

To view the attributes of a deployment plan:

- Enter `show plan`.

```
localhost/deploy > show plan
Target: foo.us.oracle.com
```

```
Step: Update Firmware
```

```
Associated Profile:
{
  firmwarePackages:
  [
    ILOM-3_0_3_31_a-Sun_Fire_X4600M2
  ]
  dryRun: False
  subType: Server
  resetSP: False
  profileDescription:
  forceDowngrade: False
  forceReinstall: False
  profileName: fwp-profile
  network: None
  ip: None
}
```

To view the structure of a plan:

- Enter `show`.

```
localhost/deploy > show
Step: Update Firmware
```

```
{
  isDryRun: Dry run?
  type: boolean

  firmwarePackages: Firmware Packages
  type: array of [
    type: java.lang.String
  ]

  resetSP: Reset SP?
  type: boolean

  forceDowngrade: Force Downgrade?
  type: boolean

  forceReinstall: Force Reinstall?
  type: boolean

  ip: IP Address
  type: java.lang.String

  network: Network Information
  {
    gateway: Gateway IP Address
    type: java.lang.String
```



```

objectName: Object Name
type: java.lang.String

userFriendlyName: User Friendly Name
type: java.lang.String

netmask: Network Address Mask
type: java.lang.String

ip: IP Address
type: java.lang.String

}
}

```

To view the results of a deployment plan:

- Enter `show results`.

Example:

```

localhost/deploy > show results
Step: Provision OS
Operating-System-1

```

Setting the Attributes of a Deployment Plan

You can modify the value of the attributes of a deployment plan using the `set` command. Some deployment plans require specific information during the deployment. See [Table 19-5](#) for the names of the attributes you must set before applying the plan.

Table 19-5 *Deployment Plans and Attributes*

Deployment Plan	Attributes
Configure M-Series Hardware, Create and Install Domain	No per-target attribute collection required.
Configure RAID	No per-target attribute collection required.
Configure Server Hardware and Install OS	No per-target attribute collection required.
Configure Service Processor	No per-target attribute collection required.
Configure and Install Dynamic System Domain	No per-target attribute collection required.
Configure and Install Logical Domains	No per-target attribute collection required.
Configure and Install Virtual Machines	No per-target attribute collection required.
Create Boot Environment	No per-target attribute collection required.
Create Dynamic System Domain	No per-target attribute collection required.

Table 19–5 (Cont.) Deployment Plans and Attributes

Deployment Plan	Attributes
Create Logical Domains	<p>guestName: name of the new logical domain. This is just the name of the virtual server guest.</p> <p>(Optional) prime: set to true to set the resource plan data so that the already set profile data on the logical domain step are used to prime the resource data.</p> <p>Logical domains creation requires the setting of resource plan structures per each logical domain to create.</p>
Create Solaris Zones	<p>ipAddress: IP for the new zone, if mode is static.</p> <p>ipMode: IP address assignment mode. Value = static or dynamic.</p> <p>ifName: network interface name to identify NIC to assign to zone.</p> <p>hostName: host name of the new zone.</p> <p>(Optional) prime: set to true to set the resource plan data so that the already set profile data on the zone step are used to prime the resource data.</p> <p>Solaris zones creation requires the setting of resource plan structures per each logical domain to create.</p>
Configure and Install Solaris Zones	The only parameters that must be set are the same as those used in creating an Oracle Solaris zones deployment plan.
Create Virtual Machines	No per-target attribute collection required.
Provision Network	No per-target attribute collection required.
Provision OS	<p>ip: set this attribute if networking option selected in the OS provision profile was None or Link Aggregation.</p> <p>dataIp: set this attribute if networking option selected in the OS provision profile was IPMP.</p>
Install Server	The only parameters that must be set are the same than in provision OS.
Software Deployment / Update	No per-target attribute collection required.
Update Solaris 11 OS	No per-target attribute collection required.
Update BIOS Configuration	No per-target attribute collection required.
Update Firmware	No per-target attribute collection required.
Update Firmware and Install Oracle VM Server for SPARC	No per-target attribute collection required.

Use the following syntax to set an attribute or list of attributes of a deployment plan:

- For simple types:

Enter set "<attribute-name> = <'string-value'>" | <int> | <true|false>

Examples:

```
localhost/deploy > set "forceReinstall = false"
localhost/deploy > set "ip = '213.123.123.23'"
```

- For attributes that are backed by other struct-like value objects, use the following syntax:

Enter set "<attribute-name> = list or nested object"

```
set "network={gateway='213.123.123.111', userFriendlyName='ip address',
ip='213.123.123.4'}"
```

- If you use JavaScript Object Notation (JSON) format, use the following syntax:

```
set
"{'resourcePlans': [{'hostName': 'myhostname', 'networks': [{'networkBindings': [{'i
pAddress': '123.123.123.1'}]}]}]"
```

Or

```
set
"resourcePlans=[{hostName='myhostname', networks=[{networkBindings=[{ipAddress='
123.123.123.1'}]}]}]"
```

Or

```
set "<json-string> "
```

- For deployment plans on multiple targets that require unique values for each target, use the `set` command as follows:

```
localhost/deploy > set "@<target-uuid or
dynamic-target>:<key>=<'string-value'>"|<int>|<true|false>|list of nested
objects"
set "@<target-uuid or dynamic-target>:<json-string>"
```

Example:

```
localhost/deploy > set @<target-UUID>:"network={gateway='213.123.123.111',
userFriendlyName='ip address', ip='213.123.123.4'}"
```

Adding Attributes to a Deployment Plan

You can add attributes to a deployment plan using the `add` command. The syntax for adding attributes is similar to the syntax for setting attributes. However, for arrays, the `add` command appends an element to the array while the `set` command modifies the attribute of an existing element within the array.

Use the following syntax to add attributes to a deployment plan:

- For simple types:

```
Enter add "<attribute-name> = {'<string-value'>"}|<int>|<true|false>}"
```

Examples:

```
localhost/deploy > add "forceReinstall = false"
localhost/deploy > add "ip = '213.123.123.23'"
```

- For attributes that are backed by other struct-like value objects, use the following syntax:

```
Enter add "<attribute-name> = list or nested object"
```

```
localhost/deploy > add "network={gateway='213.123.123.111',
userFriendlyName='ip address', ip='213.123.123.4'}"
```

- If you use JavaScript Object Notation (JSON) format, use the following syntax:

```
"{'resourcePlans': [{'hostName': 'myhostname', 'networks': [{'networkBindings': [{'i
pAddress': '123.123.123.1'}]}]}]"
```

Or

```
add
"resourcePlans=[{hostname='myhostname',networks=[{networkBindings=[{ipAddress='
123.123.123.1'}}]}]"
```

Or

```
add "<json-string> "
```

- For deployment plans on multiple targets that require unique values for each target, use the `set` command as follows:

```
add "@<target-uuid or
dynamic-target>:<key>=<'string-value'>|<int>|<true|false>|list of nested
objects"
add "@<target-uuid or dynamic-target>:<json-string>"
```

Example:

```
localhost/deploy > add @<target-UUID>:"network={gateway='213.123.123.111',
userFriendlyName='ip address', ip='213.123.123.4'}"
```

Testing a Deployment Plan

To test a deployment plan, use the `dryrun` command to create the tasks without committing the job, this ensures no errors with data.

To dry run a deployment plan:

- Enter `dryrun`.

```
localhost/deploy > dryrun
```

Applying a Deployment Plan

To execute a deployment plan, set the attributes if required and then apply the plan. Some deployment plans need to have their attributes set before applying the plan.

To apply a deployment plan to an asset:

- Enter `apply`.

```
localhost/deploy > apply
```

Discover

Discover mode allows you to set up and run discovery jobs. You can create a discovery job or use a discovery profile to discover an asset. You can also execute an existing discovery profile.

Starting Discover Mode

You must switch to discover mode to run discovery commands.

- Enter `discover` in the CLI prompt.

```
localhost > discover
localhost/discover >
```

Creating a Discovery Job

The `create` submode is used to create and launch new discovery jobs. This submode has its own set of commands.

1. Enter `create`.

```
localhost/discover > create
localhost/discover/create >
```

2. Set the type of discovery job, for example, OS, ILOM, ALOM.

Discovering an OS

1. Enter `set type=os`.
2. Enter the discovery credentials.
 - Enter `set ssh_user=<user>` to set the login user name.
 - Enter `set ssh_password=<password>` to set the login password.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_user=<root user>` to set the root user name.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_password=<root password>` to set the root password.
3. Enter `set ipaddrs=<IP address or addresses>` to specify the target IP address or addresses. The following formats can be used:
 - A single IP address.
 - A comma-separated list of IP addresses enclosed in parentheses. Example: (192.168.1.1, 192.168.1.2).
 - An IP range in Classless Inter-Domain Routing (CIDR) subnet notation. Example: 192.168.1.0/24.
 - An IP range in hyphen notation. Example: 192.168.1.1 - 192.168.1.10.
4. (Optional) Enter `set ssh_port=<port>` to specify a port.
5. (Optional) Enter `set use_default=true|false` to specify if the factory default credentials are used or not for the discovery job.
6. (Optional) Enter the `show` command to display all discovery information. If necessary, use the `unset <parameter>` command to unset a discovery parameter or `clear` to reset and revert to default values all configuration data.
7. Enter `commit` to start the configured discovery job.

The job is launched, and the job ID is displayed.

Example of OS discovery:

```
localhost/discover > create
localhost/discover/create > set type=os
localhost/discover/create > set ssh_user=sysmgmt
localhost/discover/create > set ssh_password=l0ckit
localhost/discover/create > set ipaddrs=172.20.28.202
localhost/discover/create > show
discovery type: Solaris, Linux OS
use_default=False
ipaddrs=172.20.28.202
ssh_user=sysmgmt
ssh_password=l0ckit
ssh_root_user=
ssh_root_password=
ssh_port=
localhost/discover/create > commit
```

Discovering an ILOM Service Processor

ILOM is the service processor used in some Oracle Sun servers. An ILOM discovery also discovers operating systems.

1. Enter `set type=ilom`.
2. Enter the discovery credentials.
 - Enter `set ssh_user=<user>` to set the login user name.
 - Enter `set ssh_password=<password>` to set the login password.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_user=<root user>` to set the root user name.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_password=<root password>` to set the root password.
3. Enter `set ipaddrs=<IP address or addresses>` to specify the target IP address or addresses. The following formats can be used:
 - A single IP address.
 - A comma-separated list of IP addresses enclosed in parentheses. Example: (192.168.1.1, 192.168.1.2).
 - An IP range in CIDR subnet notation. Example: 192.168.1.0/24.
 - An IP range in hyphen notation. Example: 192.168.1.1 - 192.168.1.10.
4. Enter `set ipmi_user=<IPMI user name>`.
5. Enter `set ipmi_password=<IPMI user password>`.
6. (Optional) Enter the `show` command to display all discovery information. If necessary, use the `unset <parameter>` command to unset a discovery parameter or `clear` to reset and revert to default values all configuration data.
7. Enter `commit` to start the configured discovery job.

The job is launched, and the job ID is displayed.

Discovering an ALOM Service Processor

ALOM is the service processor used in some Oracle Sun servers. An ALOM discovery also discovers operating systems and systems with ILOM service processors.

1. Enter `set type=alom`.
2. Enter the discovery credentials.
 - Enter `set ssh_user=<user>` to set the login user name.
 - Enter `set ssh_password=<password>` to set the login password.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_user=<root user>` to set the root user name.
 - (Optional) If the login user name is not the root user, enter `set ssh_root_password=<root password>` to set the root password.
3. Enter `set ipaddrs=<IP address or addresses>` to specify the target IP address or addresses. The following formats can be used:
 - A single IP address.
 - A comma-separated list of IP addresses enclosed in parentheses. Example: (192.168.1.1, 192.168.1.2).

- An IP range in CIDR subnet notation. Example: 192.168.1.0/24.
 - An IP range in hyphen notation. Example: 192.168.1.1 - 192.168.1.10.
4. Enter `set telnet_user=<telnet user name>`.
 5. Enter `set ipmi_password=<password for telnet user name>`.
 6. (Optional) Enter the `show` command to display all discovery information. If necessary, use the `unset <parameter>` command to unset a discovery parameter or `clear` to reset and revert to default values all configuration data.
 7. Enter `commit` to start the configured discovery job.
- The job is launched, and the job ID is displayed.

Viewing Discovery Profiles

You can list all discovery profiles that discover systems.

- Enter `list_profiles`.
- Use the `-i | --ids <profile IDs>` option to give a specific discovery profile ID or a list of discovery profile IDs. IDs must be separated by a comma.

Viewing Discovery Ranges

You can list all the existing discovery ranges that discover systems.

- Enter `list_ranges`.
- Use the `-i | --ids <range IDs>` option to give a specific discovery range ID or a list of discovery ranges IDs. IDs must be separated by a comma.

Viewing Driver Credentials

You can list all driver credentials used to discover systems.

- Enter `list_credentials`.
- Use the `-i | --ids <credential IDs>` option to give a specific driver credential ID or a list of driver credentials IDs. IDs must be separated by a comma.

Viewing Discovery Profile Credentials

You can list all driver credentials for a discovery profile used to discover systems.

- Enter `list_profile_credentials -i | --id <discovery profile id>`.

Viewing Discovery Ranges

You can list all driver ranges for a discovery profile used to discover systems.

- Enter `list_profile_ranges -i | --id <discovery profile id>`.

Viewing Networks

You can list all proxy networks used to discover systems.

- Enter `list_networks`.

Executing Discovery Profiles

You can execute a discovery profile to discover systems.

- Enter `execute -i | --id <discovery profile id> -t | --targets <target friendly names>`.
Use `-i | --id <ID>` to execute the discovery profile with ID <ID>. Use `-t | --targets <target friendly names>` to give a list of target friendly names. Target names must be separated by a comma.
Use the `-n | --networkuuid <network UUID>` option to give a network UUID.
Use the `-c | --credentials <discovery credentials IDs>` option to give a specific discovery credential ID or a list of discovery credential IDs. IDs must be separated by a comma.
Use the `-m | --management <management credential ID>` option to give a management credential ID.
Use the `-s | --serviceTag <servicetag credential ID>` option to give a service tag credential ID.
Use the `-a | --agentLess` option to set agentless deployment on.
Use the `-r | --ranges <discovery ranges IDs>` option to give a list of discovery ranges IDs. IDs must be separated by a comma.

FWImage

Fwimage mode allows you to view and delete firmware images.

Starting Fwimage Mode

You must switch to fwimage mode to run firmware image commands.

- Enter `fwimage` in the CLI prompt.

```
localhost > fwimage
localhost/fwimage >
```

Viewing Firmware Images

You can view a list of available firmware images.

- Enter `list`.
Use the `-l | --detail` option to display a detailed list.

Creating a Firmware Image

You can create a firmware image.

1. Enter `create`.
The image creation submode is launched.
2. Enter `add_device -v | --vendor <vendor> -m | --model <model>` to add a device for the firmware image.
Use the `-v | --validated` option to specify a validated image.
3. Enter the criteria for the image using the `set` command.
 - Enter `set name=<name>` to set the firmware image name.
 - Enter `set version=<version>` to set the firmware version.
 - Enter `set type=<type>`.

- Enter `set make=<make>`.
 - (Optional) Enter `set description=<description>`.
 - (Optional) Enter `set size=<size>`.
 - (Optional) Enter `set powerOffRequired=<True|False>`. This criteria is set to False by default.
 - (Optional) Enter `set dependeeFirmwarePackageName=<dependencies>`. This criteria is set to None by default.
 - (Optional) Enter `set updateCount=<update count>`.
4. (Optional) Use the `show` command to show the current criteria. If necessary, use the `unset <property>` command to unset a property, or use the `reset` command to clear all criteria.
 5. Enter `commit` to create the new image using your criteria.

Updating a Firmware Image

You can update an existing firmware image.

1. Enter `update -i | --image <image name>`.
The image update submode is launched.
 2. Enter the new criteria for the image using the `set` command.
 - Enter `set version=<version>` to set the firmware version.
 - Enter `set type=<type>` to set the firmware type.
 3. Enter `commit` to update the firmware image.
- ```
localhost/fwimage > update -i myimage
localhost/fwimage/update > set version=1.1
localhost/fwimage/update > set type=ALOM-CMT
localhost/fwimage/update > commit
Firmware image updated successfully
```

### Deleting a Firmware Image

You can delete any firmware image.

- Enter `delete -i | --image <firmware image name>`.  
The firmware image is deleted.

## FWProfile

Fwprofile mode allows you to view and delete firmware profiles.

### Starting FW Profile Mode

You must switch to fwprofile mode to run firmware profile commands.

- Enter `fwprofile` in the CLI prompt.
- ```
localhost > fwprofile
localhost/fwprofile >
```

Viewing Firmware Profiles

You can view the existing firmware profiles.

- Enter `list`.
Use the `-l | --detail` option to display a detailed list.
A list of firmware profiles is displayed.

Checking the Status of a Firmware Profile

You can check the status of a firmware profile.

- Enter `status -p | --profile <profile name>`.
The firmware profile status is displayed.

Creating a Firmware Profile

You can create a firmware profile using one or more existing firmware images.

- Enter `create -p | --profile <firmware profile name> -i | --images <list of images separated by comma>`.
Use the `-d | --description <description>` option to add a description to the profile.
Use the `-r | --reinstall` option to force a reinstall when provisioning the firmware.
Use the `-g | --downgrade` option to force a downgrade when provisioning the firmware.
The firmware profile is created.

Updating a Firmware Profile

You can update an existing firmware profile.

- Enter `update -p | --profile <firmware profile name> -i | --images <list of images separated by comma>`.
Use the `-d | --description <description>` option to add a description to the profile.
Use the `-r | --reinstall` option to force a reinstall when provisioning the firmware.
Use the `-g | --downgrade` option to force a downgrade when provisioning the firmware.
The firmware profile is updated.

Deleting a Firmware Profile

You can delete a firmware profile.

- Enter `delete -p | --profile <firmware profile name>`.
The firmware profile is deleted.

Gear

Use gear mode to view and manage all assets, including operating systems, hardware, and chassis. Although the term for objects that the Enterprise Manager Ops Center software manages is "asset", the CLI uses the term "gear".

Starting Gear Mode

You must switch to gear mode before using other gear subcommands.

- Enter gear in the CLI prompt.

```
localhost > gear
localhost/gear >
```

Viewing Assets

You can list the assets managed by Oracle Enterprise Manager Ops Center

- Enter list.

```
localhost/gear > list
Registered Gear:
```

UFN	Manage state	Description	Asset Type	UUID	Maintenance
v64v	MANAGED	10.79.103.14	Server	6fdg9-fc-8b-91-dd2651d	NO_MAINTENANCE
c4pr	MANAGED	10.68.102.26	Server	5a544-dc-43-82-ef8886d	NO_MAINTENANCE

Use the `-l` `--detail` option to show asset details.

Example

```
localhost/gear > list -l
Registered Gear:
```

```
-----
Name : v64v
  Description : 10.79.103.14
  Managed State : MANAGED
  Status : OK
  Lock Info : None
  Is a Zone : False
  Public UUID : 6fdg9-fc-8b-91-dd2651d
  Asset Type : Server
Name : c4pr
  Description : 10.68.102.26
  Managed State : MANAGED
  Status : OK
  Lock Info : None
  Is a Zone : False
  Public UUID : 5a544-dc-43-82-ef8886d
  Asset Type : Server
```

Use the `f` `--filter` `<filter term>` option to filter the assets list.

Example

```
localhost/gear > list --filter v64v
Registered Gear:
```

UFN	Manage state	Description	Asset Type	UUID
v64v	MANAGED	10.79.103.14	Server	6fdg9-fc-8b-91-dd2651

Updating Asset Information

You can update the name of any asset, including the Enterprise Controller. You can also update an asset's description and associated tags.

- Enter `update -g | --gear <asset name>`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use the `-n | --name <new name>` option to give the asset a new name.

Use the `-d | --description <description>` option to update the asset's description.

Use the `-t | --tags <space-separated list of tags>` option to update the asset's tags.

Viewing Jobs Associated With Assets

You can view current and historical jobs associated with an OS, server, or chassis.

- Enter `show_jobs -g | --gear <asset name>`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Managing an Asset

You can manage discovered assets.

- Enter `manage -g | --gear <asset name> -u | --user <user name with access to asset> -p | --password <user's password>`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use the `-t | --type <type>` option to give the asset type. This is useful when assets might have the same name.

Unmanaging an Asset

You can unmanage a managed asset, removing the asset from the managed asset list and removing any agent.

- Enter `unmanage -g | --gear <asset name> -u | --user <user name with access to asset> -p | --password <user's password>`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use the `-t | --type <type>` option to give the asset type. This is useful when assets might have the same name.

Setting the Maintenance Mode of an Asset

You can place an asset in maintenance mode or remove an asset from maintenance mode.

- Enter `set_maintenance -g | --gear <asset name> on|off`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use the `on` option to set the asset in maintenance mode.

Use the `off` option to unset the maintenance mode of an asset.

Deleting an Asset

You can delete discovered assets to remove it from Oracle Enterprise Manager Ops Center.

- Enter `delete -g | --gear <asset name> .`

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Reinstalling an Agent Controller on an Operating System

You can reinstall an agent on one or more managed operating systems.

- Enter `reinstall_agent -g | --gear <comma-separated asset list> -v | --version <Agent version to install> -u | --user <user name> -p | --password <password>.`

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Rebooting an Operating System

You can reboot a managed operating system.

- Enter `reboot -g | --gear <OS asset name>.`

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Updating Firmware

You can provision updated firmware on one or more systems using an existing firmware profile.

- Enter `apply_firmware -g | --gear <comma-separated asset list> -p | --profile <firmware profile>`.

Use the `-U | --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Checking Firmware Compliance

You can compare the firmware for a list of assets to a firmware profile.

- Enter `fwprofile_compliance -g | --gear <comma-separated asset list> -p | --profile <firmware profile>`.

Provisioning an OS

You can provision a new operating system to a managed system. The `provision_os` command starts a submode with its own set of subcommands.

1. Enter `provision_os -g | --gear <asset name> -p | --profile <profile>`.

```
localhost/gear > provision_os -g targetsystem -p S10
localhost/gear/provision_os# >
```

2. Use the `set` command to set the provisioning properties. The properties are:

- `profile` – The OS profile to use for the provisioning job. Required.
- `register` – Specifies whether the new OS will be registered. Values are `true` or `false`.
- `manualnetboot` – Values are `true` or `false`.
- `system_type` – The OS that is provisioned. Values are `redhat`, `suse` or `solaris`.
- `server_name` – The name of the target system. Required.
- `hostname` – The new host name. Required.
- `domain`
- `netmask`
- `gateway`
- `nameserver`
- `ip_address` – The IP address for the OS. Required.
- `boot_path`

- `boot_device`
 - `console_baud`
 - `console`
3. (Optional) Use the `show` command to show the values of all provisioning properties. If necessary, use the `unset <property>` command to unset a property or the `reset` command to reset all properties.
 4. Use the `commit` command to launch the OS provisioning job.

Changing Locator Lights

You can activate or deactivate locator lights.

- Enter `set -g| --gear <asset name> -l| --locator <on|off>`.

Use the `-U| --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Enabling or Disabling Notifications

You can enable or disable notifications for a specific OS, server, or chassis.

- Enter `set -g| --gear <asset name> -n| --notifications <on|off>`.

Use the `-U| --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Powering Off a Server or Chassis

You can deactivate a managed server or a chassis.

- Enter `poweroff -g| --gear <asset name>`.

Use the `-U| --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name.

The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use the `-f` | `--force` option to force the server or chassis to power off.

Powering On a Server or Chassis

You can activate a managed server or a chassis.

- Enter `poweron -g` | `--gear <asset name>`.

Use the `-U` | `--UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Resetting a Server or Chassis

You can reset a managed server or chassis.

- Enter `reset -g` | `--gear <server or chassis name>`.

Use the `-U` | `--UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Use `-n` | `--netboot` to cause each server to do a netboot instead of booting from disk.

Use the `-f` | `--force` option to force the server or chassis to power off.

Refreshing a Server

You can refresh the data for a server.

- Enter `refresh -g` | `--gear <server name>`.

Use the `-U| --UUID <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Use the `--attributes <attribute list>` option to give a list of asset attributes instead of the asset name. This is useful when assets might have the same name. The asset attributes that can be used are the `type`, `UUID`, `UFN`, and `ip`. Attributes must be separated by a comma. The `type` attribute must be `os`, `server` or `chassis`, and the `ip` attribute must be a valid IPv4 address.

Examples:

```
--attributes=ip=192.168.1.1,type='server'
```

```
--attributes=type='server',UFN='foo'
```

Groups

Groups mode allows you to view and manage groups.

Starting Groups Mode

You must change to groups mode to run these commands.

- Enter `groups` in the CLI prompt.

```
localhost > groups
localhost/groups >
```

Viewing Groups

You can view the list of existing groups.

- Enter `list`.

Use the `-l| --detail` option to view a list showing group details.

Use the `-b| --brief` option to limit the output to group names.

Viewing Members of a Group

You can view the members of a group.

- Enter `list_members -g| --group <group name>`.

Viewing Group Memberships

You can view the groups to which an asset belongs.

- Enter `list_membership -n| --gear <asset name>`.

Use the `-U| --uuid <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Creating a Group

You can create a group.

- Enter `create -g| --group <group> -d| --description <description> -p| --parent <parent group>`.

The parent group must already exist.

Updating Group Attributes

You can change the name, description, and parent group of a group.

- Enter `update -g | --group <group>`.
 Use the `-n | --name <new name>` option to rename the group.
 Use the `-d | --description <new description>` option to give the group a new description.
 Use the `-p | --parent <new description>` option to give the group a new parent group.

Adding an Asset To a Group

You can add an asset to a group.

- Enter `attach -n | --gear <asset name> -g | --group <group>`.
 Use the `-U | --uuid <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Removing an Asset From a Group

You can remove an asset from a group.

- Enter `detach -n | --gear <asset name> -g | --group <group>`.
 Use the `-U | --uuid <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Moving an Asset From One Group To Another Group

You can move an asset from one group to another group.

- Enter `move -n | --gear <name of gear> -f | --from_group <current group> -g | --group <destination group>`.
 Use the `-U | --uuid <UUID>` option to give the asset public UUID instead of the asset name. This is useful when assets might have the same name.

Deleting a Group

You can delete a group. Deleting a group does not affect the assets within the group.

- Enter `delete -g | --group <group>`.
 Use the `-c | --child` option to indicate if sub-groups under the specified group must be also deleted.

Powering On the Assets in a Group

You can power on assets within a group.

- Enter `poweron -g | --group <group name>`.
 Use the `-n | --netboot` option to cause each server to do a netboot instead of booting from disk.
 Deactivated hardware within the group is powered on.

Powering Off the Assets in a Group

You can power off assets within a group.

- Enter `poweroff -g | --group <group name>`.

Use the `-f | --force` option to force the hardware to power off.

Activated hardware within the group is powered off.

Resetting Servers in a Group

You can reset servers inside a group.

- Enter `reset -g | --group <group>`

Use the `-f | --force` option to force the operation.

Use the `-n | --netboot` option to cause each server to do a netboot instead of booting from disk.

Incidents

Incidents mode allows you to view and manage existing incidents.

Starting Incidents Mode

You must switch to incidents mode to run incidents commands.

- Enter incidents in the CLI prompt.

```
localhost > incidents
localhost/incidents >
```

Viewing Incidents

You can view a list of existing incidents.

- Enter `list`.

Use the `-l | --detail` option to display detailed incidents information.

Use the `-s | --state <state>` option to filter incidents by `<state>`.

Use the `-u | --user <owner>` to filter incidents by `<owner>`.

Viewing Incident Details

You can view details of the existing incidents or a specific incident.

- Enter `show -i | --incidents <incident ID or comma-separated incident IDs list>`.

Use the `-l | --detail` option to display detailed incidents information.

Viewing Alerts Associated with Incidents

You can view the alerts associated with one or more incidents.

- Enter `show_alerts -i | --incidents <incident ID or comma-separated incident IDs list>`.

Use the `-l | --detail` option to display detailed incidents information.

Acknowledging an Incident

You can indicate that you are investigating one or more incidents.

- Enter `acknowledge -i | --incidents <comma-separated incident IDs list>`

Use the `-n | --note <note>` option to add a note.

Adding an Annotation to One or More Incidents

You can add an annotation to one or more incidents.

- Enter `annotate -i | --incidents <comma-separated incident IDs list>`
 Use the `-t | --type <annotation type>` option to specify an annotation type.
 Annotation types are `comment` or `suggested_action`.
 Use the `-s | --synopsis <synopsis>` option to specify an annotation synopsis.
 Use the `-n | --note <note>` option to add a note.

Assigning Incidents

You can assign one or more incidents to a user.

- Enter `assign -i | --incidents <comma-separated incident IDs list> -u | --user <user>`.
 Use the `-n | --note <note>` option to add a note.

Closing Incident

You can close one or more incidents.

- Enter `close -i | --incidents <comma-separated incident IDs list>`.
 Use the `-d | --disable-delay <disable delay>` option to disable the incident monitoring conditions, and reactivate the monitors after `<disable delay>` minutes.
 Use the `-n | --note <note>` option to add a note.

Jobs

Use the jobs mode to view existing jobs.

Starting Jobs Mode

You must change to jobs mode to run job commands.

- Enter jobs in the CLI prompt.

```
localhost > jobs
localhost/jobs >
```

Viewing Jobs

You can view a list of jobs.

- Enter `list`.

The list of current and historical jobs is displayed.

```
localhost/jobs > list
```

Job ID	Status	Owner	Job Type
amritsar.5	SUCCESS	root	UpdateGearJob
amritsar.4	SUCCESS	root	DeleteAssetJob
amritsar.3	SUCCESS	root	AgentProvisioningJob
amritsar.2	SUCCESS	root	AgentProvisioningJob
amritsar.1	PARTIALLYSUCCESSFUL	root	DeleteAssetJob

Use the `-l | --detail` option to view detailed job information.

Use the `-f | --filter <term>` option to filter the job list.

Use the `-o | --owner <job_owner>` option to list jobs owned by `job_owner`.

Use the `-s | --status <job_status>` option to list jobs with status `job_status`.

Use the `-t | --type <job_type>` option to list jobs with job type `job_type`.

Use the `-x | --id <job_id>` option to list jobs with job id `job_id`.

Use the `-m | --matchAll` option to list jobs that match all of the criteria indicated; default is for any search criteria to match.

Use the `-r | --runid <run ID>` option to retrieve information of a job with a specific runID.

Use the `-C | --limit <limit>` option to limit the number of jobs returned.

Use the `-A | --ascend` option to sort results ascending; default is descending.

Use the `-S | --sort <job_attribute>` option to sort results on `<job_attribute>`; default is to sort on job id.

Use the `-L | --nologdetails` option to not include log details in detailed listing of jobs.

Starting a Job

You can run an existing job using its job ID.

- Enter `run -x | --id <Job ID>`.

Use the `--simulate` option to run the job in simulation mode.

Repeating a Job on Failed Targets

You can re-run a partially successful job on its failed targets.

- Enter `rerun_on_failed -x | --id <Job ID>`.

Deleting a Job

Deleting a job removes the job information from Oracle Enterprise Manager Ops Center. You can only delete jobs that are stopped or completed.

- Enter `delete -x | --id <Job ID>`.

Stopping a Job

Stopping a job ends progress on the job.

- Enter `stop -x | --id <Job ID>`.

Use the `-f | --force` option to force the operation.

Notifications

Notifications mode allows you to view notifications.

Starting Notifications Mode

You must switch to notifications mode to run notifications commands.

- Enter notifications in the CLI prompt.

```
localhost > notifications
localhost/notifications >
```

Viewing Notifications

You can view all current notifications.

- Enter `list`.

A list of notifications is displayed.

```
localhost/notifications > list
ID   Severity  Type    Message
```

```
-----
1    MEDIUM   INFO    Enterprise Controller registered with Sun. (70002)
2    MEDIUM   INFO    Proxy Controller proxyone (172.20.27.185) has b...
3    MEDIUM   INFO    Proxy Controller proxytwo (172.20.26.218) has b...
```

Use the `-l | --detail` option to display a detailed list.

Use the `-f | --filter <filter term>` option to filter the notifications.

Use the `-s | --severity <severity>` option to display only notifications of given severity [HIGH|MEDIUM|LOW].

Use the `-c | --count <max>` option to specify a maximum count of notifications to display.

Use the `-u | --username <user>` option to display only notifications belonging to a user.

Use the `-i | --notification_id <ID>` option to display notifications of a specific ID.

Use the `-r | --status <status>` option to display only notifications according to their status [READ|UNREAD]. Default is to display both.

Deleting a Notification

You can delete a notification using its ID number.

- Enter `delete -i | notification_id <notification ID number>`.

```
localhost > delete -i 43
Successfully deleted notification
```

Use the `-u | --username <user>` option to delete notifications belonging to a user.

OSImage

OS image mode allows you to view and manage OS images.

Starting OS Image Mode

You must switch to OS image mode to run OS image commands.

- Enter `osimage` in the CLI prompt.

```
localhost > osimage
localhost/osimage >
```

Viewing OS Images

You can view OS images.

- Enter `list`.

Use the `-l | --detail` option to view a detailed list.

Importing an OS Image

You can import an OS image when it is located to the Enterprise Controller's local file system.

- Enter `import -f | --filename <image file name and location> -d | --description <description> -n | --name <image name>.`

localhost/osimage > import -f folder/location/image.iso -d Solaris 10 -n
Solaris 10
JobId = EnterpriseController.47

Deleting an OS Image

You can delete OS images.

- Enter `delete -n | --name <image name>.`

OSProfile

OS profile mode allows you to view, import, and export operating system profiles.

Starting OS Profile Mode

You must switch to OS profile mode to run OS profile commands.

- Enter `osprofile` in the CLI prompt.

```
localhost > osprofile  
localhost/osprofile >
```

Viewing OS Profiles

You can view OS profiles.

- Enter `list`.

Use the `-l | --detail` option to view a detailed list.

```
localhost/osprofile > list  
Profile Name: sles10Jeff1  
Description: sles10Jeff1  
Distro Name: sles10  
OS Version: SUSE-Linux-Enterprise-Server-SP2-10.2-0  
Platform: x86
```

Importing a JET Template

You can import a JET template.

- Enter `import_jet_template -p | --profile <profile name> -i | --image <OS image> -t | --template_path <local template path>.`

Importing an OS Profile

You can import an OS profile file. If you are importing a profile that you have previously exported, change the 'Name' and 'Payload_Name' fields before importing.

- Enter `import -f | --file <profile name>`

Exporting an OS Profile

You can export an OS profile to a designated file.

- Enter `export -p | --profile <profile name> -f | --file <file name.prof>.`

Copying an OS Profile

You can make a new profile by copying an existing profile.

- Enter `clone -p | --profile <original profile name> -n | --name <new profile name>`.

Deleting an OS Profile

You can delete an existing profile.

- Enter `delete -p | --profile <profile name>`.

ServerPools

Serverpools mode allows you to view the existing server pools, and to list and remove a member from a server pool.

Starting Serverpools Mode

You must switch to serverpools mode to run server pool commands.

- Enter serverpools in the CLI prompt.

```
localhost > serverpools
localhost/serverpools >
```

Viewing Server Pools

You can view a list of existing server pools.

- Enter `list`.

Viewing Members of a Server Pool

You can view the members of a server pool.

- Enter `list_members -U | --UUID <Server Pool UUID>`.

Use the `-l | --detail` option to include detailed credential information.

Removing a Member From a Server Pool

You can remove a member or a list of members from a server pool.

- Enter `remove_members -U | --UUID <Server Pool UUID> -M | --members <Server UUID>`.

Use the `-f | --force` option to force the server removal from the server pool.

Stats

In stats mode, you view statistics about the Enterprise Controller:

- Percentage of memory used
- Free memory (MB)
- Average load (15 minutes)
- CPU utilization percentage

Starting Stats Mode

You must switch to stats mode to run stats commands.

- Enter stats in the CLI prompt.

```
localhost > stats
localhost/stats >
```

Viewing Statistics

You can view statistics about the Enterprise Controller.

- Enter show.

Name	Used Mem %	Free Mem (MB)	Avg Load 15min	CPU Util %
EnterpriseController	34.28	6646	1.74	75.00
EnterpriseController	34.28	6646	1.74	75.00
EnterpriseController	34.28	6646	1.74	75.00

Use the `-i | --interval <interval>` option to specify a repeat interval in seconds.

Update

Update mode allows you to view and manage policies, profiles, and snapshots for OS updates.

Starting Update Mode

You must change to update mode to run these commands.

- Enter update in the CLI prompt.

```
localhost > update
localhost/update >
```

Viewing OS Update Policies

You can view the available OS update policies.

- Enter `list_policies`.

Use the `-l | --detail` option to display a detailed list.

Use the `-p | --policyname <profile_name>` option to display policies with a specific policy name.

Viewing OS Update Profiles

You can view the available OS update profiles.

- Enter `list_profiles`.

Use the `-l | --detail` option to display a detailed list.

Use the `-p | --profilename <profile name>` option to display profiles with a specific profile name.

Viewing OS Distributions

You can view the available distributions. By default, this command shows only active distributions.

- Enter `list_distros`.

Use the `-a | --all` option to display inactive and active distributions.

Viewing Snapshots

You can view the snapshots for a specific asset.

- Enter `list_snapshots -g | --gear <asset name>`.

Searching Inventory

You can search an OS asset or all assets for inventory whose name matches a search string. The search string used for this command must be a regular expression. For a summary of regular expression constructs, visit

<http://docs.oracle.com/javase/1.4.2/docs/api/java/util/regex/Pattern.html#sum>

- Enter `search_inventory -s | --search_string <search string>`.

Enter the `-g | --gear <asset name>` option to search a specific asset, or use the `-a | --allgear` option to search all gear. You must use one of these two options.

Use the `-t | --snapshot <snapshot name>` option to search for a specific snapshot. This option can only be used with the `-g` option.

Use the `-D | --description` option to search descriptions.

Use the `-c | --cvs` option to output to comma-separated values (CSV).

Use the `-f | --filename <filename>` option to specify a filename for CSV output.

Use the `-d | --delimiter <delimiter>` option to specify a delimiter character. The default is the pipe (|) character.

Use the `-e | --enclosure <enclosure>` option to specify an enclosure character. The default is the quote (") character.

Use the `-p | --append` option to append the data to an existing file.

Use the `-l | --detail` option to display detailed data.

Checking Inventory

You can check an OS asset to see if inventory is available for it.

- Enter `check_inventory`.

Enter the `-g | --gear <asset name>` option to check a specific asset, or use the `-a | --allgear` option to check all assets. You must use one of these two options.

Use the `-t | --snapshot <snapshot name>` option to check a specific snapshot. This option can only be used with the `-g` option.

Use the `-c | --cvs` option to output to comma-separated values (CSV).

Use the `-f | --filename <filename>` option to specify a filename for CSV output.

Use the `-d | --delimiter <delimiter>` option to specify a delimiter character. The default is the pipe (|) character.

Use the `-e | --enclosure <enclosure>` option to specify an enclosure character. The default is the quote (") character.

Use the `-r | --header` option to include a header line.

Use the `-p | --append` option to append the asset data to an existing file.

Checking Connectivity to the Knowledge Base

You can check the connectivity to the knowledge base.

- Enter `check_guus`.
 - Use the `-u | --user <user name>` option to specify a user name.
 - Use the `-p | --password <password>` option to specify a password.
 - Use the `--proxy <proxy host name>` option to specify a proxy.
 - Use the `--port <proxy port>` option to specify a proxy port.
 - Use the `--proxyuser <proxy user name>` option to specify a proxy user name.
 - Use the `--proxypass <proxy password>` option to specify a proxy password.

Querying Job History

You can view the job history for a specific OS asset.

- Enter `job_history -g | --gear <asset list>`.
 - Use the `-l | --detail` option to show a detailed listing.
 - Use the `-f | --fromdate <from date YYYYMMDD>` option to specify a from date. The default is seven days ago.
 - Use the `-t | --todate <to date YYYYMMDD>` option to specify an end date. The default is today.
 - Use the `-i | --jobids < job ID>` option to query a specific job or jobs, you must still supply the asset name.
 - Use the `-L | --log` option to print logs.
 - Use the `--fail` option to query failed status.
 - Use the `--ok` option to query ok status.
 - Use the `--warning` option to query warning status.
 - Use the `--nostatus` option to query nostatus status.
 - Use the `--install` option to query **install** action.
 - Use the `--upgrade` option to query **upgrade** action.
 - Use the `--uninstall` option to query **uninstall** action.
 - Use the `--downgrade` option to query **downgrade** action.

Querying Job Status

You can query job status.

- Enter `get_job_status`.
 - Enter the `-i | --jobids <job ID>` option to query specific jobs, or use the `-a | --alljobs` option to query all jobs. You must use one of these two options.
 - Use the `-l | --detail` option to print out detail, including logs if available.
 - Use the `-c | --csv` option to output to comma-separated values (CSV).
 - Use the `-f | --filename <filename>` option to specify a filename for CSV output.
 - Use the `-d | --delimiter <delimiter>` option to specify a delimiter character. The default is the pipe (|) character.
 - Use the `-e | --enclosure <enclosure>` option to specify an enclosure character. The default is the quote (") character.

Use the `-r | --header` option to include a header line.

Use the `-p | --append` option to append the asset data to an existing file.

Uploading a Knowledge Base Bundle

Knowledge base bundles contain current patch data. A knowledge base bundle must be moved to the system running the Enterprise Controller before it can be uploaded.

- Enter `load_kb_bundle -f | --filename <path to KB bundle>`.

Checking the Connection Mode

You can check the connection mode.

- Enter `check_disconnected`.

Changing to Connected Mode

You can change the Enterprise Controller from disconnected to connected mode.

- Enter `set_connected_mode -c | --connected`.

Changing to Disconnected Mode

You can change the Enterprise Controller from connected to disconnected mode.

- Enter `set_connected_mode -d | --disconnected`

Modifying an OS Asset

You can update a specific OS asset by installing, upgrading, and uninstalling specific packages and patches

- Enter `modify_gear -g | --gear <assetname> -y | --policy <policy name> -a | --actual`.

Omit the `-a` option to run the job as a simulation.

Use the `--si <search string>` option to specify a search term to add to the install list.

Use the `--su <search string>` option to specify a search term to add to the upgrade list.

Use the `--sr <search string>` option to specify a search term to add to the uninstall list.

The search strings used for this command must be regular expressions. For a summary of regular expression constructs, visit

<http://docs.oracle.com/javase/1.4.2/docs/api/java/util/regex/Pattern.html#sum>

Use the `--ni <node IDs to install>` option to specify node IDs for install.

Use the `--nu <node IDs to upgrade>` option to specify node IDs for upgrade.

Use the `--nr <node IDs to uninstall>` option to specify node IDs for uninstall.

Use the `-j | --jobname <job name>` option to specify a job name.

Use the `-e | --jobdesc <job description>` option to specify a job description.

Use the `-k | --taskname <taskname>` option to specify a task name.

Use the `-f | --profilename <profile name>` option to specify a profile name.

Use the `-D` | `--description` <profile description> option to specify a profile description.

Use the `-r` | `--failure_policy` <failure policy> option to specify a failure policy.

Use the `-x` | `--execution_policy` <execution policy> option to specify an execution policy.

Use the `-o` | `--seconds` <number of seconds> option to wait a specified number of seconds before running the job.

Use the `-c` | `--case` option to run a case sensitive search.

Use the `-d` | `--desc` option to search the description.

Example:

```
localhost/update > modify_gear -g "foo.us.oracle.com" --sr SUNWbzip -y "Yes To All" -a
```

Applying an Update Profile to an OS Asset

You can apply a profile to an asset, updating all packages and patches to match the profile.

- Enter `apply_profile_to_gear -g | --gear <assetname> -y | --policy <policy name> -f | --profilename <profile name> -a | --actual`

Omit the `-a` option to run the job as a simulation.

Use the `-j` | `--jobname` <job name> option to specify a job name.

Use the `-e` | `--jobdesc` <job description> option to specify a job description.

Use the `-k` | `--taskname` <taskname> option to specify a task name.

Use the `-r` | `--failure_policy` <failure policy> option to specify a failure policy.

Use the `-x` | `--execution_policy` <execution policy> option to specify an execution policy.

Use the `-o` | `--seconds` <number of seconds> option to wait a specified number of seconds before running the job.

The following is an example of a CLI command to deploy a profile job to one or more hosts:

```
localhost/update > apply_profile_to_gear -g <space separated list of host names> -y "Yes To All" -f <profile name>
```

Creating a Historical Snapshot

You can create a historical snapshot.

- Enter `create_historical_snapshot -g | --gear <asset name> -s | --snapshotname <snapshot name> -D | --description <snapshot description>`.

Deleting a Snapshot

You can delete an existing snapshot of an asset.

- Enter `delete_snapshots -g | --gear <asset name> -s | --snapshot <snapshot name>`.

Creating a Profile

You can create a new profile by specifying packages and patches to install, upgrade, and uninstall.

- Enter `create_profile`.

You must supply a search string or node id. The search string used for this command must be a regular expression. For a summary of regular expression constructs, visit

<http://docs.oracle.com/javase/1.4.2/docs/api/java/util/regex/Pattern.html#summary>

Use the `--si <search string>` option to specify a search term to add to the install list.

Use the `--su <search string>` option to specify a search term to add to the upgrade list.

Use the `-sr <search string>` option to specify a search term to add to the uninstall list.

Use the `--ni <node IDs to install>` option to specify node IDs for install.

Use the `--nu <node IDs to upgrade>` option to specify node IDs for upgrade.

Use the `--nr <node IDs to uninstall>` option to specify node IDs for uninstall.

Use the `-f | --profilename <distribution name>` option to specify a profile name.

Use the `-o | --distro <profile name>` option to limit by distro name.

Use the `-D | --description <profile description>` option to specify a profile description.

Use the `-t | --type <profile description>` option to specify a profile type.

Use the `-a | --all` option to include all matched packages/patches.

Use the `-c | --case` option to run a case sensitive search.

Use the `-r | --replace` option to replace an existing profile if the names are the same.

Creating a Profile From Inventory

You can create a new OS profile based on an existing system.

- Enter `create_profile_from_inventory -g | --gear <gear name>`.

Use the `-t | --snapshot <snapshot name>` option to search for a specific snapshot. This option can only be used with the `-g` option.

Use the `-f | --profilename <profile name>` option to specify a profile name.

Use the `-D | --description <profile description>` option to specify a profile description.

Use the `-r | --include_removes` option to include removed packages.

Deleting a Profile

You can delete a profile.

- Enter `delete_profiles -p | --profilename <profile name>`.

Adding a File to a Distribution

You can add a file stored on the Enterprise Controller to a distribution.

- Enter `add_file_to_distro -f | --file_path <full path to distribution> -o | --distro_name <distro name> -a | --category_name <category name>.`

Use the `-s | --security` option to mark the distribution as security.

Uploading Local Software in Bulk From a Directory

You can upload the content from a directory and its subdirectories to the Updates software library.

- Enter `bulk_upload_directory -D | --distribution <distribution name of uploaded files> -d | --directory <full path to the source directory>.`

Note: For more information about how to view the available distributions, see [Viewing OS Distributions](#)

Adding a Local Action

You can add a local action.

- Enter `add_local_action -a | --actionname <action name> -f | --filename <file name> -p | --parentname <parent category name> -D | --description <description> -c | --channelname <channel name> -t | --actiontype <action type>.`

The action type must be **macro**, **postaction**, **preaction**, or **probe**.

Adding a Local Category

You can add a local category.

- Enter `add_local_category -c | --category_name <category name> -p | --parent_name <parent name> -D | --description <description> -o | --distro <distribution name>.`

Deleting a Component

You can delete a component.

- Enter `delete_component -i | --nodeid <node id> -d | --distro <distro>.`

Use the `-k | --donotcheckforinstall` option to refrain from checking for install.

Adding a Configuration File

You can add a configuration file.

- Enter `add_configuration_file -f | --filename <filename> -p | --pathname <pathname to file> -v | --version <version> -c | --categoryname <category name> -D | --description <description> -C | --channelname <channel name>.`

Setting a Component File

You can set a component file.

- Enter `set_component_file -n | --componentname <component name> -d | --distro <distribution> -f | --filename <filename>.`

Finding Nodes

You can find nodes whose names match a search string. The search string used for this command must be a regular expression. For a summary of regular expression constructs, visit

<http://docs.oracle.com/javase/1.4.2/docs/api/java/util/regex/Pattern.html#sum>

- Enter `find_nodes -s | --search_string <search string>`.

Use the `-a | --all` option to list all nodes.

Use the `-l | --detail` option to print details.

Use the `-c | --case` option to specify a case-sensitive search.

As the following examples shows, when regular expressions are not used correctly, any node is listed as a result:

```
localhost/update > find_nodes -s SMCcurl-7.21* -a
Name                               Node id   Distro           Available  Certified
-----
```

```
localhost/update > find_nodes -s SMCcurl-7.21.2-0 [1/1/2012] -a
Name                               Node id   Distro           Available  Certified
-----
```

Correct examples:

```
localhost/update > find_nodes -s SMCcurl-7.21.* -a
Name                               Node id   Distro           Available  Certified
-----
SMCcurl-7.21.0-0 [1/1/2012]      41025315  SOLARIS10_SPARC  True       True
SMCcurl-7.21.1-0 [1/1/2012]      41125695  SOLARIS10_SPARC  True       True
SMCcurl-7.21.2-0 [1/1/2012]      42327548  SOLARIS10_SPARC  True       True
```

```
EnterpriseController/update > find_nodes -s SMCcurl-* -a
Name                               Node id   Distro           Available  Certified
-----
SMCcurl                           41002397  SOLARIS10_SPARC  False      True
SMCcurl                           46108251  SOLARIS9_SPARC   False      True
SMCcurl                           52104822  SOLARIS10_X86    False      True
```

User

User mode allows you to view and manage local and remote users and their roles.

Starting User Mode

You must switch to user mode to run these commands.

- Enter `user` in the CLI prompt.

```
localhost > user
localhost/user >
```

Viewing Users

You can view the existing users in the Enterprise Controller.

- Enter `show_local_users`.

A list of users is shown.

Use the `-r| --roles <role name>` option to view users with `<role name>` role. A list of users and their roles are shown.

Viewing Role Types

You can view the available role types in the Enterprise Controller.

- Enter `show_all_local_roles`.

Viewing User Roles

You can view the roles assigned to a specific user in the Enterprise Controller.

- Enter `show_local_user_roles -u| --user <user name>`.

Adding a User to the Enterprise Controller

You can add a user to the Enterprise Controller and grant a role to the new user. The user must already exist in the host where the Enterprise Controller is installed.

- Enter `add_local_user -u| --user <user name> -r| --role <role or comma-separated roles list>`.

For information about the available role types you can grant to a user, see [Viewing Role Types](#).

Granting Roles to a User

You can grant a user role to a user of the Enterprise Controller.

- Enter `grant_local_role -u| --user <user name> -r| --role <role or comma-separated roles list>`.

For information about the available role types you can grant to a user, see [Viewing Role Types](#).

Replicating User Roles

You can replicate the roles from a source user to a recipient user. The current roles and privileges of the recipient users are overwritten with the roles and privileges of the source user.

- Enter `replicate_user_roles -u|--user <source user> -r| --role <recipient user>`.

Removing Roles From a User

You can remove a role from a user of the Enterprise Controller. For more information about how to see the roles of a user, see [Viewing User Roles](#).

- Enter `revoke_local_role -u| --user <user name> -r| --role <role or list of comma-separated roles to remove>`.

Adding a Directory Server

You can add a directory server.

1. Enter `configureds`.

```
localhost/user > configureds
localhost/user/configureds >
```

2. Enter `set name=<name>` to set the directory server name.

3. Enter `set hostname=<host name>` to set a fully qualified host name.
4. Enter `set use_ssl=<false|true>` if you want to enable or disable SSL. Default is `false`.
5. Enter `set port=<port>` to set a port.
6. Enter `set user=<user name>` to set a user name to connect with.
7. Enter `set password=<password>` to set a password.
8. Enter `set root_suffix=<root suffix>` to set the root suffix.
9. Enter `set user_dn=<user DN>` to set the user search DN.
10. Enter `set user_scope=<user scope>` to set the user search scope.
11. Enter `set search_filter=<search filter>` to set the user search filter.
12. (Optional) Use the `show` command to show the values of all directory server properties. If necessary, use the `unset <property>` command to unset a property or the `reset` command to reset all properties.
13. Use the `commit` command to finish configuring the directory server.

Viewing Directory Servers

You can view a list of directory servers and their IDs.

- Enter `get_directory_servers`.
Use the `-d| --details` option to show directory server details.

Viewing Remote Users

You can get a list of directory server users and their roles.

- Enter `get_ds_users -d| --directory <directory server name>`.
Use the `-r| --roles <role>` option to filter users with `<role>` role.

Synchronizing All Remote Users

You can synchronize all directory server users and their roles with the cache on the Enterprise Controller.

- Enter `sync_all_ds_users -d| --directory <directory server name>`.

Synchronizing a Single User

You can synchronize a specific user on the directory server with the cache on the Enterprise Controller.

- Enter `sync_ds_user -d| --directory <directory server name> -u| --user <user>`.

Removing a Directory Server

You can remove a directory server from the Enterprise Controller.

- Enter `remove_directory_service -d| --directory <directory server name>`.

Deleting a User

You can delete a local user in the Enterprise Controller.

- Enter `delete_local_user -u| --user <user name>`.

Virtualization

The virtualization mode manages virtual hosts and virtual machines (Oracle Solaris zones only).

Starting Virtualization Mode

You must switch to virtualization mode to run virtualization commands.

- Enter `virtualization` in the CLI prompt.

```
localhost > virtualization
localhost/virtualization >
```

Viewing Hosts

You can view the list of guest containers (global zones).

- Enter `list_hosts`.

Viewing Guests

You can view the list of non-global zones for a specific global zone.

- Enter `list_guests -C | --container-on <global zone name>`.

Migrating a Non-Global Zone

You can detach a non-global zone from its non-functioning global zone and attach the non-global zone to a different global zone.

- Enter `startup (-z zonename -S ObjectName | -Z ObjectName) -D ObjectName`.

Use the `-z | --zonename` option to specify a guest (non-global zone). If you select this option you must specify a source global zone object name.

Use the `-S | --source-on` option if you specify a zonename before. Using this option you specify a source global zone object name.

Use the `-Z | --zone-on` option in case you did not select the `-z | --zonename` option. Using this option you indicate the host (global zone) object name.

Use the `-D | --destination` option to indicate the destination (global zone).

Finding the Equivalent Browser User Interface Actions For the Mode Commands

You can use the Oracle Enterprise Manager Ops Center browser user interface (BUI) instead of using the command-line interface. You can find in the browser user interface many, but not all, equivalent actions for the CLI mode commands.

The following charts show the mode commands that have a corresponding action in the browser user interface:

- [Credentials Mode](#)
- [Deploy-Setup Mode](#)
- [Deploy Mode](#)
- [Discover Mode](#)
- [FWImage Mode](#)
- [FWProfile Mode](#)
- [Gear Mode](#)

- Groups Mode
- Incidents Mode
- Jobs Mode
- Notifications Mode
- OSImage Mode
- OSProfile Mode
- ServerPools Mode
- Update Mode
- User Mode
- Virtualization Mode

Controllers Mode

Table 19–7 shows the commands included in the Controllers mode and their equivalent actions in the BUI.

Table 19–6 Controllers Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Open the alert indicating that a Proxy Controller has failed> Migrate Assets> Auto Balance across Proxy Controllers> Migrate
<code>migrate_assets</code>	Open the alert indicating that a Proxy Controller has failed> Migrate Assets> Auto Balance across Proxy Controllers> Migrate

Credentials Mode

Table 19–7 shows the commands included in the Credentials mode and their equivalent actions in the BUI.

Table 19–7 Credentials Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Administration> Enterprise Controller> Credentials
<code>modify</code>	Navigation Pane> Administration> Enterprise Controller> Credentials> Edit Credentials

Deploy-Setup Mode

Table 19–8 shows the commands included in the Deploy-Setup mode and their equivalent actions in the BUI.

Table 19–8 Deploy-Setup Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>plan</code>	Navigation Pane> Plan Management> Deployment Plans
<code>target</code>	Navigation Pane> Plan Management> Deployment Plans> Select a type of plan> Select a plan> Action Pane> Apply Deployment Plan

Deploy Mode

Table 19–9 shows the commands included in the Deploy mode and their equivalent actions in the BUI.

Table 19–9 Deploy Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>show plan</code>	Navigation Pane> Plan Management> Deployment Plans> Select a type of plan> Select a plan> Center Pane> Details tab
<code>show</code>	NA
<code>show results</code>	Navigation Pane> Plan Management> Deployment Plans> Select a type of plan> Select a plan> Center Pane> Results tab
<code>set</code>	Navigation Pane> Plan Management> Deployment Plans> Select a type of plan> Select a plan> Actions Pane> Edit Deployment Plan
<code>add</code>	NA
<code>dryrun</code>	NA
<code>apply</code>	Navigation Pane> Plan Management> Deployment Plans> Select a type of plan> Select a plan> Actions Pane> Apply Deployment Plan

Discover Mode

Table 19–10 shows the commands included in the Discover mode and their equivalent actions in the BUI.

Table 19–10 Discover Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>create</code>	NA
<code>list_profiles</code>	Navigation Pane> Plan Management> Profiles and Policies> Discovery
<code>list_ranges</code>	NA
<code>list_credentials</code>	Navigation Pane> Administration> Enterprise Controller> Credentials
<code>list_profile_credentials</code>	Navigation Pane> Plan Management> Profiles and Policies> Discovery> Select a Discovery Profile> Discovery Credentials table in the Center Pane.
<code>list_profile_ranges</code>	NA
<code>list_networks</code>	NA
<code>execute</code>	Navigation Pane> Plan Management> Profiles and Policies> Discovery> Select a Discovery Profile> Actions Pane> Add Assets

FWImage Mode

Table 19–11 shows the commands included in the FWImage mode and their equivalent actions in the BUI.

Table 19–11 FWImage Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	NA

Table 19–11 (Cont.) FWImage Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>create</code>	NA
<code>update</code>	Navigation Pane> Libraries> Software Libraries> Initial EC Library> Actions Pane> Upload Firmware
<code>delete</code>	NA

FWProfile Mode

Table 19–12 shows the commands included in the FWProfile mode and their equivalent actions in the BUI.

Table 19–12 FWProfile Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Plan Management> Profiles and Policies> Firmware
<code>status</code>	NA
<code>create</code>	Navigation Pane> Plan Management> Profiles and Policies> Firmware> Actions Pane> Create Profile
<code>update</code>	Navigation Pane> Plan Management> Profiles and Policies> Firmware> Center Pane> Edit
<code>delete</code>	Navigation Pane> Plan Management> Profiles and Policies> Firmware> Center Pane> Delete

Gear Mode

Table 19–13 shows the commands included in the Gear mode and their equivalent actions in the BUI.

Table 19–13 Gear Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Assets> All Assets
<code>update</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Edit Attributes or Edit Tags
<code>show_jobs</code>	Navigation Pane> Assets> All Assets> Jobs tab
<code>manage</code>	NA
<code>unmanage</code>	NA
<code>set_maintenance</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Place in Maintenance Mode or Remove From Maintenance Mode
<code>delete</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Delete Assets
<code>reinstall_agent</code>	Navigation Pane> Assets> All Assets> Select an OS> Actions Pane> Switch Management Access
<code>reboot</code>	Navigation Pane> Assets> All Assets> Select an OS> Actions Pane> Reboot
<code>apply_firmware</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Update Firmware

Table 19–13 (Cont.) Gear Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>fwprofile_compliance</code>	NA
<code>provision_os</code>	Navigation Pane> Plan Management> Deployment Plans> Provision OS> Select a Plan> Actions Pane> Apply Deployment Plan
<code>set</code> (locator lights)	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Locator Lights On or Locator Lights Off
<code>set</code> (notifications)	NA
<code>poweroff</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Power Off
<code>poweron</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> Power On
<code>reset</code>	Navigation Pane> Assets> All Assets> Select a server Actions Pane> Reset Server(s)
<code>refresh</code>	Navigation Pane> Assets> All Assets> Select a server Actions Pane> Refresh

Groups Mode

Table 19–14 shows the commands included in the Groups mode and their equivalent actions in the BUI.

Table 19–14 Groups Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code> (Predefined Groups)	Navigation Pane> Assets> Server, Chassis, Operating Systems or All User Defined Groups
<code>list_members</code>	Navigation Pane> Assets> Server, Chassis, Operating Systems or All User Defined Groups
<code>list_membership</code>	Navigation Pane> Assets> Server, Chassis, Operating Systems or All User Defined Groups
<code>create</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Create Group
<code>update</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Edit Group
<code>attach</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Add Asset to Group
<code>detach</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Remove Asset from Group
<code>move</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Move Asset to Group
<code>delete</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Delete Group
<code>poweron</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Power On
<code>poweroff</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Power Off
<code>reset</code>	Navigation Pane> Assets> All User Defined Groups> Actions Pane> Reset Server (s)

Incidents Mode

Table 19–15 shows the commands included in the Incidents mode and their equivalent actions in the BUI.

Table 19–15 Incidents Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others
<code>show</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others
<code>acknowledge</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others> Acknowledge Incident(s)
<code>annotate</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others> Add Annotation to Incident(s)
<code>assign</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others> Assign Incident(s)
<code>close</code>	Navigation Pane> Message Center> Unassigned Incidents, My Incidents or Incidents Assigned to Others> Close Incident(s)

Jobs Mode

Table 19–16 shows the commands included in the Jobs mode and their equivalent actions in the BUI.

Table 19–16 Jobs Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Jobs Pane
<code>run</code>	Jobs Pane> Select a job> Re-Run Selected Jobs
<code>rerun_on_failed</code>	Jobs Pane> Select a failed job> Re-Run Selected Jobs
<code>delete</code>	Jobs Pane> Select a job> Delete Selected Jobs
<code>stop</code>	Jobs Pane> Select a job> Stop Selected Jobs

Notifications Mode

Table 19–17 shows the commands included in the Notifications mode and their equivalent actions in the BUI.

Table 19–17 Notifications Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Message Center>
<code>delete</code>	Navigation Pane> Message Center> Select a Notification> Delete

OSImage Mode

Table 19–18 shows the commands included in the OSImage mode and their equivalent actions in the BUI.

Table 19–18 OSImage Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Libraries> Software Libraries> Initial EC Library> Actions Pane> Import Image
<code>import</code>	Navigation Pane> Libraries> Software Libraries> Initial EC Library> Actions Pane> Import Image
<code>delete</code>	NA

OSProfile Mode

Table 19–19 shows the commands included in the OSProfile mode and their equivalent actions in the BUI.

Table 19–19 OSProfile Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Plan Management> Profiles and Policies> OS Provisioning
<code>import_jet_template</code>	Navigation Pane> Plan Management> Profiles and Policies> OS Provisioning> Create Profile> Select JET Template in the Subtype
<code>import</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Profiles> Center Pane> Import Profile
<code>export</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Profiles> Center Pane> Export Profile
<code>clone</code>	Navigation Pane> Plan Management> Profiles and Policies> OS Provisioning> Center Pane> Copy Profile
<code>delete</code>	Navigation Pane> Plan Management> Profiles and Policies> OS Provisioning> Center Pane> Delete Profile

ServerPools Mode

Table 19–23 shows the commands included in the ServerPools mode and their equivalent actions in the BUI.

Table 19–20 ServerPools Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list</code>	Navigation Pane> Assets> Server Pools
<code>list_members</code>	Navigation Pane> Assets> Server Pools> Select a server pool
<code>remove_members</code>	Navigation Pane> Assets> Server Pools> Select a server pool> Select a server> Actions Pane> Remove from Server Pool

Update Mode

Table 19–21 shows the commands included in the Update mode and their equivalent actions in the BUI.

Table 19–21 Update Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list_policies</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Policies

Table 19–21 (Cont.) Update Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list_profiles</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Profiles
<code>list_distros</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Content tab> Search Criteria Menu> Distribution list
<code>list_snapshots</code>	NA
<code>search_inventory</code>	NA
<code>check_inventory</code>	NA
<code>check_guus</code>	NA
<code>job_history</code>	Navigation Pane> Assets> Select an asset> Jobs tab
<code>get_job_status</code>	Jobs Pane> Status column
<code>load_kb_bundle</code>	Navigation Pane> Administration> Enterprise Controller> Actions Pane> Setup Connection Mode
<code>check_disconnected</code>	Navigation Pane> Administration> Enterprise Controller> Actions Pane> Setup Connection Mode
<code>set_connected_mode</code>	Navigation Pane> Administration> Enterprise Controller> Actions Pane> Setup Connection Mode> Switch to Connected Mode or Switch to Disconnected Mode.
<code>modify_gear</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> New Update OS Job
<code>apply_profile_to_gear</code>	Navigation Pane> Assets> All Assets> Select an asset> Actions Pane> New Update OS Job
<code>create_historical_snapshot</code>	NA
<code>delete_snapshots</code>	NA
<code>create_profile</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Profiles> Actions Pane> New Profile
<code>create_profile_from_inventory</code>	NA
<code>delete_profiles</code>	Navigation Pane> Plan Management> Profiles and Policies> Update Profiles> Select a profile> Delete Profile
<code>add_file_to_distro</code>	NA
<code>bulk_upload_directory</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Bulk Upload Packages
<code>add_local_action</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Upload Local Action
<code>add_local_category</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Add Local Category
<code>delete_component</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Delete Local Component
<code>add_configuration_file</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Upload Local Configuration

Table 19–21 (Cont.) Update Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>set_component_file</code>	Navigation Pane> Libraries> Software Libraries> Linux, Solaris 8-10 Software Update Library> Actions Pane> Edit Local Component File
<code>find_nodes</code>	NA

User Mode

Table 19–22 shows the commands included in the User mode and their equivalent actions in the BUI.

Table 19–22 User Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>show_local_users</code>	Navigation Pane> Administration> Enterprise Controller> Local Users
<code>show_all_local_roles</code>	Navigation Pane> Administration> Enterprise Controller> Local Users> Role Permission Mapping tab
<code>show_local_user_roles</code>	Navigation Pane> Administration> Enterprise Controller> Local Users
<code>add_local_user</code>	Navigation Pane> Administration> Enterprise Controller> Local Users > Actions Pane> Add User
<code>grant_local_role</code>	Navigation Pane> Administration> Enterprise Controller> Local Users> Manage User Roles icon
<code>replicate_user_roles</code>	Navigation Pane> Administration> Enterprise Controller> Local Users> Replicate User Roles icon
<code>revoke_local_role</code>	Navigation Pane> Administration> Enterprise Controller> Local Users> Manage User Roles icon
<code>configureds</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves> Add Directory Server
<code>get_directory_servers</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves>
<code>get_ds_users</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves> Select a Directory Server from the list> User tab
<code>sync_all_ds_users</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves> Select a Directory Server from the list> Actions Pane> Sync all remote users and roles
<code>sync_ds_user</code>	NA
<code>remove_directory_service</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves> Select a Directory Server from the list> Delete Directory Server
<code>delete_local_user</code>	Navigation Pane> Administration> Enterprise Controller> Directory Serves> Select a Directory Server from the list> User tab> Select a user> Delete User

Virtualization Mode

Table 19–23 shows the commands included in the Virtualization mode and their equivalent actions in the BUI.

Table 19–23 Virtualization Mode Commands and Their Equivalent Actions in the BUI

CLI Command	BUI Action
<code>list_hosts</code>	Navigation Pane> Assets> All Assets> Servers
<code>list_guests</code>	Navigation Pane> Assets> All Assets> Servers
<code>startup</code>	Navigation Pane> Assets> All Assets> Select a non-global zone> Actions Pane> Migrate Zone

Part VI

Engineered Systems

Part VI contains the following chapters:

- [Chapter 20, "Oracle Engineered Systems"](#)

Oracle Engineered Systems

The following information is included in this section:

- [Introduction to Oracle Engineered Systems](#)
- [Oracle Exalogic Elastic Cloud](#)
- [SPARC SuperCluster T4-4](#)
- [Related Resources](#)

Introduction to Oracle Engineered Systems

Traditionally, IT departments have been required to install and configure software that is hosted on hardware that was itself installed and configured separately. Oracle Engineered Systems relieve this burden from the IT department by providing a pre-defined combination of hardware and software. Oracle Engineered System is a complete set of integrated hardware and software designed to reach a predetermined level of capability, capacity, and scale. This relatively new concept of an integrated hardware and software solution dedicated to provide a specific service, changes the way one must think about architecture in IT environments.

In addition to removing the burden of installation and configuration, Oracle Engineered Systems provide significant cost savings, this is one of the many advantages delivered by Oracle Engineered Systems. An even larger advantage is the optimization that pre-defined hardware and software enables. As the hardware and software are engineered together to form a complete system, there are multiple opportunities to improve the overall system performance. The unified and integrated monitoring and management of the Engineered System also provides cost savings through simplification of the overall environment. Thus, Oracle engineered systems allow datacenter services to be delivered more efficiently via modular, dedicated systems. This greatly simplifies the entire purchase, deploy, configure, monitor, and manage lifecycle of the provided services.

The whole impetus for cloud computing is to provide specific services (whether IaaS, PaaS, or SaaS) as efficiently as possible; therefore, the increased efficiency delivered by Oracle Engineered Systems make them ideal candidates for the building blocks of a cloud computing environment.

In addition to creating a cloud computing environment, these building blocks address a variety of enterprise architecture requirements to provide a path for enterprise maturation alongside the move to cloud adoption. Each building block can be selected as needed to meet specific maturation needs, and thus ultimately they provide an increase in overall architectural flexibility. For example, an architect may choose to integrate an Exadata Storage Server and SPARC SuperCluster or Oracle VM Blade Cluster early on to begin the process of consolidation into a scalable architecture, and

then expand the system later to meet future demand. Shifting the overall environment into a standardized architecture opens possibilities for future shift into cloud adoption and/or continuation down the traditional enterprise computing maturation continuum while providing for future scale.

The Oracle Exalogic Elastic Cloud and SPARC SuperCluster T4-4 Engineered Systems are described in detail in the following sections.

Oracle Exalogic Elastic Cloud

The Oracle Exalogic Elastic Cloud comprises of a complete stack of hardware and software, computing, storage, and network, all engineered to work optimally together to provide a private cloud leveraging fusion middleware. Oracle Enterprise Manager Ops Center is a key part of the Exalogic appliance control system and serves as the operating system of the Exalogic platform.

SPARC SuperCluster T4-4

The SPARC SuperCluster T4-4 comprises of a complete stack of hardware and software, computing, storage, and network, all engineered to work optimally together to provide a consolidated platform for running database, middleware, or third party applications. Oracle Enterprise Manager Ops Center is closely integrated with SPARC SuperCluster and provides hardware management, provisioning, and virtualization management.

Datacenter Management

Using Oracle Enterprise Manager Ops Center, you can discover and manage Oracle Engineered Systems. You can view and access multiple Engineered Systems from a single datacenter through Oracle Enterprise Manager Ops Center. You can also view the Engineered System's assets, incidents reported from the Engineered Systems, and Service Requests from the same UI. You can generate reports for all datacenter assets, including Engineered Systems.

The discovery and management process is performed using one of the following methods:

- [Finding Assets](#)
- [Adding Assets Using a Discovery Profile](#)

Finding Assets

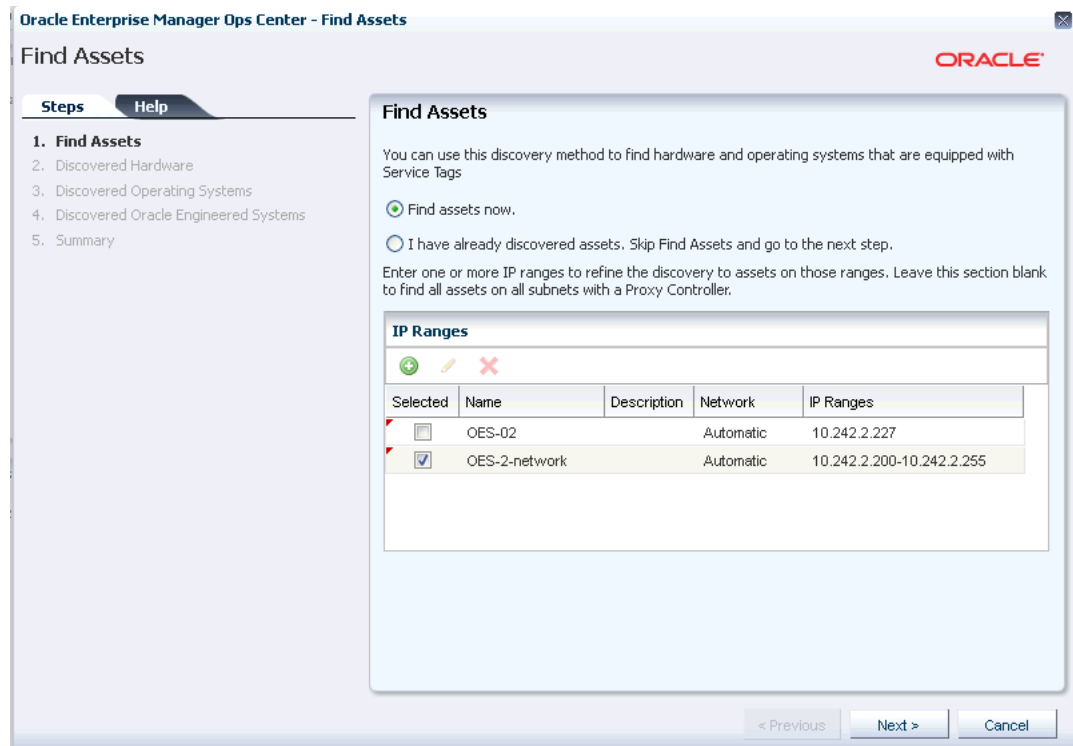
You can find assets using their service tags. Service tags are small XML files containing product information. Many Oracle systems come equipped with service tags. If you have hardware assets equipped with service tags, you can discover them using the Find Assets Wizard. This method lets you discover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

The Find Assets Wizard searches IP ranges or networks for service tags, then uses credentials that you specify to discover and manage the discovered assets.

Products without service tags cannot be discovered using this method.

To find assets, perform the following steps:

1. On the Navigation pane, under Assets, select All Assets.
2. In the Actions pane, click Find Assets. The Find Assets Wizard opens.

Figure 20–1 Find Assets

3. Select Find Assets Now and click Next.

When the initial discovery is complete, select the assets in each category (hardware, operating systems, and engineered systems) that you want to manage and provide necessary credentials.

Note: When adding a network range, enter the values without any space between the IP addresses. Use the format *start_ip-end_ip*.

4. Click Finish. The manage asset job is launched.

See [Discovering and Managing Assets](#) for more information.

Creating a Discovery Profile

To discover an asset, a discovery profile has to be created. After it has been created, you run a discovery job using the profile.

To create a discovery profile, perform the following steps:

1. On the Navigation pane, select Plan Management.
2. Under Profiles and Policies, select Discovery.
3. In the Actions pane, click Create Profile.
4. Enter a name for the discovery profile and select the type as Oracle Engineered Systems. Click Next.
5. If the service tag parameters have been modified for the target assets, enter the Service Tag parameters:

- Service tag passphrase - Necessary if a service tag has been configured to be encrypted.
 - Service tag port - Necessary if a service tag has been configured to use a port other than the default of 6481.
 - Service tag timeout - The default value is 20 seconds.
6. Enter the information for the target assets.
 - Host names or IPs - Enter a comma-separated list of host names or IP addresses. If you want to target an IP range, use the IP Ranges field instead.
 - Network - The managed network associated with the host names or IP addresses to route the discovery to the correct Proxy Controller. Select Automatic to route the job to the most appropriate Proxy Controller. The IP address of a target must resolve to only one known network for automatic routing to succeed.
 - IP ranges - Click the plus icon to add one or more IP ranges. Enter a name, description (optional), network, and IP range for each. If you want to target specific host names or IP addresses, you must enter them when you run the discovery profile.
 7. Enter credentials for each selected discovery profile.
 8. View the summary and click Finish to create the discovery profile.

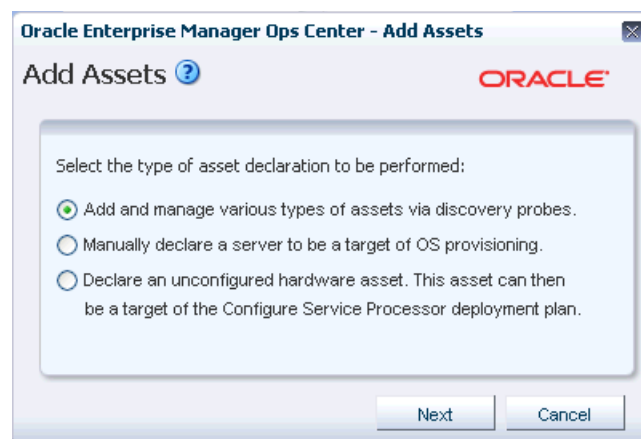
Adding Assets Using a Discovery Profile

After you have created a discovery profile, you run a discovery job and manage Engineered System's assets. This lets you discover assets using a pre-existing set of protocols and credentials, and manage assets consistently.

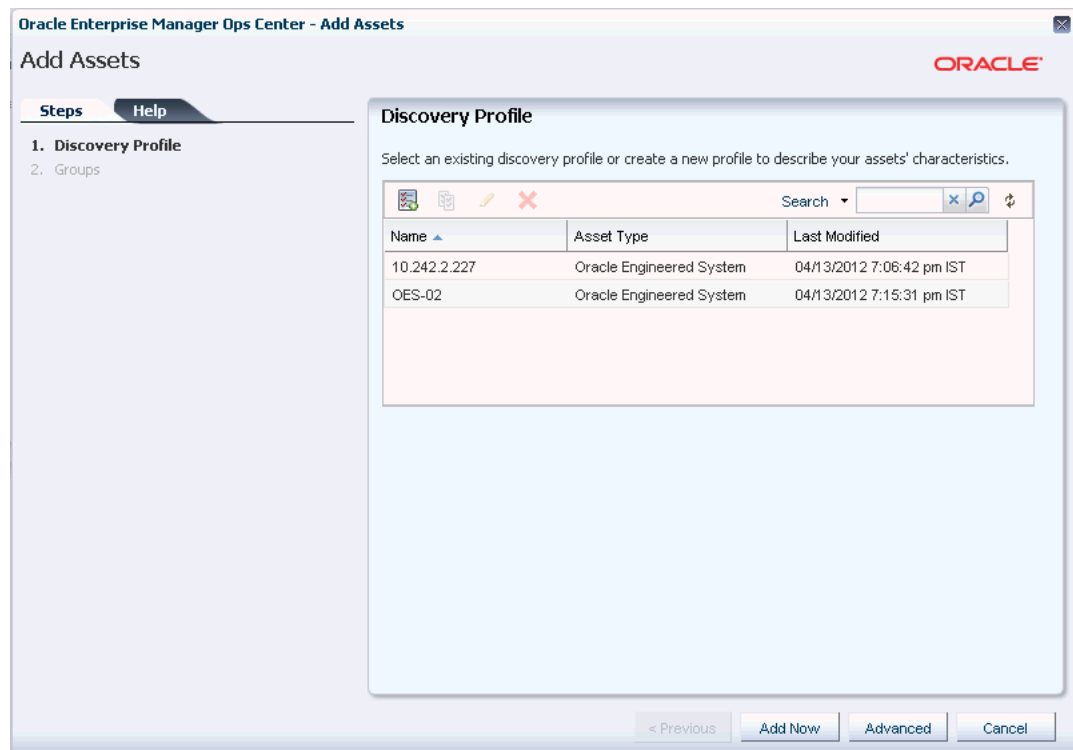
To declare unconfigured assets for service processor configuration, perform the following steps:

1. On the Navigation pane, under Assets, select All Assets.
2. In the Actions pane, click Add Assets. The Add Assets Wizard opens.

Figure 20–2 Add Assets



3. Select the option, Add and manage various types of assets via discovery probes. The Discovery profile window opens.

Figure 20–3 Select Discovery Profile

4. Select a discovery profile.
5. Add or edit the IP addresses, host names, and credentials for the targets.
6. Click Add Now.

Viewing the Asset Content Pane

You can view Engineered System's assets in the Datacenter Enterprise Controller BUI.

To view the Engineered System's assets in the Datacenter Enterprise Controller BUI, perform the following steps:

1. On the Navigation pane, under Assets, select Oracle Engineered Systems.
2. Select the Engineered System. The Dashboard, Details, Rack Info, Incidents, Service Request, and Monitoring tabs are displayed in the center pane.
 - Dashboard Tab – Displays the summary of the selected Engineered System, its membership graph, and the status of the system.
 - Details Tab – Displays the name of the selected Engineered System, description, master subnet manager address, number of racks, compute nodes, storage nodes, switches, and PDUs in the system.
 - Rack Info Tab – Displays a photorealistic view of the selected Engineered System.
 - Incidents Tab – Displays all incidents reported from the selected Engineered System.
 - Service Request Tab – Displays the service request for the selected Engineered System.

- **Monitoring Tab** – Displays the alert monitoring rules and service dependencies of the selected Engineered System.

To perform active management actions on any asset of Engineered Systems, launch the web interface of the selected Engineered System by clicking **Launch Web Interface** on the **Actions** pane.

Relayed Incidents

Relayed Incidents are datacenter incidents that are reported from Oracle Engineered Systems attached to Datacenter Enterprise Controller. You can only view and delete the relayed incidents. No actions can be performed on these incidents.

When you close the relayed incidents from the Datacenter Enterprise Controller, the incident is closed in the respective Engineered System as well.

To view all datacenter incidents, do the following:

On the **Navigation** pane, under **Message Center**, select **Relayed Incidents**.

All the relayed incidents are displayed.

See [Incidents](#) for more information.

Relayed Service Requests

You can view the relayed service requests related to the Oracle Engineered Systems. You can also file service requests.

To view the relayed service requests, do the following:

On the **Navigation** pane, under **Message Center**, select **Relayed Service Requests**.

All the relayed service requests are displayed. Details such as summary of the request, SR number, serial number, severity, status, contact person, and last updated date are displayed.

See [Oracle Services and Service Requests](#) for more information.

Baseline Check

Baseline Check is a feature in Oracle Enterprise Manager Ops Center Engineered Systems in which the factory default configuration of eighth, quarter, half, and full rack configurations are considered as a normal or baseline setup. When the assets are discovered and associated with the rack, this setup is compared with the default factory configuration. For any discrepancy, the baseline check reports a question mark at the slot in the rack where the local configuration differs from the factory default configuration.

Hover the cursor over the question mark to view details of the warning, for example, missing asset, asset mismatch, or unknown asset. Based on the warning, you can decide to adjust the assets.

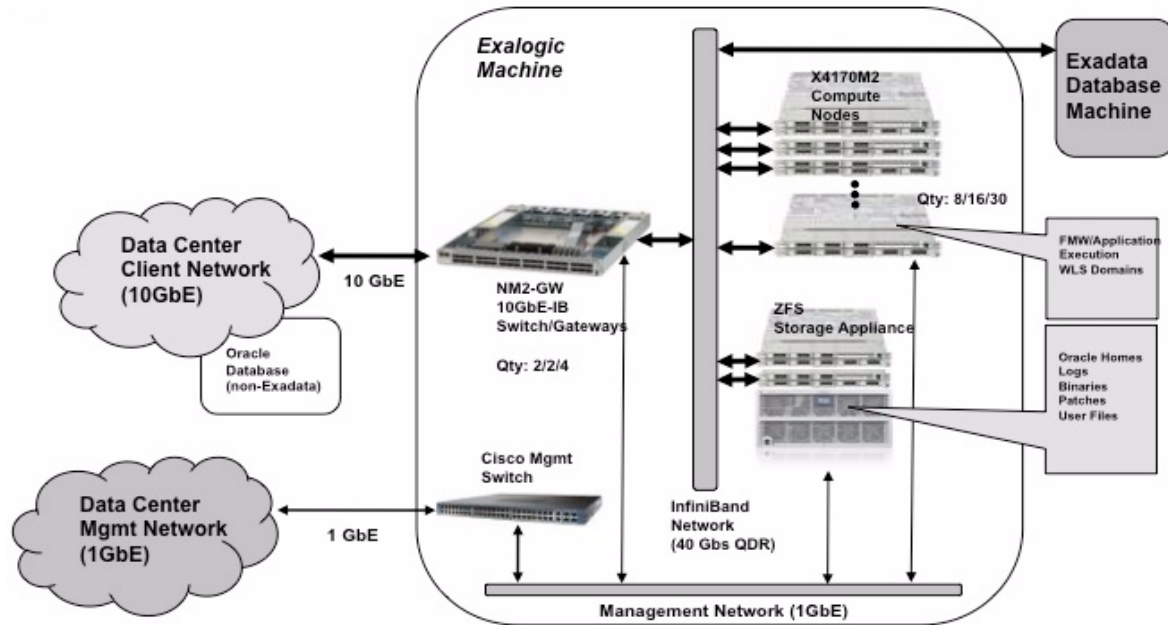
Oracle Exalogic Elastic Cloud

The Oracle Exalogic Elastic Cloud Engineered System delivers a standalone Oracle Enterprise Manager Ops Center that is preinstalled on a factory packaged Exalogic system. The embedded Oracle Enterprise Manager Ops Center is an integral part of Exalogic Control software and offers full set of Oracle Enterprise Manager Ops Center built-in functionality to provide Exalogic platform administration operations. It also provides features to view the Exalogic machine as an appliance, monitor and manage

the Exalogic hardware, and cover the replacement part of the fault management lifecycle.

Figure 20–4 is a pictorial representation of the Exalogic System.

Figure 20–4 Exalogic System Overview



Note: Quantities are specified as Quarter/Half/Full rack configurations.

Oracle Enterprise Manager Ops Center Configuration

The Exalogic Control software components are deployed in HA configuration to ensure that the platform management system is reliable and available.

To conserve the limited Exalogic resources and simplify the deployment model, Exalogic Control deploys a single database in HA configuration that is shared among other Exalogic Control software components such as Oracle Enterprise Manager Ops Center, Oracle VM Manager, and LDAP.

Understanding User Roles

In Oracle Enterprise Manager Ops Center, users can be assigned several roles such as Asset Admin, Cloud Admin, Exalogic System Admin, and many more. Each role grants the user a set of permissions. A particular permission can be granted by more than one role such as Asset Management, Network Management, and so on.

You can add users to Oracle Enterprise Manager Ops Center from the local authentication subsystem of the Enterprise Controller's operating system. Each user can be given a different role which grants or denies access to the different functions of Oracle Enterprise Manager Ops Center. You can view the existing users from the Users tab of the Administration section.

Exalogic System Admin Role

The Exalogic System Administrator is responsible for overall monitoring and management of the engineered system, including hardware subsystems, network management, and user management. The Exalogic System Administrator also has privileges to manage Oracle VM Servers.

The Exalogic System Admin role is a default role for the management of the Exalogic appliance and tasks related to the above functionality. As an Exalogic System Admin, you can perform most of the operations that are allowed in the context of the Exalogic system, except tasks related to cloud management.

For tasks related to cloud management, the Cloud Admin and Cloud User roles are available.

Prerequisites: The user must be familiar with the use of Oracle Enterprise Manager Ops Center and be familiar with hardware management and OS management.

Cloud Administrator

The Cloud Administrator's responsibilities include setting up of infrastructure and resource allocation so that cloud users can deploy their application onto authorized accounts. They also manage the cloud users accessing the accounts and their authorization.

Prerequisites: The user must be trained on Oracle Enterprise Manager Ops Center, installation and configuration, and the continual maintenance of the product.

Cloud User

Cloud users create virtual servers and deploy applications. Cloud users are restricted to virtual datacenter infrastructure activities and are presented only with the required options on the Oracle Enterprise Manager Ops Center UI.

Prerequisites: The user must be familiar with the use of Oracle Enterprise Manager Ops Center and be familiar with hardware management and OS management in general.

Adding a User

Users with the User Admin role can add other users. New user information, such as passwords for new users are retrieved from the local authentication subsystem.

Perform the following steps to add a user:

1. Select Administration in the Navigation pane.
2. Click Local Users. The Users page is displayed.
3. Click the Add User icon. The Add User window is displayed.
4. Enter the user name.
5. Add one or more roles to the list of Selected Roles.
6. Click Add User. The new user is created.

Viewing the Exalogic System

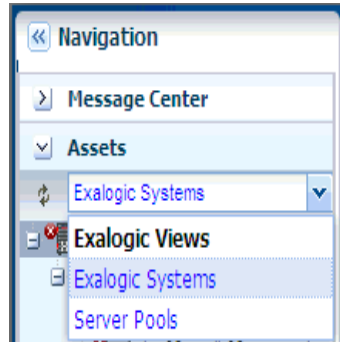
Viewing the Exalogic system basically means how you view the virtual Exalogic system using the four tabs on the Center pane, namely Dashboard, Details, Controls, and Networks. You can also perform relative actions by clicking the respective actions on the Actions pane. The actions that are allowed to be performed from the Actions

pane are context sensitive to the selected assets on the Navigation pane. The actions are also role based, not all users are allowed to perform all actions. For more information on Roles, see [Understanding User Roles](#).

To view the Exalogic system, perform the following steps:

1. Log in to Oracle Enterprise Manager Ops Center using the Exalogic Admin Role.
2. On the Navigation pane, under Assets, select Exalogic Systems from the drop-down list.

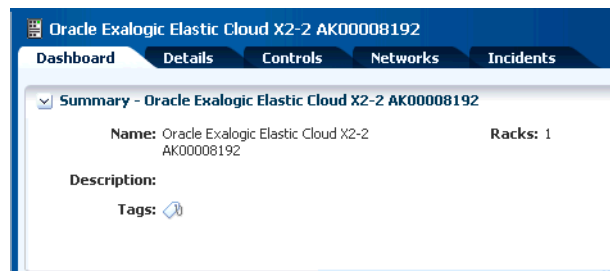
Figure 20–5 Navigation pane



3. Select the Exalogic system that you want to view.

The Dashboard, Details, Controls, Networks, and Incidents tabs are displayed in the center pane.

Figure 20–6 Center Pane With Dashboard Tab Selected



4. Click the tab on the Center pane to view more details.
 - **Dashboard Tab** displays the summary of the Exalogic system, its membership graph, and the status of the system.
 - **Details Tab** displays the name of the Exalogic system, description, master subnet manager address, number of racks, compute nodes, storage nodes, switches, and PDUs present in the system.
 - **Controls Tab** displays the Exalogic Control Software details such as name of the software, version number, description, name of the virtual machine, and server name.
 - **Networks Tab** displays the infrastructure network table and network connectivity table.
 - **Incidents Tab** displays all the incidents and alerts reported in the Exalogic System.

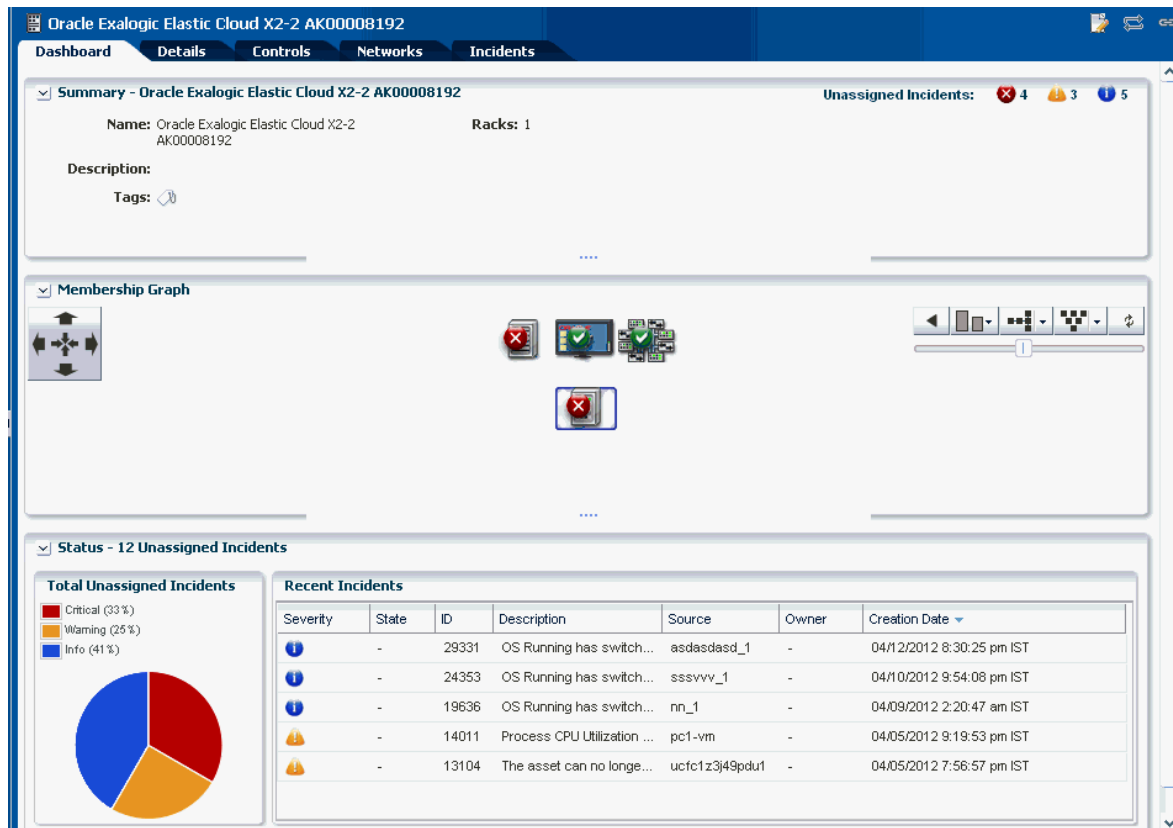
Dashboard Tab

The dashboard has three sections - Summary, Membership Graph, and Status.

In the Dashboard Tab, the following panes are displayed:

- [Summary Pane](#)
- [Membership Graph Pane](#)
- [Status Pane](#)

Figure 20–7 Dashboard Tab



Summary Pane The Summary section displays the name of the Exalogic system, description or name of the system identifier, number of racks that are part of the system, and also the number of unassigned incidents. The Unassigned Incidents icon in the summary pane includes incidents resulting from hardware faults. The three icons depict Critical Incidents, Warnings, and Information Incidents respectively. These incidents are from the assets that belong to all Exalogic racks.

Membership Graph Pane The Membership Graph pane displays the Exalogic system as a hierarchy of its components showing the relationship between the Oracle VM Manager, physical InfiniBand fabric, and the hardware that is grouped in the rack. You can navigate to any asset by double-clicking on the asset in the membership graph.

Using the controls on the top right of the graphic pane, you can change the view of the graph to either a horizontal or a vertical orientation. You can also refresh the view by clicking the Refresh icon. You can also change the graph depths or size of the images.

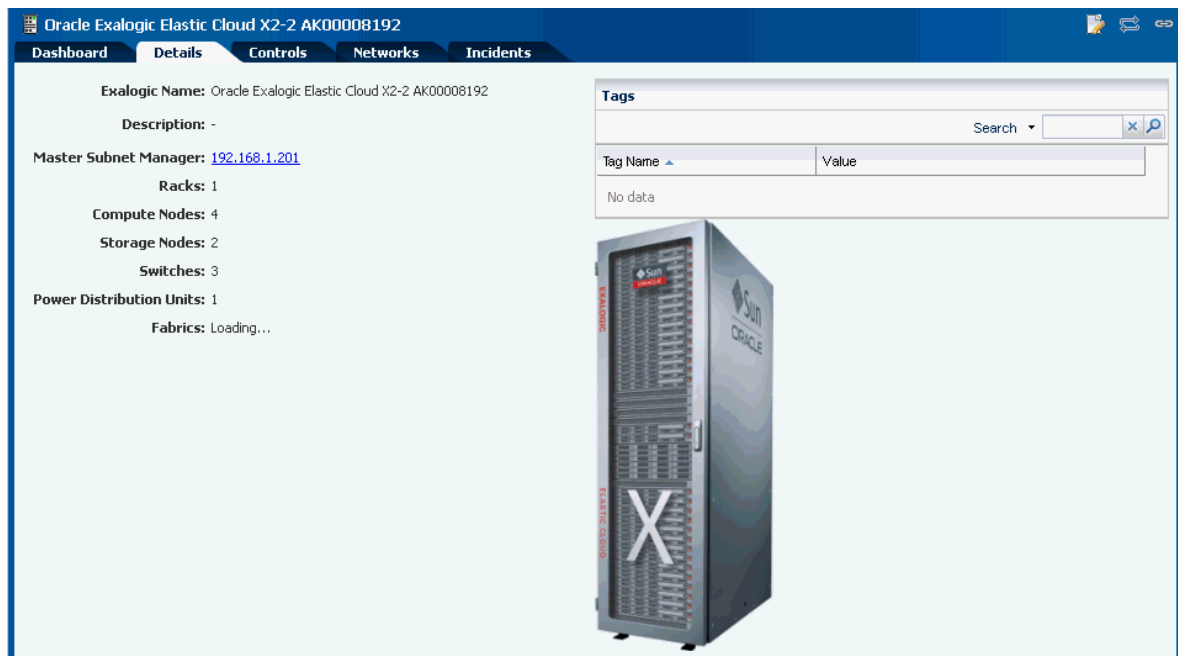
Status Pane The Status Pane displays the total unassigned incidents in chart format as well as the recent incidents encountered by the Exalogic system. In the Recent Incidents section, the following details are displayed:

- Severity – Displays the severity of the incident.
- State – Displays the state of the incident.
- ID – Displays the incident ID.
- Description – Displays the incident description.
- Source – Displays the source of the incident.
- Owner – Displays the name of the owner who reported the incident.
- Creation Date – Displays the date on which the incident was created.

Details Tab

In the Details Tab, the following details are displayed:

Figure 20–8 Details Tab



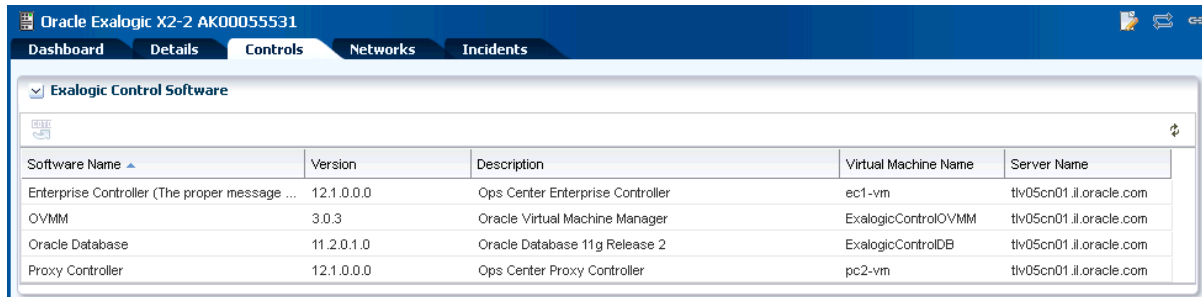
- Exalogic Name
- Description
- Model
- Number of Racks
- Number of Compute Node
- Number of Storage Nodes
- Number of Switches

In the Tags pane, the tag names and their respective values are displayed. You can search for a particular tag using the Search feature.

Controls Tab

The Controls tab displays the Exalogic Control Software details including the database, Oracle VM Manager, Proxy Controllers, and Enterprise Controller. The table displays the name of the control component, version number of the control component, description, virtual machine name, and server name.

Figure 20–9 Controls Tab



Software Name	Version	Description	Virtual Machine Name	Server Name
Enterprise Controller (The proper message ...	12.1.0.0.0	Ops Center Enterprise Controller	ec1-vm	tlv05cn01.il.oracle.com
OVMM	3.0.3	Oracle Virtual Machine Manager	ExalogicControlOVMM	tlv05cn01.il.oracle.com
Oracle Database	11.2.0.1.0	Oracle Database 11g Release 2	ExalogicControlDB	tlv05cn01.il.oracle.com
Proxy Controller	12.1.0.0.0	Ops Center Proxy Controller	pc2-vm	tlv05cn01.il.oracle.com

You can select a control component from the table and double-click it to view the Dashboard of the respective control component. The values in the table give an idea as to where the Exalogic control stack runs so that the user can quickly navigate to that compute node just by clicking on it.

Use the Refresh icon on the Exalogic Control Software pane to refresh the list as and when needed.

Networks Tab

The Networks Tab displays the Infrastructure Networks table and Network Connectivity table.

Figure 20–10 Networks Tab

Infrastructure Networks

Network Name	Network CIDR	Partition Key	IP Range	Roles
EolB-external-mgmt	10.242.0.0/21	32773	-	-
IPolB-admin	192.168.20.0/24	32769	-	-
IPolB-ovm-mgmt	192.168.23.0/24	32772	-	-
IPolB-storage	192.168.21.0/24	32770	-	-
IPolB-virt-admin	172.16.0.0/16	32771	-	-
el01-eth-admin	192.168.1.0/24	-	-	-

Network Connectivity

Network Name	Asset Type	IPolB-virt-admin	IPolB-admin	IPolB-ovm-mgmt	IPolB-storage	EolB-external-mgmt	el01-eth-admin
Rack: Rack: Oracle Exalogic X2-2 AK00008192 (10)							
192.168.1.115		-	192.168....	-	192.168.21.9	-	192.168.1.15,192.168...
192.168.1.116		-	192.168....	-	192.168.21.9	-	192.168.1.15,192.168...
192.168.1.200		-	-	-	-	-	192.168.1.200
192.168.1.201		-	-	-	-	-	192.168.1.201
192.168.1.202		-	-	-	-	-	192.168.1.202
el01.cn05.localdomain		172.16....	192.168....	192.168.23....	192.168.21.45	-	192.168.1.5
el01.cn06.localdomain		172.16....	192.168....	192.168.23.46	192.168.21.46	-	192.168.1.106,192.16...
el01.cn07.localdomain		172.16....	192.168....	192.168.23....	192.168.21.47	-	192.168.1.7,192.168...
el01.cn08.localdomain		172.16....	192.168....	192.168.23.48	192.168.21.48	-	192.168.1.8,192.168...
ucfc1z3j49pdu1		-	-	-	-	-	192.168.1.210

Infrastructure Networks Table

The Infrastructure Networks table displays the infrastructure networks that are used inside the Exalogic system for communication between the Exalogic control components and cloud virtual machines.

- Network Name - displays the name of the managed network.
- Network CIDR - displays the type of the asset.
- Partition Key - displays the partition of the port.
- IP Range - specifies the minimum and maximum boundaries of the IP addresses assigned to the network.
- Roles - displays the role.

Network Connectivity Table

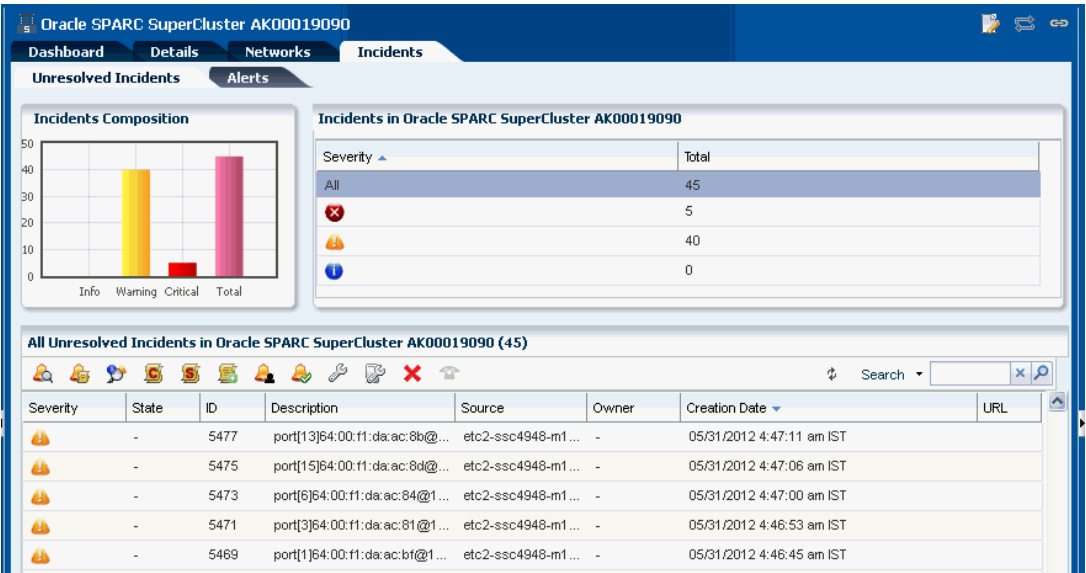
The Network Connectivity Table displays the infrastructure networks with IPs assigned to the individual hardware components of Exalogic system.

- Network Name
- Asset Type

Incidents Tab

The Incidents Tab displays all the unassigned incidents and alerts reported in the Exalogic system.

Figure 20–11 Incidents Tab



Permitted Actions in the Exalogic System View

The following are the permitted actions that can be performed in the Exalogic system view. Refer to *Oracle Enterprise Manager Ops Center Feature Reference Guide* for steps on how to perform these actions.

- Create Rack
- Edit Tags
- Add Assets
- Find Assets

Viewing the Exalogic System Rack

This section describes how to view the Exalogic System rack, visualization of the rack physical layout, aggregated rack components, energy data, and other rack details. As an administrator, you can drill down to an asset contained in the rack (server, storage node, switch, PDU) or even further.

To view the Exalogic system rack, perform the following steps:

1. On the Navigation pane, under Assets, select Exalogic Systems.
2. Select a rack that you want to view.

The Dashboard, Details, Firmware, Incidents, Charts, Energy, and Jobs tabs are displayed in the center pane.

Figure 20–12 Tabs of the Exalogic System Rack



- Dashboard Tab – Displays the summary, membership graph, and status of the rack.

- Details Tab – Displays the rack info with a photorealistic view of the rack.
 - Firmware Tab – Displays the name, description, and version of the firmware for each component and the slot for each component
 - Incidents Tab – Displays all unresolved incidents in the rack.
 - Charts 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.
 - 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.
 - Jobs Tab – Displays the current and historical jobs.
3. Select any of the tabs to view more details.

Refer to [Hardware Monitoring](#) for information on the Tabs and Views of the rack.

Photorealistic View of the Rack

The photorealistic view represents the physical rack so that you can see the placement of the rack and its components.. The front and rear views of the rack are displayed in this view. All slots and the respective assets are displayed. Positions within the rack are displayed in a 3-dimensional view.


Each asset in the rack is represented by an image. The health status of assets such as OK, Warning, or Critical is displayed in the form of colored lights as seen on the physical rack itself. OK status is identified by the green color. Warning and critical status are identified by the yellow color. Hover the mouse over the slots in the rack and view details about the assets such as slot number, asset name and description, type of asset, model number of the asset, and its health status.

A question mark in any of the slots indicates a discrepancy from the baseline check. See [Baseline Check](#) for more information.

To see the photorealistic view of the rack, perform the following steps:

1. On the Navigation pane, under Assets, select Exalogic Systems.
2. Select a rack that you want to view.
3. In the Center pane, click the Details tab.

A photorealistic view of the rack is displayed.

Figure 20–13 Photorealistic View of the Rack

Permitted Actions in Rack View

The following are the permitted actions that can be performed in the Exalogic system Rack view. Refer to *Oracle Enterprise Manager Ops Center Feature Reference Guide* for steps on how to perform these actions.

- Create Rack
- Delete Rack
- Place/Remove Assets in Rack
- Place/Remove PDUs in Rack
- Edit Tags
- Add Asset to Group
- Add Assets
- Find Assets

Creating and Viewing Reports

In general, using reports, you can view the configuration for each rack, including the asset details related to the firmware. Reports provide information about your assets,

such as job history, firmware, OS updates, and problems. Reports are created in PDF and CSV formats. Reports can be exported or used to launch jobs on targeted assets.

To create a Report, see *How to Create an Exalogic Report* in the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

You can either view existing reports that have been created or you can create a new report for yourself and view it.

In this section, the Exalogic System Report is described in detail. The following sections of the report are described below:

- [Report Parameters](#)
- [Summary Table](#)
- [System Control Software Table](#)
- [Compute Nodes Table](#)
- [Switches Table](#)
- [Storage Appliances Table](#)
- [Power Distribution Units Table](#)
- [Validation Table](#)

Report Parameters

Report Parameters displays the name and model of the Exalogic system. It also displays the number of assets in the rack by their types, such as compute nodes, switches, and storage nodes in the system.

Figure 20–14 Report Parameters

Report Name: SIR01-All
 Description: All Assets
 Run Date: Mon Apr 09 18:49:07 EDT 2012
 Report Type: Exalogic System Report

Report Parameters

Exalogic Name	Oracle Exalogic Elastic Cloud X2-2 AK00008192
Model	2.0.1.0.0
Racks	1
Compute Nodes	4
Switches	3
Storage Nodes	2

Summary Table

The Summary Table is a table of all tables in the report. It displays the number of records for each table.

Figure 20–15 Summary Table

Name of Detailed Table	Records
Virtual Datacenter Status Report	4
Rack Oracle Exalogic X2-2 AK00008192 - Compute Nodes	4
Rack Oracle Exalogic X2-2 AK00008192 - Switches	3
Rack Oracle Exalogic X2-2 AK00008192 - Storage Appliances	2
Rack Oracle Exalogic X2-2 AK00008192 - Power Distribution Units	1
Rack Oracle Exalogic X2-2 AK00008192 - Validation Table	12

System Control Software Table

The System Control Software Table lists the name of the software (control components – Proxy Controller, Enterprise Controller, LDAP, Oracle VM Manager, and database), version, description, server name, and virtual machine name. These are the same attributes as seen on the UI of control components.

Compute Nodes Table

The Compute Nodes Table displays the setup of the racks such as the type of asset, name of the asset (compute node), firmware version (such as BIOS, service processor, and HCA adaptor), Oracle VM Server version, and the location of the asset in the rack with the slot number where it is placed.

Switches Table

The Switches Table displays the name of the switch node, type of switch, firmware version, and location of the switch in the rack with the slot number where it is placed.

Storage Appliances Table

The Storage Appliances Table displays the name of the compute node, version of the service processor firmware, and the location of the appliance with the slot number where it is placed in the rack.

Power Distribution Units Table

The Power Distribution Units Table displays the PDUs and their firmware version.

Validation Table

The Validation Table displays the validity result of the rack. A baseline check is performed against the known schema for the rack. See [Baseline Check](#) for more information. The Expected Component column in the table displays the expected component if the placed component in the slot does not match the schema when the baseline check is performed.

Figure 20–16 Validation Table**Rack: Oracle Exalogic X2-2 AK00008192 - Validation Table**

Location	Type	Name	Validation Result	Expected Component
Slot 25	Switch	192.168.1.200	OK	
Slot 23	Switch	192.168.1.202	OK	
Slot 22	Storage	192.168.1.116	OK	
Slot 21	Storage	192.168.1.115	OK	
Slot 17	DiskShelf	192.168.1.15	OK	
Slot 16	Switch	192.168.1.201	OK	
Slot 5	Server	el01cn08.localdomain	OK	
Slot 4	Server	el01cn07.localdomain	OK	
Slot 3	Server	el01cn06.localdomain	OK	
Slot 2	Server	el01cn05.localdomain	OK	
-	PDU	ucfc1z3j49pdu1	OK	
-	-	-	MISSING	PDU, Oracle Rack Power Distribution Unit

SPARC SuperCluster T4-4

SPARC SuperCluster is an Oracle Engineered System which integrates SPARC compute nodes, a Sun ZFS Storage Appliance, InfiniBand switches, PDUs, and Exadata Storage Servers into a multi-rack system.

The SPARC SuperCluster hardware comprises of two or four SPARC T4-4 servers, three Sun Datacenter InfiniBand Switch 36 switches as InfiniBand backplane, a Cisco 4948 48-port 1Gb Ethernet Switch for external connectivity, and a Sun Rack II.

SPARC SuperCluster is available in two configurations, half rack and full rack. The half rack includes three Exadata Storage Servers and the full rack includes six Exadata Storage Servers. In addition, for additional storage capacity, connecting to an Oracle Exadata Storage Expansion Rack is supported.

Note: Oracle Sun ZFS Storage 7320 and 7420 Appliances provide a two-node cluster configuration. To discover the storage appliance, the administrative interfaces of both nodes must be private so that each node has a different static IP address. To verify that the appliance's nodes are using private administrative interfaces, you must use the appliance's user interface. For steps on how to determine if the interface is private, see *Verify the Administrative Interfaces*.

SPARC SuperCluster Configuration

The SPARC SuperCluster uses logical domains (LDMs) so that you can create specialized virtual systems within a single hardware platform. The logical domains have direct access to storage and network resources..

LDM configurations supported on SPARC SuperCluster have the following types:

- 11gR2 Database (DB) Domains
- Oracle Solaris 10 General Purpose Domains
- Oracle Solaris 11 General Purpose Domains

There are nine basic configuration options to choose from for the SPARC SuperCluster (Configurations A - I), with each configuration option containing one or more of the LDOM configurations listed above. Each configuration is based on a mirrored two-server model, so any one of the nine basic configuration options can be used in a SPARC SuperCluster Half Rack, with some restrictions. For a SPARC SuperCluster Full Rack, you can use different combinations of these nine basic configuration options, with some restrictions.

Note: In this document, it is assumed that the SPARC SuperCluster system is configured and discovered in Oracle Enterprise Manager Ops Center.

Understanding User Roles

In Oracle Enterprise Manager Ops Center, users can be assigned several roles such as Asset Admin, Cloud Admin, SuperCluster System Admin, and many more. Each of such role grants the user a set of permissions; a particular permission can be granted by more than one role such as Asset Management, Network Management, and so on.

You can add users to Oracle Enterprise Manager Ops Center from the local authentication subsystem of the Enterprise Controller's operating system. Each user can be given a different role which grants or denies access to the different functions of Oracle Enterprise Manager Ops Center. You can view the existing Users from the Users tab of the Administration section.

SuperCluster System Admin Role

The SuperCluster System Administrator is responsible for overall monitoring and management of the SPARC SuperCluster system. The SuperCluster System Administrator also has privileges to manage the Oracle Solaris Cluster, Virtual Pools, and Storages. As a SuperCluster System Administrator, you can perform most of the operations which are allowed in the context of the SPARC SuperCluster appliance system.

The SuperCluster System Administrator role is a default role highly recommended for the management of the SPARC SuperCluster appliance and tasks related to the above functionality.

Prerequisites: The user must be familiar with the use of Oracle Enterprise Manager Ops Center and be familiar with hardware management and OS management in general.

Viewing the SPARC SuperCluster System

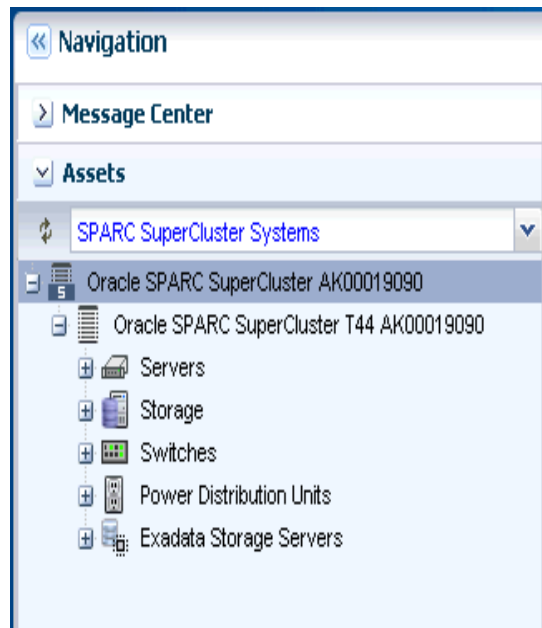
Viewing the SPARC SuperCluster system means how you view the virtual SPARC SuperCluster system using the tabs on the center pane, namely Dashboard, Details, Networks, and Incidents. You can also perform actions by clicking the respective actions on the Actions pane. The actions available on the Actions pane depend on the selected asset on the Navigation pane. The actions are also role based, not all users are allowed to perform all actions. For more information on Roles, see [Understanding User Roles](#).

To view the SPARC SuperCluster system, perform the following steps:

1. Login to Oracle Enterprise Manager Ops Center using the SuperCluster System Admin role.

2. On the Navigation pane, under Assets, select SuperCluster Systems from the drop-down list.

Figure 20–17 Navigation pane

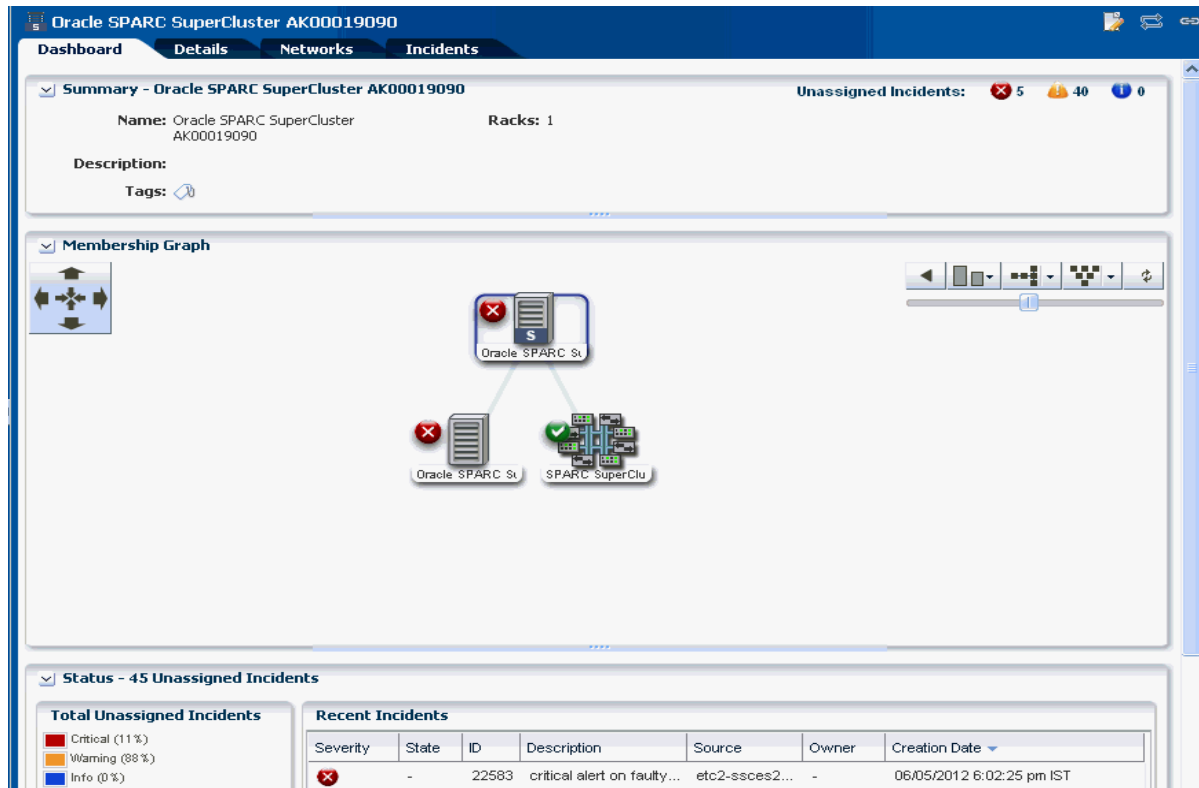


3. Select the Oracle SPARC SuperCluster system that you want to view.
The Dashboard, Details, Controls, Networks, and Incidents tabs are displayed in the center pane.
4. Click a tab to view more details.
 - **Dashboard Tab** displays the summary of the SPARC SuperCluster system, membership graph, and the status of the system.
 - **Details Tab** displays the name of the SPARC SuperCluster system, description, master subnet manager address, number of racks, compute nodes, storage nodes, Exadata Storage Servers, switches, and PDUs in the rack.
 - **Networks Tab** displays the Infrastructure Network Table and Network Connectivity Table.
 - **Incidents Tab** displays all the incidents and alerts reported in the SPARC SuperCluster System.

Dashboard Tab

The dashboard comprises of three sections - Summary, Membership Graph, and Status.

- **Summary Pane**
- **Membership Graph Pane**
- **Status Pane**

Figure 20–18 Dashboard Tab

Summary Pane The Summary section displays the name of the SPARC SuperCluster system, description or the name of the system identifier, number of racks that are part of the system, and the number of unassigned incidents. The Unassigned Incidents icon in the summary pane includes incidents resulting from hardware faults. The three icons depict Critical Incidents, Warnings, and Information Incidents respectively. These incidents are from the assets that belong to all SPARC SuperCluster systems.

Membership Graph Pane

The Membership Graph pane displays the SPARC SuperCluster system as a hierarchy of its components showing the relationship between the SPARC SuperCluster system and fabrics that are grouped in the rack. You can instantly navigate to any asset by double-clicking on the asset in the membership graph.

Using the controls on the top right of the graphic pane, you can change the view of the graph to either a horizontal or a vertical orientation. You can also refresh the view by clicking the Refresh icon. You can also change the graph depths or size of the images.

Status Pane

The Status Pane displays the total unassigned incidents in chart format as well as the recent incidents encountered by the SPARC SuperCluster system. In the Recent Incidents section, the following details are displayed:

- Severity - displays the severity of the incident.
- State - Displays the state of the incident.
- ID - Displays the incident ID.
- Description - Displays the incident description.

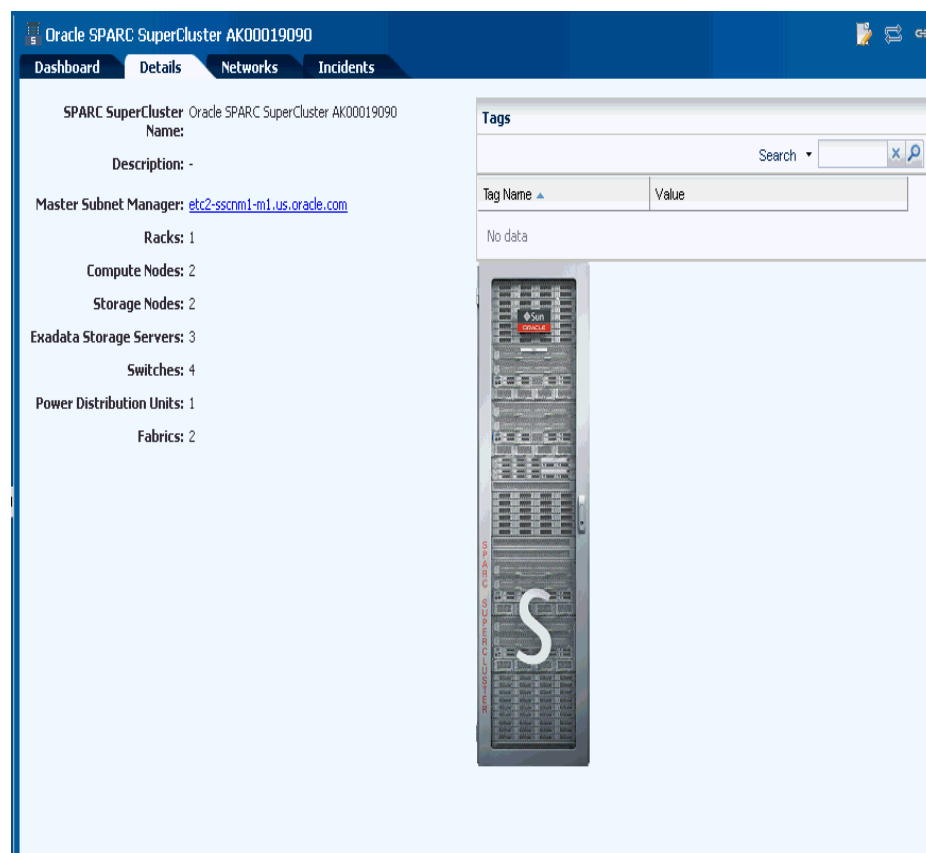
- Source – Displays the source of the incident.
- Owner – Displays the name of the owner who reported the incident.
- Creation Date – Displays the date on which the incident was created.

Details Tab

In the Details Tab, the following details are displayed:

- SPARC SuperCluster System Name
- Description
- Master Subnet Manager
- Number of Racks
- Number of Compute Nodes
- Number of Storage Nodes
- Number of Exadata Storage Servers
- Number of Switches
- Number of PDUs

Figure 20–19 Details Tab

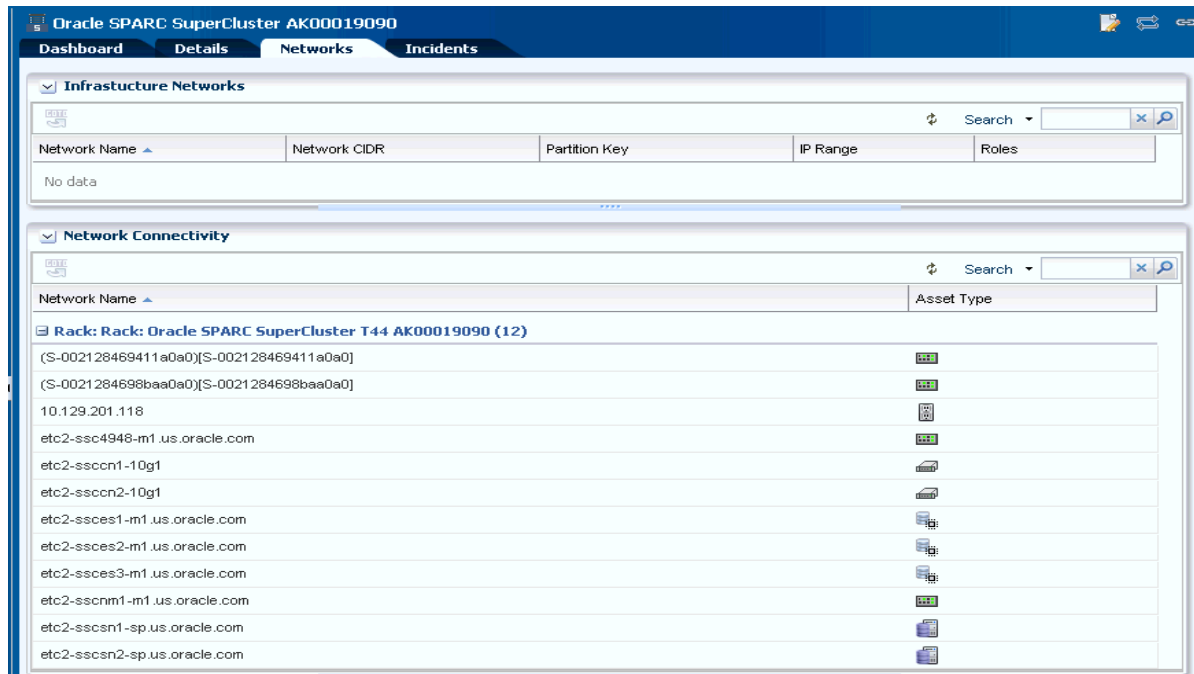


In the Tags pane, the tag names and their values are displayed. You can search for a particular tag using the Search feature.

Networks Tab

The Networks Tab displays the Infrastructure Networks Table and Network Connectivity Table.

Figure 20–20 Networks Tab



Infrastructure Networks Table

The Infrastructure Networks Table displays the infrastructure networks that are defined and used inside the SPARC SuperCluster system for communication between the SPARC SuperCluster control components.

- Network Name - displays the name of the managed network.
- Network CIDR - displays the type of the asset.
- Partition Key - displays the partition of the port.
- IP Range - specifies the minimum and maximum boundaries of the IP addresses assigned to the network.
- Roles - displays the role.

Network Connectivity Table

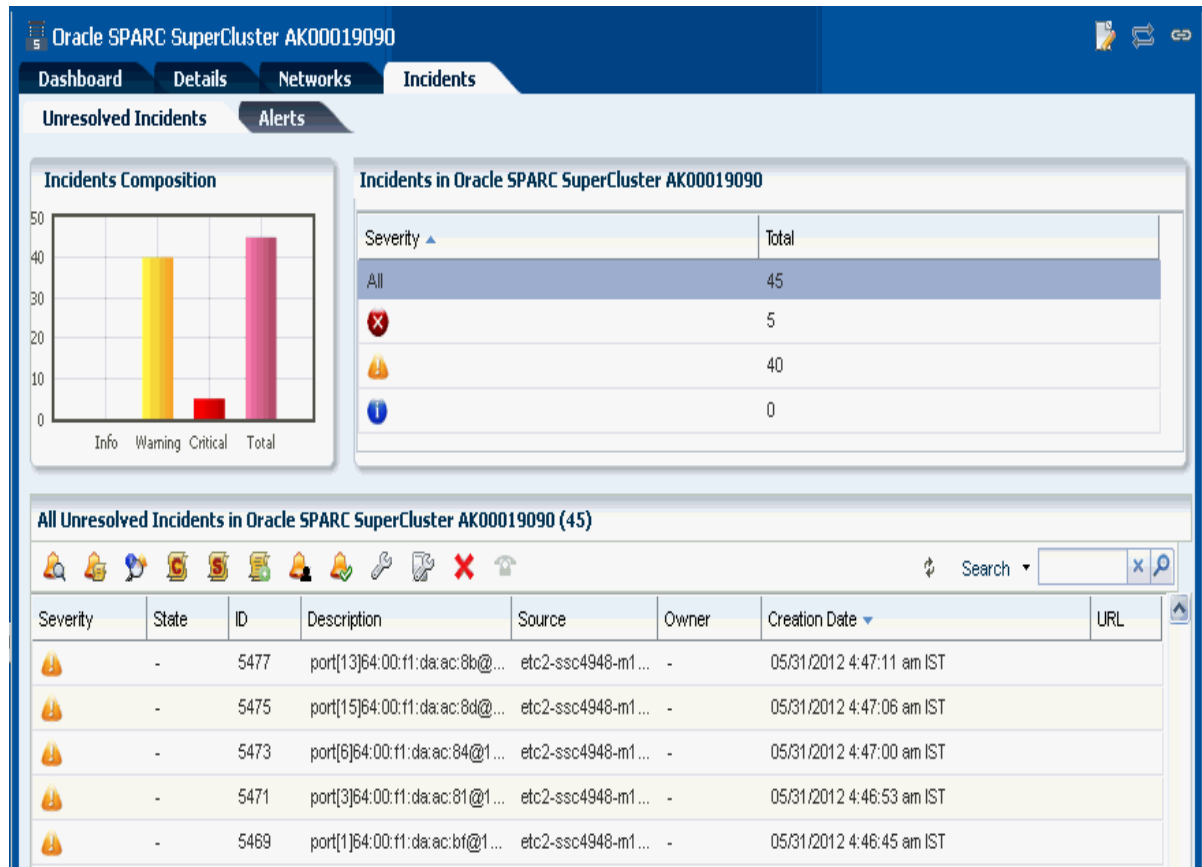
The Network Connectivity Table displays the infrastructure networks with IPs assigned to the individual hardware components of the SPARC SuperCluster system.

- Network Name
- Asset Type

Incidents Tab

The Incidents Tab displays all the unresolved incidents and alerts reported in the SPARC SuperCluster System.

Figure 20–21 Incidents tab



Permitted Actions in the SPARC SuperCluster System View

You can perform the following actions.

To see the procedures for completing these tasks, refer

Oracle Enterprise Manager Ops Center Feature Reference Guide

- Create Rack
- Edit Tags
- Add Assets
- Find Assets

Viewing the SPARC SuperCluster System Rack

This section describes how to view the SPARC SuperCluster System rack, visualization of the rack physical layout, aggregated rack components, energy data, and other rack details. As an administrator, you can drill down to an asset contained in the rack (server, storage node, switch, Exadata Cells) or even further.

To view the SPARC SuperCluster System rack, perform the following steps:

1. On the Navigation pane, under Assets, select SuperCluster Systems.
2. Select a rack that you want to view.

The Dashboard, Details, Firmware, Incidents, Charts, Energy, and Jobs tabs are displayed in the center pane.

Figure 20–22 Tabs of the SuperCluster System Rack

- Dashboard Tab – Displays the summary, membership graph, and status of the rack.
 - Details Tab – Displays the rack info with a photorealistic view of the rack.
 - Firmware Tab – Displays the name, description, and version of the firmware for each component and the slot for each component
 - Incidents Tab – Displays all unresolved incidents in the rack.
 - Charts 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.
 - 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.
 - Jobs Tab – Displays the current and historical jobs.
3. Select any of the tabs to view more details of the rack.

Refer To [Hardware Monitoring](#) section in the *Hardware Management Guide* for information on the Tabs and Views of the rack.

Photorealistic View of the SPARC SuperCluster System Rack

The photorealistic view represents the physical rack so that you can see the placement of the rack and its components.. The front and rear views of the rack are displayed in this view. All slots and the respective assets are displayed. Positions within the rack are displayed in a 3-dimensional view.


Each asset in the rack is represented by an image. The health status of assets such as OK, Warning, or Critical is displayed in the form of colored lights as seen on the physical rack itself. OK status is identified by the green color. Warning and critical status are identified by the yellow color. Hover the mouse over the slots in the rack and view details about the assets such as slot number, asset name and description, type of asset, model number of the asset, and its health status.

A question mark in any of the slots indicates a discrepancy from the baseline check. See [Baseline Check](#) for more information.

To see the photorealistic view of the rack, perform the following steps:

1. On the Navigation pane, under Assets, select SuperCluster Systems.
2. Select a rack that you want to view.
3. In the center pane, click the Details tab.

A photorealistic view of the SPARC SuperCluster System rack is displayed.

Figure 20–23 Photorealistic view of SPARC SuperCluster

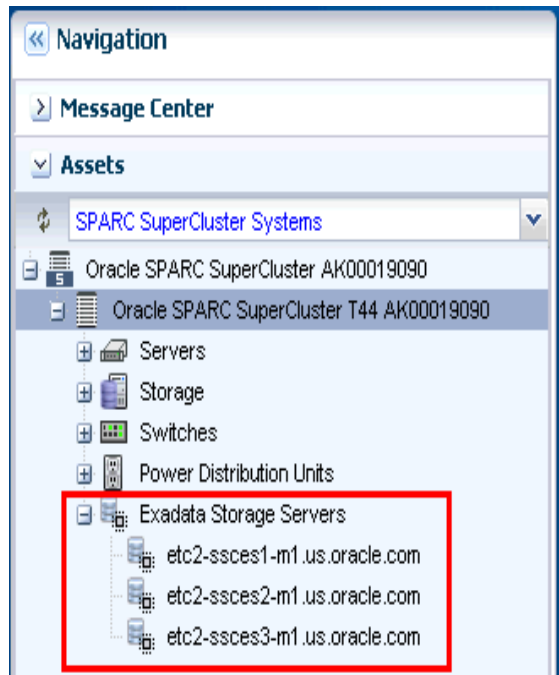
Viewing the Exadata Storage Server

The Exadata Storage Server provides storage for the SPARC SuperCluster system. It is a sub-type of the Linux server installed to be a database node.

The Exadata Storage Servers are grouped together in the Navigation pane. In the same way that you view the SPARC SuperCluster system and its rack, you can also view the Exadata Storage Servers.

To view the Exadata Storage Server, do the following:

1. On the Navigation pane, under Assets, select SuperCluster Systems from the drop-down menu.

Figure 20–24 Exadata Storage Server

2. Expand Exadata Storage Servers, and select an Exadata Storage Server. The server details are displayed on the center pane. Select the respective tabs on the center pane to view or perform actions.

Refer To [Hardware Monitoring](#) section in the *Hardware Management Guide* for information on the different views and tasks that can be performed on the server cells.

Discovering a SPARC SuperCluster Component

Manual discovery is necessary if a component has been replaced or some properties of a component has changed autonomously. The component can either be a compute node, storage, or switch.

To manually discover a SPARC SuperCluster component, do the following:

1. On the Navigation pane, select Plan Management.
2. Select Profiles and Policies, then select Discovery.

The available discovery profiles are listed in the center pane.

3. Select the required stored discovery profile and edit the profile and metadata if required.
4. Click Finish. The job is run and the component is discovered.

Add Assets to the SPARC SuperCluster Rack

After the asset is discovered, the new asset moves to the available asset poll. The asset must be added to the SPARC SuperCluster rack.

To add an asset to the SPARC SuperCluster rack, perform the following steps:

1. On the Navigation pane, under Assets, select SuperCluster Systems from the drop-down list.

2. Select the rack to which you want to add the asset.
3. On the center pane, click the Details tab. A photorealistic view of the rack is displayed.
4. Click the Add/Remove Asset in the rack.
5. Search for the right type of asset.
6. Assign or Remove the asset in the rack.
7. Save the rack changes. A rack modification job is created.

After the rack modification job is successful, the manually discovered asset appears in the SPARC SuperCluster rack.

Creating and Viewing Reports

In general, using reports, you can view the rack setup for each of the rack within the system including the asset details related to the rack. Reports provide information about your assets, such as job history, firmware, OS updates, and incidents. Reports are created in PDF and CSV formats. You can export reports or use reports to launch jobs on targeted assets.

To create a Report, see the How To library at http://docs.oracle.com/cd/E27363_01/nav/howto.htm.

You can either view existing reports that have been created or you can create a new report for yourself and view it.

In this section, the SPARC SuperCluster System Report is described in detail. The following sections of the report are described below:

- [Report Parameters](#)
- [Summary Table](#)
- [Compute Nodes Table](#)
- [Switches Table](#)
- [Storage Appliances Table](#)
- [Exadata Storage Servers Table](#)
- [Power Distribution Units Table](#)
- [Validation Table](#)

Report Parameters

Report Parameters displays the name and model of the SPARC SuperCluster system. It also displays the number of assets in the rack by their types, such as compute nodes, switches, storage nodes, and Exadata Storage Servers present in the system.

Summary Table

The Summary Table represents a table of contents of all tables in the report. It displays the number of records for each table.

Compute Nodes Table

The Compute Nodes Table displays the setup of the racks such as name of the asset (compute node), firmware version, hypervisor version, OBP version, LDOM manager version, and the location of the asset in the rack with the slot number where it is placed.

Switches Table

The Switches Table displays the name of the switch node, type of switch (infiniband or ethernet), firmware version, and location of the switch in the rack with the slot number where it is placed.

Storage Appliances Table

The Storage Appliances Table displays the name of the compute node, version of the service processor firmware, and the location of the appliance with the slot number where it is placed in the rack.

Exadata Storage Servers Table

The Exadata Storage Servers table displays the name of the compute node, version of the service processor firmware, and the location of the appliance with the slot number where it is placed in the rack.

Power Distribution Units Table

The Power Distribution Units Table displays the name and firmware version of the PDUs.

Validation Table

The Validation Table displays the location, type, name, validation result, and expected components of the rack. A baseline check is performed against the known schema for the rack. See [Baseline Check](#) for more information. The expected component column in the table displays the expected component if the placed component in the slot does not match the schema when the baseline check was performed.

Related Resources

For instructions on performing actions or to learn more about Oracle Engineered Systems, see the following resources.

- *Enabling Single Root Input Output Virtualization in Exalogic Elastic Cloud* at http://docs.oracle.com/cd/E27363_01/doc.121/e35262/toc.htm.
- *Discovering and Managing a Power Distribution Unit* at http://docs.oracle.com/cd/E27363_01/doc.121/e27338/toc.htm.
- *Creating a SPARC SuperCluster Report* at http://docs.oracle.com/cd/E27363_01/doc.121/e29308/toc.htm.

Part VII

Reference

Part VII contains the following chapters:

- [Appendix A, "Asset Attributes"](#)
- [Appendix B, "Oracle Solaris Cluster"](#)
- [Appendix C, "Logs and Directories"](#)
- [Appendix D, "JumpStart Enterprise Toolkit"](#)
- [Appendix E, "Icons"](#)
- [Appendix F, "Oracle Enterprise Manager Ops Center API"](#)
- [Glossary](#)

Asset Attributes

Oracle Enterprise Manager Ops Center records the attributes of the assets it manages. You can use these attributes to administer groups. You can use the time zone values to adjust the reporting of assets.

Monitoring Attributes

The attributes that are used for asset monitoring vary across asset types. The list of attributes for each asset is available in a set of Javadocs within the SDK package. This package is located in the following directory:

- Linux OS: *Installation directory*/xvmoc_full_bundle/Linux_i686/Product/components/packages/sun-xvmoc-sdk-12.0.0.noarch.rpm
- Oracle Solaris OS for SPARC: *Installation directory*/xvmoc_full_bundle/SunOS_sparc/Product/components/packages/SUNWxvmoc-sdk.pkg
- Oracle Solaris OS for x86: *Installation directory*/xvmoc_full_bundle/SunOS_i386/Product/components/packages/SUNWxvmoc-sdk.pkg

This package installs the Javadocs in the `/opt/sun/xvm/sdk/xvm_oc/doc/javadoc/` directory. Open the Javadoc file for an asset type to view the attributes for that asset type.

Grouping Attributes

[Table A-1](#) lists the attributes that you can use in group rules.

Table A-1 Group Asset Attributes

Attribute	Supported Values
Appliance IP Address	Any IP address
Appliance Name	Any string
Architecture	na
Asset Provisioning Supported	na

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Asset Type	GlobalZone NonGlobalZone LDomHost LDomGuest OperatingSystem Server Network ServerContainer MSeriesChassis VirtualPool SolarisCluster SolarisClusterNode SolarisClusterZoneClusterGroup
Auto Boot	on off
Capability Entry	na
Control Component Info Description	na
Control Component Info Version	na
Control Component Server Hostname	na
Control Web Console URI	na
CPU Architecture	POWER_PC SPARC SPARC-SUN4V X86
CPU Core Count	Any number
CPU Info Thread Count	Any number
CPU Mode	compatible auto
CPU Model	Any string
CPU Speed	Any number
CPU Type	Any string
Debug User SSH Access	TRUE FALSE
Description	Any string
Diag Level	max none min
DNS Configured	TRUE FALSE
DNS Domain Name	Any string

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Domain Degredation Policy	fru system xsb
Ethernet Port Fabric Tag	
Ethernet Port Info IP Address	Any IP address
Ethernet Port Infos Management	TRUE FALSE
Fabric Type	
Firmware Info Type	Any string
Firmware Description	Any string
FirmwareInfo Compliant	TRUE FALSE
FirmwareInfo Provider	Any string
Firmware Version	Any string
Hard Disk Size (Bytes)	Any number
Host ID	Any string
Hostname	Any string
HTTP Enabled	na
IbP Key	na
In Ucast Packets	na
Interface Info Enabled	TRUE FALSE
Interface Info Fabric Tag	
Interface Info Subnet Mask	Any number
Interface Infos Mac Address	Any number
Interface IP Address	Any IP address
IP Address	Any IP address
IPSLicense	Any string
IPS Readme	Any string
Keep Alive Trap Time Seconds	na
LDAP Profile Name	Any string
LDAP Proxy DN	Any string
LDAP Server	Any string
Local DNS Address	na

Table A–1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Locale	C POSIX en_CA en_CA.ISO8859-1 en_CA.UTF-8 en_US en_US.ISO8859-1 en_US.ISO8859-15 en_US.ISO8859-15@euro es es_MX es_MX.ISO8859-1 es_MX.UTF-8 fr fr_CA fr_CA.ISO8859-1 fr_CA.UTF-8 iso_8859_1
Location	na
Locator Light On	TRUE FALSE
MAC Address	na
Machine Hardware Class	InfiniBand Ethernet FibreChannel EthernetGateway InfiniBandGateway Unknown TOR NEM
Maintenance Mode	na
Managed State	HIDDEN MANAGED MANAGED_OR_PROVISIONED PROVISIONING UNMANAGED UNMANAGED_OR_PROVISIONED
Management IP	na
Manufactured Service Tag	na
Manufacturer	Any string
MBean Interface Name	Any string

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Member Info Secure Mode	on off
Member Info Set Type	DOMAIN NEM SERVER
Memor Size	Any number
Memor Type	Any string
Model	Any string
Monitored	TRUE FALSE
Monitoring Policy	na
MTU	na
Multi Path Enabled	TRUE FALSE
Name	
Name Service Domain Name	Any string
Network CIDR	Any string
NFS Map ID Domain	Any string
NIS Name Server By IP Address	Any IP address
NIS Name Server By Name	Any string
Notifications Enabled	TRUE FALSE
Operator Panel Switch Status	Any string
Ops Center Version	na
OSRunning	TRUE FALSE
Out Ucast Pkts	na
PDU Modules	na
Platform	na
Port Info Address	Any string
Port Info Connected Peer	Any string
Port Connector Type	BOARD_INTERNAL CHASSIS_BLADE CHASSIS_FABRIC CXP QSFP SFP SFP_PLUS UNKNOWN

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Port Description	Any string
Port Fabric Tag	ETHERNET FIBRE_CHANNEL INFINIBAND UNKNOWN
Port Info In Discards	Any number
Port Info In Errors	Any number
PortInfos.inNUcastPkts	Any number
Port Info In Octets	Any number
Port Info Last Tick	Any number
Port Info Link Type	Any string
Port Info Local IP Address	Any string
Port Info Management	TRUE FALSE
Port Info MTU	Any number
Port Info Node Address	na
Port Info Oper Status	Any string
Port Info Out Discards	Any number
Port Info Out Errors	Any number
Port Info Out NU CastPackets	Any number
Port Info Out Octets	Any number
Port Info Peer Address	Any number
Port Info Port Type	HOST ROUTER SWITCH TARGET
Port Info Remote DNS Address	Any string
Port Info Remote IP Address	Any string
Port Info Speed	Any string
Port Info Type	ETHERNET FIBRE_CHANNEL INFINIBAND POWER SERIAL VLAN UNKNOWN
Port Infos In Unknown Proto	na
Port Infos NUcast Packets	na

Table A–1 (Cont.) Group Asset Attributes

Attribute	Supported Values
PoweredOn	TRUE FALSE
Product Defined Instance ID	na
Product Name	na
Product URN	na
Product Vendor	na
Product Version	na
Public UUID	na
Rack ID	na
Rack Info Asset Domain Type	na
Rack Info Asset Gear Type	na
Rack Info Asset ON	na
Rack Info Asset Status	na
Rack Info Asset User Friendly Name	na
Rack Info Assets	na
Rack Info Assets Model	na
Rack Info Assets Occupancy	na
Rack Info Position X	na
Rack Info Position Y	na
Rack Info Position Z	na
Rack Info Assets User Friendly Name	na
Rack Info Configuration	na
Rack Info Number of Slots	na
Rack Info Type	na
Rcv Error Rate	Any string
RcvThroughput	Any string
Repeat Trap Number	na
Repeat Trap Time Seconds	na
Revision	na
RoutingMode	DYN_OFF DYN_ON DYN_AUTO
Satellite	Any string
Semantic Tag Key	na
Semantic Tag Key & Value	na
Sensor Notification Enabled	TRUE FALSE

Table A–1 (Cont.) Group Asset Attributes

Attribute	Supported Values
Serial Console URL	Any string
Serial Number	na
Service Tag Agent Version	Any string
Service Tag AssetLabel	Any string
Service Tag CPU Manufacturer	Any string
Service Tag Customer Asset Tag	Any string
Service Tag Host	Any string
Service Tag Instance URN	Any string
Service Tag IP Address	Any IP address
Service Tag Platform Architecture	Any string
Service Tag Port	Any number
Service Tag Product Hierarchy	Any string
Service Tag Release	Any string
Service Tag Serial Number	Any string
Service Tag Sub-Status	Any string
Service Tag Type	NATIVE NATIVE_NOT_FOUND MANUFACTURED
SNMP Enabled	na
Solaris Non-global zone Exclusive	na
Solaris Non-Global Zone.	na
Static Route Info Set	na
Status	FAULTED NON_RECOVERABLE MAINTENANCE DECONFIGURED DEGRADED CRITICAL WARNING INFO OK UNINITIALIZED UNKNOWN UNCONFIGURED NEEDS_POWERON
Subnet Manager	TRUE FALSE
Subnet Manager Address	Any string
Subnet Mask	Any string

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
System Identifier	na
System Name	na
System Version	na
Time Zone	See Time Zone Values list below
Total Rcv Errors	Any string
Total Rcv Octets	Any string
Total Xmit Errors	Any string
Total Xmit Octets	Any string
User Friendly Description	Any string
User Friendly Name	Any string
Version	na
Virtualization Type	ldom xen zone
VLAN ID	Any string
Web Console URI	Any string
Web Console URL	Any string
Xmit Error Rate	Any string
Xmit Throughput	Any string
XSB Assignment Status	Assigned Available Unavailable
XSB Connection Status	y n
XSB Diag Status	Failed Passed Testing Unknown Unmount

Table A-1 (Cont.) Group Asset Attributes

Attribute	Supported Values
XSB Domain ID	00
	01
	02
	03
	04
	05
	06
	07
	08
	09
	10
	11
	12
	13
	14
	SP
XSB Dynamic Reconf Status	Any string
XSB ID	Any string
XSB Fault Status	Degraded
	Faulted
	Normal
XSB Info Set	na
XSB Power Status	y
	n
XSB PSB ID	Any string
XSB Incorporation Status	y
	n
XVMServerVersion	Any string

Time Zone Values

- Africa/Abidjan
- Africa/Accra
- Africa/Addis_Ababa
- Africa/Algiers
- Africa/Asmara
- Africa/Asmera
- Africa/Bamako
- Africa/Bangui
- Africa/Banjul

- Africa/Bissau
- Africa/Blantyre
- Africa/Brazzaville
- Africa/Bujumbura
- Africa/Cairo
- Africa/Casablanca
- Africa/Ceuta
- Africa/Conakry
- Africa/Dakar
- Africa/Dar_es_Salaam
- Africa/Djibouti
- Africa/Douala
- Africa/El_Aaiun
- Africa/Freetown
- Africa/Gaborone
- Africa/Harare
- Africa/Johannesburg
- Africa/Kampala
- Africa/Khartoum
- Africa/Kigali
- Africa/Kinshasa
- Africa/Lagos
- Africa/Libreville
- Africa/Lome
- Africa/Luanda
- Africa/Lubumbashi
- Africa/Lusaka
- Africa/Malabo
- Africa/Maputo
- Africa/Maseru
- Africa/Mbabane
- Africa/Mogadishu
- Africa/Monrovia
- Africa/Nairobi
- Africa/Ndjamena
- Africa/Niamey
- Africa/Nouakchott

- Africa/Ouagadougou
- Africa/Porto-Novo
- Africa/Sao_Tome
- Africa/Timbuktu
- Africa/Tripoli
- Africa/Tunis
- Africa/Windhoek
- America/Adak
- America/Anchorage
- America/Anguilla
- America/Antigua
- America/Araguaina
- America/Argentina/Buenos_Aires
- America/Argentina/Catamarca
- America/Argentina/ComodRivadavia
- America/Argentina/Cordoba
- America/Argentina/Jujuy
- America/Argentina/La_Rioja
- America/Argentina/Mendoza
- America/Argentina/Rio_Gallegos
- America/Argentina/Salta
- America/Argentina/San_Juan
- America/Argentina/San_Luis
- America/Argentina/Tucuman
- America/Argentina/Ushuaia
- America/Aruba
- America/Asuncion
- America/Atikokan
- America/Atka
- America/Bahia
- America/Bahia_Banderas
- America/Barbados
- America/Belem
- America/Belize
- America/Blanc-Sablon
- America/Boa_Vista
- America/Bogota

- America/Boise
- America/Buenos_Aires
- America/Cambridge_Bay
- America/Campo_Grande
- America/Cancun
- America/Caracas
- America/Catamarca
- America/Cayenne
- America/Cayman
- America/Chicago
- America/Chihuahua
- America/Coral_Harbour
- America/Cordoba
- America/Costa_Rica
- America/Cuiaba
- America/Curacao
- America/Danmarkshavn
- America/Dawson
- America/Dawson_Creek
- America/Denver
- America/Detroit
- America/Dominica
- America/Edmonton
- America/Eirunepe
- America/El_Salvador
- America/Ensenada
- America/Fort_Wayne
- America/Fortaleza
- America/Glace_Bay
- America/Godthab
- America/Goose_Bay
- America/Grand_Turk
- America/Grenada
- America/Guadeloupe
- America/Guatemala
- America/Guayaquil
- America/Guyana

- America/Halifax
- America/Havana
- America/Hermosillo
- America/Indiana/Indianapolis
- America/Indiana/Knox
- America/Indiana/Marengo
- America/Indiana/Petersburg
- America/Indiana/Tell_City
- America/Indiana/Vevay
- America/Indiana/Vincennes
- America/Indiana/Winamac
- America/Indianapolis
- America/Inuvik
- America/Iqaluit
- America/Jamaica
- America/Jujuy
- America/Juneau
- America/Kentucky/Louisville
- America/Kentucky/Monticello
- America/Knox_IN
- America/La_Paz
- America/Lima
- America/Los_Angeles
- America/Louisville
- America/Maceio
- America/Managua
- America/Manaus
- America/Marigot
- America/Martinique
- America/Matamoros
- America/Mazatlan
- America/Mendoza
- America/Menominee
- America/Merida
- America/Mexico_City
- America/Miquelon
- America/Moncton

- America/Monterrey
- America/Montevideo
- America/Montreal
- America/Montserrat
- America/Nassau
- America/New_York
- America/Nipigon
- America/Nome
- America/Noronha
- America/North_Dakota/Center
- America/North_Dakota/New_Salem
- America/Ojinaga
- America/Panama
- America/Pangnirtung
- America/Paramaribo
- America/Phoenix
- America/Port-au-Prince
- America/Port_of_Spain
- America/Porto_Acre
- America/Porto_Velho
- America/Puerto_Rico
- America/Rainy_River
- America/Rankin_Inlet
- America/Recife
- America/Regina
- America/Resolute
- America/Rio_Branco
- America/Rosario
- America/Santa_Isabel
- America/Santarem
- America/Santiago
- America/Santo_Domingo
- America/Sao_Paulo
- America/Scoresbysund
- America/Shiprock
- America/St_Barthelemy
- America/St_Johns

- America/St_Kitts
- America/St_Lucia
- America/St_Thomas
- America/St_Vincent
- America/Swift_Current
- America/Tegucigalpa
- America/Thule
- America/Thunder_Bay
- America/Tijuana
- America/Toronto
- America/Tortola
- America/Vancouver
- America/Virgin
- America/Whitehorse
- America/Winnipeg
- America/Yakutat
- America/Yellowknife
- Antarctica/Casey
- Antarctica/Davis
- Antarctica/DumontDUrville
- Antarctica/Macquarie
- Antarctica/Mawson
- Antarctica/McMurdo
- Antarctica/Palmer
- Antarctica/Rothera
- Antarctica/South_Pole
- Antarctica/Syowa
- Antarctica/Vostok
- Arctic/Longyearbyen
- Asia/Aden
- Asia/Almaty
- Asia/Amman
- Asia/Anadyr
- Asia/Aqtau
- Asia/Aqtobe
- Asia/Ashgabat
- Asia/Ashkhabad

- Asia/Baghdad
- Asia/Bahrain
- Asia/Baku
- Asia/Bangkok
- Asia/Beirut
- Asia/Bishkek
- Asia/Brunei
- Asia/Calcutta
- Asia/Choibalsan
- Asia/Chongqing
- Asia/Chungking
- Asia/Colombo
- Asia/Dacca
- Asia/Damascus
- Asia/Dhaka
- Asia/Dili
- Asia/Dubai
- Asia/Dushanbe
- Asia/Gaza
- Asia/Harbin
- Asia/Ho_Chi_Minh
- Asia/Hong_Kong
- Asia/Hovd
- Asia/Irkutsk
- Asia/Istanbul
- Asia/Jakarta
- Asia/Jayapura
- Asia/Jerusalem
- Asia/Kabul
- Asia/Kamchatka
- Asia/Karachi
- Asia/Kashgar
- Asia/Kathmandu
- Asia/Katmandu
- Asia/Kolkata
- Asia/Krasnoyarsk
- Asia/Kuala_Lumpur

- Asia/Kuching
- Asia/Kuwait
- Asia/Macao
- Asia/Macau
- Asia/Magadan
- Asia/Makassar
- Asia/Manila
- Asia/Muscat
- Asia/Nicosia
- Asia/Novokuznetsk
- Asia/Novosibirsk
- Asia/Omsk
- Asia/Oral
- Asia/Phnom_Penh
- Asia/Pontianak
- Asia/Pyongyang
- Asia/Qatar
- Asia/Qyzylorda
- Asia/Rangoon
- Asia/Riyadh
- Asia/Riyadh87
- Asia/Riyadh88
- Asia/Riyadh89
- Asia/Saigon
- Asia/Sakhalin
- Asia/Samarkand
- Asia/Seoul
- Asia/Shanghai
- Asia/Singapore
- Asia/Taipei
- Asia/Tashkent
- Asia/Tbilisi
- Asia/Tehran
- Asia/Tel_Aviv
- Asia/Thimbu
- Asia/Thimphu
- Asia/Tokyo

- Asia/Ujung_Pandang
- Asia/Ulaanbaatar
- Asia/Ulan_Bator
- Asia/Urumqi
- Asia/Vientiane
- Asia/Vladivostok
- Asia/Yakutsk
- Asia/Yekaterinburg
- Asia/Yerevan
- Atlantic/Azores
- Atlantic/Bermuda
- Atlantic/Canary
- Atlantic/Cape_Verde
- Atlantic/Faeroe
- Atlantic/Faroe
- Atlantic/Jan_Mayen
- Atlantic/Madeira
- Atlantic/Reykjavik
- Atlantic/South_Georgia
- Atlantic/St_Helena
- Atlantic/Stanley
- Australia/ACT
- Australia/Adelaide
- Australia/Brisbane
- Australia/Broken_Hill
- Australia/Canberra
- Australia/Currie
- Australia/Darwin
- Australia/Eucla
- Australia/Hobart
- Australia/LHI
- Australia/Lindeman
- Australia/Lord_Howe
- Australia/Melbourne
- Australia/NSW
- Australia/North
- Australia/Perth

- Australia/Queensland
- Australia/South
- Australia/Sydney
- Australia/Tasmania
- Australia/Victoria
- Australia/West
- Australia/Yancowinna
- Brazil/Acre
- Brazil/DeNoronha
- Brazil/East
- Brazil/West
- CET
- CST6CDT
- Canada/Atlantic
- Canada/Central
- Canada/East-Saskatchewan
- Canada/Eastern
- Canada/Mountain
- Canada/Newfoundland
- Canada/Pacific
- Canada/Saskatchewan
- Canada/Yukon
- Chile/Continental
- Chile/EasterIsland
- Cuba
- EET
- EST
- EST5EDT
- Egypt
- Eire
- Etc/GMT
- Etc/GMT+0
- Etc/GMT+1
- Etc/GMT+10
- Etc/GMT+11
- Etc/GMT+12
- Etc/GMT+2

- Etc/GMT+3
- Etc/GMT+4
- Etc/GMT+5
- Etc/GMT+6
- Etc/GMT+7
- Etc/GMT+8
- Etc/GMT+9
- Etc/GMT-0
- Etc/GMT-1
- Etc/GMT-10
- Etc/GMT-11
- Etc/GMT-12
- Etc/GMT-13
- Etc/GMT-14
- Etc/GMT-2
- Etc/GMT-3
- Etc/GMT-4
- Etc/GMT-5
- Etc/GMT-6
- Etc/GMT-7
- Etc/GMT-8
- Etc/GMT-9
- Etc/GMT0
- Etc/Greenwich
- Etc/UCT
- Etc/UTC
- Etc/Universal
- Etc/Zulu
- Europe/Amsterdam
- Europe/Andorra
- Europe/Athens
- Europe/Belfast
- Europe/Belgrade
- Europe/Berlin
- Europe/Bratislava
- Europe/Brussels
- Europe/Bucharest

- Europe/Budapest
- Europe/Chisinau
- Europe/Copenhagen
- Europe/Dublin
- Europe/Gibraltar
- Europe/Guernsey
- Europe/Helsinki
- Europe/Isle_of_Man
- Europe/Istanbul
- Europe/Jersey
- Europe/Kaliningrad
- Europe/Kiev
- Europe/Lisbon
- Europe/Ljubljana
- Europe/London
- Europe/Luxembourg
- Europe/Madrid
- Europe/Malta
- Europe/Mariehamn
- Europe/Minsk
- Europe/Monaco
- Europe/Moscow
- Europe/Nicosia
- Europe/Oslo
- Europe/Paris
- Europe/Podgorica
- Europe/Prague
- Europe/Riga
- Europe/Rome
- Europe/Samara
- Europe/San_Marino
- Europe/Sarajevo
- Europe/Simferopol
- Europe/Skopje
- Europe/Sofia
- Europe/Stockholm
- Europe/Tallinn

- Europe/Tirane
- Europe/Tiraspol
- Europe/Uzhgorod
- Europe/Vaduz
- Europe/Vatican
- Europe/Vienna
- Europe/Vilnius
- Europe/Volgograd
- Europe/Warsaw
- Europe/Zagreb
- Europe/Zaporozhye
- Europe/Zurich
- Factory
- GB
- GB-Eire
- GMT
- GMT+0
- GMT-0
- GMT0
- Greenwich
- HST
- Hongkong
- Iceland
- Indian/Antananarivo
- Indian/Chagos
- Indian/Christmas
- Indian/Cocos
- Indian/Comoro
- Indian/Kerguelen
- Indian/Mahe
- Indian/Maldives
- Indian/Mauritius
- Indian/Mayotte
- Indian/Reunion
- Iran
- Israel
- Jamaica

- Japan
- Kwajalein
- Libya
- MET
- MST
- MST7MDT
- Mexico/BajaNorte
- Mexico/BajaSur
- Mexico/General
- Mideast/Riyadh87
- Mideast/Riyadh88
- Mideast/Riyadh89
- NZ
- NZ-CHAT
- Navajo
- PRC
- PST8PDT
- Pacific/Apia
- Pacific/Auckland
- Pacific/Chatham
- Pacific/Easter
- Pacific/Efate
- Pacific/Enderbury
- Pacific/Fakaofu
- Pacific/Fiji
- Pacific/Funafuti
- Pacific/Galapagos
- Pacific/Gambier
- Pacific/Guadalcanal
- Pacific/Guam
- Pacific/Honolulu
- Pacific/Johnston
- Pacific/Kiritimati
- Pacific/Kosrae
- Pacific/Kwajalein
- Pacific/Majuro
- Pacific/Marquesas

- Pacific/Midway
- Pacific/Nauru
- Pacific/Niue
- Pacific/Norfolk
- Pacific/Noumea
- Pacific/Pago_Pago
- Pacific/Palau
- Pacific/Pitcairn
- Pacific/Ponape
- Pacific/Port_Moresby
- Pacific/Rarotonga
- Pacific/Saipan
- Pacific/Samoa
- Pacific/Tahiti
- Pacific/Tarawa
- Pacific/Tongatapu
- Pacific/Truk
- Pacific/Wake
- Pacific/Wallis
- Pacific/Yap
- Poland
- Portugal
- ROK
- Singapore
- Turkey
- UCT
- US/Alaska
- US/Aleutian
- US/Arizona
- US/Central
- US/East-Indiana
- US/Eastern
- US/Hawaii
- US/Indiana-Starke
- US/Michigan
- US/Mountain
- US/Pacific

- US/Pacific-New
- US/Samoa
- UTC
- Universal
- W-SU
- WET
- Zulu

Oracle Solaris Cluster

Use Oracle Enterprise Manager Ops Center to install a new Oracle Solaris Cluster and to maintain existing Oracle Solaris Clusters.

Upgrading an Oracle Solaris Cluster

The cluster update procedures operate on complete clusters. Do not attempt to upgrade only a cluster node.

To upgrade a cluster to a newer version of the Oracle Solaris Cluster software, use a profile that defines the new version. Either select the target cluster and then the deployment plan or select the deployment plan and then the target. In either case, you use the Solaris Cluster Upgrade Job Wizard to specify how the deployment operates.

Before You Begin

Import the cluster profile.

See XREF or detailed procedure about importing a cluster profile.

To Upgrade the Cluster

1. Expand Plan Management in the Navigation pane.
2. In the Profiles and Policies section, click Update Profiles.
3. Select **Oracle Solaris Upgrade Job** in the Actions pane. The Oracle Solaris Upgrade Job Wizard starts.
4. Enter a name for the upgrade job.
5. Click display the cluster targets.
6. Choose one or more clusters and click Select/Add to Target List.
7. Select Dual Partition Mode.
8. Select the policy to use if a task in the upgrade cannot be completed.
9. Select the profile that you imported.
10. Review the summary of the job and click Finish. The upgrade job is submitted. To follow the progress of the job, click the View Job Details icon in the jobs pane.

Cluster Profiles

You use cluster profiles in a deployment plan to perform the following operations:

- Install Oracle Solaris Cluster software

- Upgrade Oracle Solaris Cluster software

The profiles and deployment plans rely on pre-action and post-action scripts to suspend cluster operations and save configuration information during the update process.

The following cluster-specific profiles and scripts are located in Enterprise Manager Ops Center's Local Content software library. As the Oracle Solaris Cluster software is revised, new profiles are made available to you through the Java.net site's Ops Center Cluster Profiles project. You download these new profiles and scripts, then import them into the Local Content library. See [Obtaining the Cluster Profiles and Scripts](#) for instructions to get these files.

- Profiles for Provisioning Oracle Solaris Cluster Software version 3.2u3
 - SPARC
 - * sc-3.2u3-core-sparc
 - * sc-3.2u3-manager-sparc
 - * sc-3.2u3-agents-sparc
 - x86
 - * sc-3.2u3-core-x86
 - * sc-3.2u3-manager-x86
 - * sc-3.2u3-agents-x86
- Profiles for Upgrading Version 3.2u3 to Version 3.3 or for Provisioning an Oracle Solaris Cluster Software Version 3.3
 - SPARC
 - * sc-3.3-core-sparc
 - * sc-3.3-manager-sparc
 - * sc-3.3-agents-sparc
 - x86
 - * sc-3.3-core-x86
 - * sc-3.3-manager-x86
 - * sc-3.3-agents-x86
- Scripts for pre-action and post-action
 - Post-action script for provisioning either version 3.2 or 3.3: SolarisCluster-Post.ksh
 - Post-action script for upgrading version 3.2 to 3.3: SolarisCluster3.21109-Post
 - Pre-action script for upgrading version 3.2 to 3.3: SolarisCluster3.21109-Pre

Obtaining the Cluster Profiles and Scripts

To obtain profiles and scripts for a new release of Oracle Solaris s Cluster software, download them to the system that is running the Enterprise Controller.

Note: When you download, transfer, upload, or import these files, do not change the file name. Keep the same name throughout the procedure.

To Obtain the Cluster Profiles and Scripts

1. Go to <http://java.net/projects/oc-cluster-profiles>
2. Click the Downloads button.
3. Select the folder for the profiles you want:
 - To provision Oracle Solaris Cluster Software version 3.2u3, select the Version-3.2u3 folder.
 - To upgrade version 3.2u3 to version 3.3, select the Version-3.3 folder.
 - To provision Oracle Solaris Cluster Software version 3.3, select the Version-3.3 folder.
4. In the version folder, select the platform you want, either SPARC or x86.
The platform folder contains three files: Agents, Core, and Manager. You need all three files to complete upgrade or provision operations.
5. Use the browser window's Save feature to save each file, keeping its file name.
6. Return to the home page for the project and select the Action-Scripts folder.
7. Select the scripts that complete the core profiles:
 - To upgrade version 3.2u3 to version 3.3, select the Pre-Upgrade and Post-Upgrade scripts.
 - To provision version 3.2u3 or version 3.3, select the Post-Provision script.
8. Upload the profiles and scripts, according to [Uploading the Cluster Profiles and Scripts](#).

Uploading the Cluster Profiles and Scripts

The files containing the profiles and scripts must reside on the same system that is running the browser and they must reside in the same directory. This operation uploads all files in the selected directory.

1. Move the profiles and scripts to the appropriate system and directory.
2. Expand Libraries in the Navigation pane.
3. Click Local Content in the Solaris/Linux OS Updates library.
4. Click **Bulk Upload Packages and Patches** in the Actions pane.
5. Click Distribution to select the distribution that applies to these files.
6. Select Upload from Directory.
7. Specify the path to the directory or click Browse to locate and select it.
All files in the directory and its subdirectories are uploaded.
8. Click Submit.
9. Import the profiles into Enterprise Manager Ops Center, according to [Importing Cluster Profiles and Scripts](#).

Importing Cluster Profiles and Scripts

Before You Begin

- To import profiles, use the browser on the same system or transfer the profiles to the system on which you are running the browser.
- To import scripts, transfer the files to the Enterprise Controller's system.

Note: When you download, transfer, upload, or import these files, do not change the file name. Keep the same name throughout the procedure.

To Import a Cluster Profile

1. Expand Plan Management in the Navigation pane.
2. Click Update Profiles.
3. Click the Import Profile icon.
4. For each profile and script:
 - a. Enter the name of the file you downloaded. Do not change the name.
 - b. Click Import.
5. Expand Libraries in the Navigation pane.
6. Select Solaris/Linux OS Updates.
7. Select Local Content.
8. Click **Upload Local Action**.
9. For each profile:
 - a. Enter the name of the profile as you want it to be displayed in the Plan Management section.
 - b. Select the distribution of the OS that uses the profile.
 - c. Enter the name of the file that contains the profile. Do not change the name.
 - d. Click Upload.
10. For each script:
 - a. Select action type, the matching parent category, and the OS distribution.
 - b. Enter the name of the file that contains the script. Do not change the name.
 - c. Click Upload.
11. For provisioning profiles, edit the core profiles to include the Post-Provision scripts, according to [Editing the Core Profile for Provisioning](#)

To upgrade a cluster, no modification is needed because the scripts are available to the upgrade profile after the import operation.

Editing the Core Profile for Provisioning

1. Expand Plan Management in the Navigation pane.
2. Click Update Profiles.
3. In the OS Update Profiles table, select the core profile.

4. Click the Edit Profile icon.
5. In the Edit OS Update Profile window, navigate from Local to either Post-actions or Pre-action.
6. Select the script and click Required. For example, to edit the core profile for provisioning, navigate to Post-actions, select SolarisCluster-Post.ksh, and click Required. The file is now listed in the Profile Contents section.
7. Click the Saved as Named Profile button.

Logs and Directories

Each Enterprise Manager Ops Center component has some auditing capability. Follow audit advice in this document and regularly monitor audit records.

Enterprise Manager Ops Center performs each action as a job. The details of a job show the order of operations in the job and the managed assets that were targets of the job. You can view the details of a job from either the browser or the command-line interface. Each job is stored until it is deleted explicitly.

In addition to the jobs record, log files can be a source of activity records. Logs are written during operations and can provide additional detail about system activity. Log files are protected by file permissions and therefore requires a privileged user to access them.

Installation

- Log of the most recent installation or uninstallation:
`/var/tmp/installer.log.latest`
- Log of previous installation or uninstallation operations:
`/var/tmp/installer.log.xxxx`
- Log of an agent installation: `/var/scn/install/log`

General Events

- Messages: `/var/adm/messages*`
- Sessions with IP address and port: `/var/opt/sun/xvm/logs/audit-logs*`
- BUI: `/var/opt/sun/xvm/logs/emoc.log`
- Actions of the BUI and remote clients on the Enterprise Controller:
 - On Oracle Solaris: `/var/cacao/instances/oem-ec/audits/`
 - On Linux: `/var/opt/sun/cacao/instances/oem-ec/audits/`
- Events between controllers and agents:
 - On an Oracle Solaris Enterprise Controller:
`/var/cacao/instances/oem-ec/logs`
 - On a Linux Enterprise Controller:
`/var/opt/sun/cacao/instances/oem-ec/logs`
 - On each Oracle Solaris Proxy Controller:
`/var/cacao/instances/scn-proxy/logs/cacao.n`

- On each Linux Proxy Controller:
`/var/opt/sun/cacao/instances/scn-proxy/logs/cacao.n`

High Availability

When Oracle Enterprise Manager Ops Center is in a High Availability configuration, each Enterprise Controller is a Clusterware node. The Clusterware resource activity is logged and the following log file is updated each time the active Enterprise Controller's resource action script's check() function is executed. The default interval is 60 seconds.

On Oracle Solaris: `/var/opt/sun/xvm/ha/EnterpriseController.log`

Software Updates

The Software Update component has its own server with its own logs. The following logs provide information on the activity for this server:

- Audit Log
 - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/audit.log`
 - On Linux: `/usr/local/uce/server/logs/audit.log`
- Errors
 - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/error.log`
 - On Linux: `/usr/local/uce/server/logs/error.log`
 - Log of errors in download jobs: `/opt/SUNWuce/server/logs/SERVICE_CHANNEL/error.log`
- Job Log
 - On Oracle Solaris: `/var/opt/sun/xvm/uce/var.opt/server/logs/job.log`
 - On Linux: `/usr/local/uce/server/logs/job.log`

Agents

- `/var/scn/update-agent/logs` directory.
- `/var/opt/sun/xvm/logs`

Local Database

- On the Enterprise Controller:
 - For installation events:
`/var/opt/sun/xvm/oracle/cfgtoollogs/dbca/OCDB/*`
`/var/tmp/installer.log.latest`
 - For operational events, reported by the `ecadm sqlplus` utility:
`/var/opt/sun/xvm/oracle/diag/rdbms/ocdb/OCDB/alert/log.xml.*`
`/var/opt/sun/xvm/oracle/diag/rdbms/ocdb/OCDB/trace/alert_OCDB.log.*`
`/var/opt/sun/xvm/oracle/diag/tnslsnr/hostname/oclistener/alert/log.xml.*`
`/var/opt/sun/xvm/oracle/diag/tnslsnr/hostname/oclistener/trace/listener.log.*`

- For schema changes:
 - `/var/opt/sun/xvm/log/satadmsqlplus.log`
 - `/var/opt/sun/xvm/logs/alter_oracle_schema.out`
 - `/var/opt/sun/xvm/logs/alter_oracle_storage.out`
- For backup, restore, and migrate operations:
 - `/var/tmp/sat-backup-<`date +%F-%T`>.log`
 - `/var/tmp/sat-restore-<`date +%F-%T`>.log`
 - `/var/opt/sun/xvm/logs/migrate.log`
- On the Proxy Controller: `/var/opt/sun/xvm/proxydb/*`
- On each agent: `/var/opt/sun/xvm/agentdb/*`

Upgrade

The log of upgrade actions for the Enterprise Controller and its co-located proxy controller is in the `/var/scn/update-saved-state/update_satellite_bundle_12.1.n.xxxx/updatelog.txt` file.

The log of upgrade actions for a Proxy Controller that is not co-located is in the `/var/scn/update-saved-state/update_proxy_bundle_12.1.n.xxxx/updatelog.txt` file.

After an upgrade, the state of the database is saved in the `/var/scn/update-saved-state/update_satellite_bundle_12.1.n.xxxx/dblogs` folder.

Event Logs for Jobs

Event logs for jobs are generated by tasks as they progress. See [Viewing Jobs](#) for instructions on viewing job event logs.

Controlling the Number of Common Agent Container Log Files

The Common Agent Container cacao is a common Java container for JDMX/JMX management and handles the interactions between controllers and agents. All events are recorded in the cacao log files and events that are above the level of INFO are also logged in the syslog. You can view the contents of the current log file using the Ops Center BUI.

- On an Oracle Solaris Enterprise Controller:
 - `/var/cacao/instances/oem-ec/logs/cacao.n`
- On a Linux Enterprise Controller:
 - `/var/opt/sun/cacao/instances/oem-ec/logs/cacao.n`
- On each Oracle Solaris Proxy Controller:
 - `/var/cacao/instances/scn-proxy/logs/cacao.n`
- On each Linux Proxy Controller:
 - `/var/opt/sun/cacao/instances/scn-proxy/logs/cacao.n`

The maximum file size is 1 MB. When the limit is reached, the current log file is closed and a new one created. The default number of log files is three. You can change the

number of log files Ops Center retains, using the Common Agent Container's management utility, `cacaoadm`.

To view the current number of log files maintained for the Enterprise Controller, issue the following command on the system where the Enterprise software is running:

```
# cacaoadm get-param log-file-count -i oem-ec
log-file-count=3
```

To view the number of log files maintained for a Proxy Controller, issue the following command on the system where the proxy controller software is running:

```
# cacaoadm get-param log-file-count -i scn-proxy
log-file-count=3
```

To change the number of log files on the Enterprise Controller, a Proxy Controller, or both:

1. Verify that there are no active jobs.
2. Stop the Common Agent Container service on the Enterprise Controller.
 - On Oracle Solaris:

```
# /opt/SUNWxvmoc/bin/satadm stop -w -v
```
 - On Linux:

```
# /opt/sun/xvmoc/bin/satadm stop -w -v
```

Stop the Common Agent Container service on a Proxy Controller:

 - On Oracle Solaris:

```
# /opt/SUNWxvmoc/bin/proxyadm stop -w -v
```
 - On Linux:

```
# /opt/sun/xvmoc/bin/proxyadm stop -w -v
```
3. Specify the maximum number of log files to be retained in addition to the current log file. In the following example, the count of 10 specifies that nine log files of previous events are retained in addition to the log file for current events. On the Enterprise Controller:

```
# cacaoadm set-param log-file-count=10 -i default
```

On a Proxy Controller:

```
# cacaoadm set-param log-file-count=10 -i scn-proxy
```
4. Start the Enterprise Controller.
 - On Oracle Solaris:

```
# /opt/SUNWxvmoc/bin/satadm start -w -v
```
 - On Linux:

```
# /opt/sun/xvmoc/bin/satadm start -w -v
```
5. Verify that the Enterprise Controller has been restarted completely before attempting other operations. For example, if you have stopped both the Enterprise Controller and a Proxy Controller, wait for the Enterprise Controller to restart before restarting each Proxy Controller. On an Oracle Solaris Proxy Controller:

```
# /opt/SUNWxvmoc/bin/proxyadm start -w -v
```

On a Linux Proxy Controller:

```
# /opt/sun/xvmoc/bin/proxyadm start -w -v
```

- 6.** Verify that all controllers have restarted completely before attempting other operations.

JumpStart Enterprise Toolkit

Use JumpStart Enterprise Toolkit (JET) to extend the JumpStart installation functionality provided within the Oracle Solaris 9 and 10 operating systems.

JET is a framework designed to simplify and extend the JumpStart installation capabilities for provisioning, the Oracle Solaris 9 or 10 operating system. JET provides a set of helper scripts to simplify the use of Jumpstart for the installation of Solaris 10 and earlier on both SPARC and x86 servers. Oracle Solaris 11 does not use JET, instead, it uses the automated installer (AI).

The SUNWjet and JetFLASH packages are installed on the Proxy Controller during installation. See the JET page on the Oracle Technology Network at <http://www.oracle.com/technetwork/systems/jet-toolkit/jet-toolkit-1614844.html> for more information on JET, including additional packages available for download and a link to the user documentation.

Note: JET must run on a Proxy Controller that is running on an Oracle Solaris 9 or 10 operating system.

JumpStart Enterprise Toolkit Configuration File Location

The JET module parameters are available for use in OS profiles. See the `module.conf` configuration files that are associated with JET modules for information about parameters for specific JET modules. The configuration files are located in the `/opt/SUNWjet/Products` directory on the Proxy Controller. For example, the configuration files for the custom module is located in the following directory on the Proxy Controller: `/opt/SUNWjet/Products/custom/custom.conf`. You can review the parameters for these modules by looking at the `sample.template` file in the `/opt/SUNWjet/Templates` directory on a Proxy Controller.

SUNWjet Parameters

The main JET framework is supplied in a single SVR4 package called SUNWjet. This package contains everything necessary to do a standard Oracle Solaris installation using either bootp or dhcp.

When you specify JET parameters with an OS profile, the following parameters from the `base_config` JET module are automatically updated within the OS profile and must not be modified:

- `base_config_ClientArch`
- `base_config_ClientEther`

- `base_config_client_allocation`
- `base_config_sysidcfg_network_interface`
- `base_config_sysidcfg_ip_address`
- `base_config_sysidcfg_netmask`
- `base_config_sysidcfg_nameservice`
- `base_config_sysidcfg_system_locale`
- `base_config_sysidcfg_terminal`
- `base_config_sysidcfg_timeserve`
- `base_config_sysidcfg_timezone`
- `base_config_sysidcfg_root_password`
- `base_config_sysidcfg_security_policy`
- `base_config_sysidcfg_protocol_ipv6`

The following list describes the parameters that are associated with the `base_config JET` module. These parameters provide basic operating system configuration information. Values for many of these parameters use the term `targetableComponent` to represent the target system.

- `base_config_client_allocation`: The mechanism used to build this client. By default, the options listed in `/opt/SUNWjet/etc/jumpstart.conf` are used. Leave the value blank unless you need to do something different from the default for this specific client. If you are provisioning the Oracle Solaris 10 1/06 x86 release, set the value of this variable to `GRUB` to enable GRUB-based booting and installation.
- `base_config_ClientArch`: Kernel architecture, such as `sun4u` or `x86`. By default, this is set to the kernel architecture of the `targetableComponent`.
Default Value: `[targetableComponent:kernel_arch]`
- `base_config_ClientEther`: Ethernet MAC address. By default, this is set to the Ethernet MAC address of the `targetableComponent`.
Default Value: `[targetableComponent:ethernet_mac_address]`
- `base_config_ClientOS`: Version of the OS to be provisioned.
Example: `Solaris9_u7_sparc`
- `base_config_dedicated_dump_device`: If set, the `dumpadm` utility configures the partition as a Dedicated Dump Device. See `dumpadm(1M)` for supported Operating Environments.
- `base_config_defaultrouter`: Value to use for `/etc/defaultrouter`.
- `base_config_disable_sysid_probe`: If set, skip the `sysid` step on the first reboot. This can significantly increase provisioning efficiency on systems that have many unused network adapters.
Default Value: `yes`
- `base_config_dns_disableforbuild`: Delay DNS configuration until later. If DNS is not available in the build environment, set this variable to `yes`.
- `base_config_dns_domain`: DNS domain entry for the `/etc/resolv.conf` file.
- `base_config_dns_nameservers`: Space-separated list of IP addresses to use for DNS name server entries in the `/etc/resolv.conf` file.

- `base_config_dns_searchpath`: List of entries to go in the DNS search line in `/etc/resolv.conf` file.
- `base_config_dumpadm_minfree`: Set a limit so that crash dumps do not fill up the dump file system. See the `dumpadm(1M)` `-m` option for possible values.

Example: 20000k

- `base_config_enable_altbreak`: If set, enable alternate break sequence.
- `base_config_enable_rootftp`: If set to any value, enable root FTP access.
- `base_config_enable_rootlogin`: If set to any value, enable network root login from telnet, rsh, and ssh.
- `base_config_enable_savecore`: If set to any value, enable save core for Solaris 2.6 systems.

Default Value: yes

- `base_config_grub_append`: For Oracle Solaris 10 1/06 x86 systems, specifies additional options or arguments to pass to the GRUB bootloader.
- `base_config_ipmp_networkifs`: Space-separated list of interfaces to be defined under IPMP control. For each interface listed, define sets of variables to provide the netgroup, mode, test1, test2, netmask, host name, log-ip, host name2, and log-ip2 for the interface.

Example: qfe0_qfe4!database-net 1 10.0.0.1 10.0.0.2 24 oracle-db
10.0.0.3 apache 10.0.0.4

- `base_config_networkifs`: Space-separated list of additional network interfaces to be defined. For logical interfaces, use underscores (_) rather than colons (:). Use the format `format cntndn`. For each interface listed, define sets of variables to provide the netname, netmask, host name, and IP address for the interface.

Example: le1!netB 255.255.255.0 myhost-netB 192.168.1.0

- `base_config_nfs_mounts`: Space-separated list of remote NFS mount points. Use ? to separate the mount source from the mount target, as shown in the example.

Example: fs?1.1.1.1:/fs

- `base_config_nfsv4_domain`: Set up the NFSv4 domain to prevent being prompted at first reboot. If not set, look first for the entry in `base_config_dns_domain`, and second for the domain value in `/etc/default/nfs`.
- `base_config_noautosshutdown`: If set to any value, disable power management.

Default Value: pm_disabled

- `base_config_nodename`: Value to use for `/etc/nodename` if not the default host name.
- `base_config_notrouter`: If set to y, then disable IPv4 forwarding and create the `/etc/notrouter` file.
- `base_config_ntp_servers`: Space-separated list of names or IP addresses for the NTP servers. The first server has a prefer tag. This section places lines of the form: `server [prefer]` into the `/etc/inet/ntp.conf` file. For additional NTP control, use the custom module to deploy your own custom `ntp.conf` file.
- `base_config_patchdir`: Path to the patches. If blank, use information from the `jumpstart.conf` file and the IP address of the JET server. If your patch files are not stored on the JET server, then provide an NFS-style path to the location of the patches.

- `base_config_poweroff_afterbuild`: If set, shut down the system after the build completes.
- `base_config_productdir`: Path to the products. If blank, use information from the `jumpstart.conf` file and the IP address of the JET server. If your package files are not stored on the JET server, then provide an NFS-style path to the location of the packages.
- `base_config_products`: JET modules to provision.
- `base_config_profile`: Create a custom JumpStart profile. By default, if you leave this variable blank, the OS provisioning plug-in creates the `/opt/SUNWjet/Clients/hostname/profile` based on the other `base_config_profile` variables. Alternatively, you can create your own custom JumpStart profile. To use the profile that you created manually, set the `base_config_profile` variable to the name of the created profile. By default, the OS provisioning plug-in looks for the profile in the `/opt/SUNWjet/Clients/hostname` directory. To direct the plug-in to a profile in another directory, provide an absolute path name in the `base_config_profile` variable.

Note: If you are provisioning Oracle Solaris OS on x86 target hosts, you must create a custom JumpStart profile that deletes any existing partitions and point to that profile in the `base_config_profile` variable.

- `base_config_profile_add_clusters`: Space-separated list of cluster packages to add.
- `base_config_profile_add_geos`: Comma-separated list of geographical regions to add.
Example: `N_Europe, C_Europe`
- `base_config_profile_add_locales`: Comma-separated list of locales to add.
Example: `fr_FR, ja_JP.UTF-8`
- `base_config_profile_add_packages`: Space-separated list of packages to add.
- `base_config_profile_additional_disks`: A list of disks to use and configure in addition to the boot disk. Use the format `cnt:rdn`. For each disk listed, define sets of variables for each slice to identify the mount point and the size.
- `base_config_profile_cluster`: Oracle Solaris software group package.
 - **Default Value:** `SUNWCreq`
 - **Example:** `SUNWCreqSUNWCuserSUNWCprogSUNWCallSUNWCXallSUNWCrnet`
- `base_config_profile_del_clusters`: Space-separated list of cluster packages to remove.
Example: `SUNWCpm SUNWCpmx SUNWCdial SUNWCdialx`
- `base_config_profile_del_geos`: Comma-separated list of geographical regions to delete.
- `base_config_profile_del_locales`: Comma-separated list of locales to delete.
- `base_config_profile_del_packages`: Space-separated list of packages to remove. To prevent interactive installations on Solaris x86 headless target hosts, set this value to `SUNWxwssu SUNWxwscf`.

- `base_config_profile_dontuse`: A comma-separated list of disks that must not be used. Use the format *cntndn*. This variable applies only if `base_config_profile_usedisk` is not set.
- `base_config_profile_root`: Root space (free, or size in Megabytes)
Default Value: free.
- `base_config_profile_s3_mtpt`: Mount path to slice 3.

Note: If you are using VxVM and you want your boot disk to look like the mirror, then leave slices 3 and 4 empty.

- `base_config_profile_s3_size`: Size of slice 3 (in Megabytes).
- `base_config_profile_s4_mtpt`: Mount path of slice 4.
- `base_config_profile_s4_size` – Size of slice 4 (in Megabytes).
- `base_config_profile_s5_mtpt`: Mount path of slice 5.
Default Value: /var
- `base_config_profile_s5_size`: Size of slice 5 (in Megabytes).
- `base_config_profile_s6_mtpt`: Mount path of slice 6.
Default Value: /usr
- `base_config_profile_s6_size`: Size of slice 6 (in Megabytes).
- `base_config_profile_s7_mtpt`: Mount path of slice 7.
Default Value: /opt

Note: If you are using Oracle Solaris Volume Manager (SVM), the default behavior is to use slice 7 as a location for metastate databases. If you are using the SVM default configuration, do not use slice 7 for data.

- `base_config_profile_s7_size`: Size of slice 7 (in Megabytes).
- `base_config_profile_swap`: Swap space (in Megabytes).
Default Value: 256
- `base_config_profile_usedisk`: Defines the boot disk onto which the OS will be loaded. Use the format *cntndn* or the keyword `rootdisk`. If the value is `rootdisk`, then the current boot disk is used.
Default Value: `rootdisk`
- `base_config_shutup_sendmail`: If set, create an alias host name to disable `sendmail`.
Default Value: yes
- `base_config_sysidcfg_default_route`: Router IP address to use during JumpStart for Solaris 9 or later environments. If blank, JumpStart uses value from the `defaultrouter_base_config` variable. If that is also blank, or for another net interface, JumpStart `sysidcfg` gets a router IP from the JET server.

- `base_config_sysidcfg_ip_address`: IP address to use at initial boot. By default, this is set to the IP address of the targetable component.
Default Value: `[targetableComponent:ethernet_ip_address]`
- `base_config_sysidcfg_nameservice`: Name service to configure at initial boot.
Default Value: NIS
- `base_config_sysidcfg_netmask`: Netmask to use at initial boot. By default, this is set to the netmask of the targetable component.
Default Value: `[targetableComponent:ethernet_netmask]`
- `base_config_sysidcfg_network_interface`: Network interface to use at initial boot.
Default Value: NONE
- `base_config_sysidcfg_protocol_ipv6`: Whether to use IPv6 protocol at initial boot.
Default Value: no
- `base_config_sysidcfg_root_password`: Encrypted root password.
- `base_config_sysidcfg_security_policy`: Kerberos security policy to use at initial boot.
Default Value: NONE
- `base_config_sysidcfg_system_locale`: System locale to use at initial boot.
Example: `n_US.ISO8859-1`
- `base_config_sysidcfg_terminal`: Terminal emulator to set at initial boot.
Default Value: `vt100`
- `base_config_sysidcfg_timeserver`: Where to get system time for initial boot. If blank, system time comes from the JET server. Alternatively, you can set this variable to `localhost` to get the system time from the hardware clock on the client.
- `base_config_sysidcfg_timezone`: System time zone to use for initial boot.
Example: `US/Pacific`
- `base_config_sysidcfg_x86_kdmfile`: For Solaris x86 systems, specifies the name of a keyboard, display, and mouse configuration file to append to the `sysidcfg` file.
Default Value: `/sysidcfg-addon-file`
- `base_config_ufs_logging_filesys`: For Oracle Solaris 7 and later systems, a space-separated list of mount points to use for logging. To enable logging on all UFS file systems, use the keyword `all`. Oracle Solaris 9 09/04 enables logging by default. To disable logging on a specific file system, add a hyphen in front of the mount point. To disable logging on all file systems, use the keyword `none`.
Default Value: `all`

Note: You cannot mix keywords and mount points. You can specify the root file system (`/`), although the root file system is included as part of the `all` and `none` keywords.
- `base_config_update_terminal`: If set, put the `sysidcfg` terminal type into `inittab`.

Default Value: yes

- `base_config_x86_confflags`: For Oracle Solaris 9 x86 systems, specifies arguments to be used with the `confflags` attribute of the `add_install_client` command.

Example: `-f -P /boot/solaris/dca`

- `base_config_x86_console`: For x86 systems, set the console to the correct tty port if you are not going to connect a keyboard and monitor to the client. Setting this variable enables you to perform installs through the serial port. For b1600, v20z, and v40z systems, use `ttya`. For lx50, v60x, and v65x systems, use `ttyb`.
- `base_config_x86_disable_acpi`: For x86 systems, any value disables ACPI. Disabling ACPI might make the installation process proceed better due to how the interrupts are handled.
- `base_config_x86_disable_kdmconfig`: For Oracle Solaris x86 systems, disables the `kdmconfig` interactive utility for configuring the keyboard, display, and mouse of the target host. If you are installing an Oracle Solaris OS with the GRUB bootloader, set this variable value to `yes`.
- `base_config_x86_nowin`: For x86 systems, prevents Oracle Solaris from trying to run Windows during the install.

Default Value: yes

- `base_config_x86_safetoreboot`: For x86 systems, controls whether the system automatically reboots. If your PXE boot is a one-time option, and the next reboot attempts to boot from disk, you must set this option to `yes`.

Downloading Additional JET Packages

Additional JET packages are available. To download JET packages and the JET User Guide, go to

<http://www.oracle.com/technetwork/systems/jet-toolkit/index.html>.

An icon is a small picture or symbol on a user interface (UI) that represents an action or a state.

Oracle Enterprise Manager Ops Center has a variety of icons to represent an asset, the current state of an asset, and actions. A tool tip appears when you place the mouse over an icon. Tool tips provide brief descriptions of the icons that appear in Oracle Enterprise Manager Ops Center.

When an icon is grayed out, it means that the feature is not functional for the current scenario or you have not been assigned the role to perform the action.

Administration Icons

Oracle Enterprise Manager Ops Center provides you with a variety of tools for understanding and managing your datacenter. The Administration icons are associated with actions that you use to configure and manage Oracle Enterprise Manager Ops Center. You can only enter or edit the cost of power in Administration. Following are the icons that are associated with Administration.

Figure E-1 Register Enterprise Controller



Figure E-2 Ops Center Downloads



Figure E-3 Set Enterprise Controller Storage Library



Figure E-4 Setup Connection Mode



Figure E-5 Edit Authentications



Figure E–6 Proxy Deployment Wizard



Figure E–7 Configure Local Agent



Figure E–8 Self Diagnosis



Figure E–9 Unconfigure Enterprise Controller



Figure E–10 Change HTTP Proxy



Figure E–11 Edit Energy Cost



Figure E–12 Add User



Figure E–13 DHCP Config



Figure E–14 Subnets



Figure E–15 External DHCP Server



Figure E–16 Navigate to Selected Oracle VM Manager



Figure E–17 Discover Oracle VM Managers



Figure E–18 Discover Oracle VM Server**Figure E–19 Discover File Server****Figure E–20 Discover iSCSI Storage**

Badges

Badges are special types of icons that appear in the corner of asset icons to show the status of the assets.

Figure E–21 Critical Badge**Figure E–22 Warning Badge****Figure E–23 Info Badge****Figure E–24 Configure Badge****Figure E–25 Shutdown Badge****Figure E–26 Update Badge****Figure E–27 Lock Badge****Figure E–28 Unreachable Badge**

Boot Environment Icons

A boot environment is an instance of a bootable Oracle Solaris image plus additional software packages that are installed onto the image, and the set of all file systems and devices (disk slices and mount points) that are required to operate an Oracle Solaris OS instance.

Following are the icons that are associated with Boot Environment.

Figure E–29 Create New Boot Environment



Figure E–30 Activate Boot Environment and Reboot



Figure E–31 Active Boot Environment



Figure E–32 Active Boot Environment Upon Reboot



Figure E–33 Delete New Boot Environment



Figure E–34 Synchronize Boot Environment



Connectivity Icons

Connectivity icons display the session mode and connectivity. To determine if you run in Connected Mode and have access to My Oracle Support Services, check the icons in the upper right corner of the UI. If an icon is not colored, you are not in connected mode. The icons indicate:

- Whether or not Oracle Enterprise Manager Ops Center is connected to the Internet.
- Whether or not Oracle Enterprise Manager Ops Center is connected to the Oracle Knowledge Base.
- Whether or not Oracle Enterprise Manager Ops Center is connected to My Oracle Support Services.

Figure E–35 My Oracle Support Services Connected



Figure E–36 My Oracle Support Services Disconnected



Figure E–37 Internet Connected



Figure E–38 Internet Disconnected**Figure E–39 Knowledge Base Connected****Figure E–40 Knowledge Base Disconnected**

Group Icons

Group icons are icons that are associated with administrative structures that organize assets so that you can locate an asset quickly or perform operations on all assets of the same type.

Figure E–41 Create Group**Figure E–42 Add Assets to Groups****Figure E–43 Remove Asset From Group****Figure E–44 Move Assets to Group****Figure E–45 Add Attribute Icon****Figure E–46 Add Rule Icon****Figure E–47 Delete Attribute Icon****Figure E–48 Delete Rule Icon**

Hardware Icons

Hardware icons represent the various types of hardware that is used to manage the datacenter.

Figure E-49 Blade Icon



Figure E-50 Chassis Icon



Figure E-51 CPU Icon



Figure E-52 Disk Controller Icon



Figure E-53 Disk Expander



Figure E-54 Ethernet Port Icon



Figure E-55 Fan Icon



Figure E-56 Fan Tray Icon



Figure E-57 Generic Hardware Icon



Figure E-58 Hard Disk Icon



Figure E-59 I/O Board Icon



Figure E-60 I/O Module Icon



Figure E-61 Memory Icon



Figure E-62 Memory Board Icon



Figure E-63 PCI Icon



Figure E-64 Power Control Setting Icon



Figure E-65 Power Supply Setting Icon



Figure E-66 System Icon



Figure E-67 Utilization CPU



Figure E-68 Utilization File System



Figure E-69 Utilization



Figure E-70 Utilization IO



Figure E-71 Utilization Memory



Figure E-72 Utilization Network



Figure E-73 Utilization Swap Icon



Figure E–74 Utilization System Load Icon



Figure E–75 XSCF Icon



Incident Icons

Incident icons are icons that reflect the current state of the incidents.

Figure E–76 View Alerts



Figure E–77 View Annotations



Figure E–78 View Possible Impacts and Causes



Figure E–79 View Comments



Figure E–80 View Suggested Actions



Figure E–81 Add Annotations



Figure E–82 Assign Incidents



Figure E–83 Acknowledge Incident



Figure E–84 Take Action on Incidents



Figure E–85 Mark Incident as Repaired



Figure E-86 Close Incident

Job Icons

Every action creates a job. Many jobs can run at the same time and some jobs take longer than others to complete. To follow the progress of a particular job, you can display the job details.

Figure E-87 All Jobs**Figure E-88 Jobs In Progress****Figure E-89 Jobs Waiting For User Input****Figure E-90 Jobs Stopped****Figure E-91 Jobs Scheduled****Figure E-92 Jobs Partially Successful****Figure E-93 Jobs Suspended****Figure E-94 Jobs Failed****Figure E-95 Jobs Successful****Figure E-96 View Job Details**

Figure E–97 Re-run Selected Jobs



Figure E–98 Display Selected Target Details



Figure E–99 Debug Jobs Using OCMonitor



Figure E–100 Mode Actual



Figure E–101 Delete Selected Jobs



Kernel Icons

The kernel is a program that constitutes the central core of a computer operating system. It has complete control over everything that occurs in the system.

Figure E–102 Tree Node Icon



Library Icons

Libraries store and manage cached data, images, packages, and metadata. A software library stores images for provisioning operations. A Storage Library is a file system on the virtual host's server or it can be accessed through an NFS server or SAN network. Guest configurations for its operating system, data, CPU, memory, and network as metadata are stored in the storage library associated with the virtual host.

Figure E–103 New Local Software Library



Figure E–104 New NAS Software Library



Figure E–105 New SAN Storage Library



Figure E-106 Set Enterprise Controller Storage Library

Message Center Icons

You can create, view, and manage service requests from the Oracle Enterprise Manager Ops Center user interface.

Figure E-107 Open Service Requests**Figure E-108 Service Request Opened By Others****Figure E-109 My Service Requests****Figure E-110 Notifications****Figure E-111 Delete All Notifications****Figure E-112 My Incidents****Figure E-113 Unassigned Incidents**

Monitoring Icons

Monitoring rules and policies define the monitoring parameters. Monitoring rules define the alerting conditions. Rules are associated with, and determined by, the type of managed resource. Monitoring rules state the values and boundaries for an asset's activity. A monitoring policy defines alert configurations to be performed on one or more managed resources. It contains the information needed to monitor a managed asset with user-defined alert configurations, including defined thresholds and alert monitors.

Below are the icons that are associated with monitoring rules and policies.

Figure E-114 Add Alert Monitoring Rule

Figure E–115 Deselect All Rules



Figure E–116 Select All Rules



Figure E–117 Edit Alert Monitoring Rule



Figure E–118 Disable Alert Monitoring Rule



Figure E–119 Delete Alert Monitoring Rule



Figure E–120 Navigate to Monitoring Policy



Figure E–121 Applying a Monitoring Policy



Figure E–122 Extract a Monitoring Policy



Figure E–123 Create a New Monitoring Policy



Figure E–124 Copy Monitoring Policy



Figure E–125 Delete Monitoring Policy



Figure E–126 Set as Default Policy



Figure E–127 View Associated Assets



Network Icons

Oracle Enterprise Manager Ops Center manages networks for its virtual hosts. Guests in the network communicate with each other or with the Internet through these virtual hosts. Use networks to:

- Manage individual hosts
- Connect hosts to the Proxy Controller
- Allow guests to communicate with each other or with the Internet
- Connect remote JMX with the public API

Below are the icons that are associated with Network.

Figure E-128 Create Network Domain



Figure E-129 Delete Network Domain



Figure E-130 Connect Guests



Figure E-131 Create Public Network



Figure E-132 Define Network



Figure E-133 Edit Attributes



Operating System Icons

Operating system are icons that depict the type of operating system and the actions that are associated with the operating system.

Figure E-134 Linux Operating System Icon



Figure E-135 Oracle Solaris Icon



Figure E–136 *Oracle Solaris Cluster Icon*



Figure E–137 *Microsoft Windows Icon*



Figure E–138 *New Update OS Job*



Figure E–139 *Deploy/Update Software*



Figure E–140 *Reboot*



Figure E–141 *Attach Network*



Figure E–142 *Place in Maintenance*



Figure E–143 *Update Operating System*



Oracle VM Server Icons

These are icons that are related to the functions for Oracle VM Server.

Figure E–144 *Oracle VM Server Pool*



Figure E–145 *x86 Server Pool*



Figure E–146 *LDOM SPARC*



Figure E-147 LDOM Server Pool



Figure E-148 LDOM SPARC Server Pool



Figure E-149 Global Zone



Process Icons

These are icons associated to the processes of Oracle Enterprise Manager Ops Center.

Figure E-150 CPU Tree Node Icon



Figure E-151 Memory Tree Node Icon



Figure E-152 Tree Node Icon



Report Icons

Reports let you gather information about your assets, such as job history, firmware, operating system updates, and incidents. You can export reports or use them to start jobs on targeted assets.

Figure E-153 Create Windows Host Compliance Report



Figure E-154 Create Windows Incident Compliance Report



Figure E-155 View Host Compliance Report Result



Figure E-156 Rerun Report



Figure E–157 Delete Incident Compliance Report



Figure E–158 Run Report Template



Figure E–159 Report Pending to Run



Figure E–160 Create System Information Report



Figure E–161 Create Incident Summary Report



Figure E–162 Create Incident Detail Report



Figure E–163 Create Firmware Report



Figure E–164 Create Change History Report



Figure E–165 Create Baseline Report



Figure E–166 Create System Catalog Report



Figure E–167 Create CVE report



Figure E–168 Create Recommended Software Configuration Report



Figure E–169 Create Solaris Update Compliance Report



Figure E–170 Create Profile Report



Figure E–171 Create Exalogic System Report



Figure E–172 Compare Catalog Report



Figure E–173 Distribution Update Report



Figure E–174 Package Compliance Report



Figure E–175 Service Pack Compliance Report



Figure E–176 View CSV Output Icon



Figure E–177 View Interactive Output Icon



Figure E–178 View PDF Output Icon



SuperCluster Icons

SPARC SuperCluster is an Oracle Engineered System, which integrates SPARC compute nodes, a ZFS storage appliance, InfiniBand switches, and Exadata database Cells into a multi-rack system. Below are the icons that are associated with SuperCluster.

Figure E–179 SuperCluster



Figure E–180 SuperCluster Cluster



Figure E–181 SuperCluster Node



Figure E–182 SuperCluster Partnership



Figure E–183 SuperCluster Quorum



Figure E–184 SuperCluster Zone Cluster



Figure E–185 SuperCluster Zone Cluster Node



Figure E–186 Oracle Solaris Cluster Upgrade



Switch Icons

Switch icons represent the hardware devices that join multiple computers within one local area network (LAN).

Figure E–187 General Switch



Figure E–188 InfiniBand Icon



Figure E-189 10 GIG E**Figure E-190 Switch Reset****Figure E-191 Launch Switch UI**

Universal Icons

Universal icons are icons that are associated with the most common functions or actions that are performed in the application.

Figure E-192 Add Icon**Figure E-193 Delete Icon****Figure E-194 Deselect All Icon****Figure E-195 Select All Icon****Figure E-196 Power Off Icon****Figure E-197 Refresh Icon****Figure E-198 Start Icon****Figure E-199 Suspend Icon**

Figure E–200 View Details Icon**Figure E–201 Go To Icon****Figure E–202 Shutdown**

Update Policy Icons

Policies are lists of actions that are explicitly approved or denied. Policies are specific to operating system update and define the amount of user interaction you want when applying operating system patches and packages. Update policies define the component configuration of the system you want to update.

Figure E–203 New Policy**Figure E–204 Import Policy**

Update Profile Icons

Update profile icons are icons that are associated to the update actions that are carried out on a profile in a defined order.

Figure E–205 Copy Icon**Figure E–206 New Profile Icon****Figure E–207 Delete Icon****Figure E–208 Import Profile****Figure E–209 Profile Install Packages**

Figure E-210 Profile Uninstall Packages



vDC Icons

Virtual Datacenter or vDC is a collection of the server pools that share the common storage and network resources.

Figure E-211 Reserved IP Address



Figure E-212 Allocated IP Address



Figure E-213 Used IP Address



Oracle Enterprise Manager Ops Center API

Oracle Enterprise Manager Ops Center exposes certain APIs. These APIs allow external systems management solutions to integrate with Oracle Enterprise Manager Ops Center. The APIs are remotely available using TCP/IP, through one of two secure protocols: JMXRemoting over TLS and WS-Management over https.

The Oracle Enterprise Manager Ops Center Software Developer's Kit (SDK), that is shipped with the Oracle Enterprise Manager Ops Center software, includes some examples that show how to use the public APIs.

To use the examples included in the package:

1. Locate the `SUNWxvmoc-sdk.pkg` package in the `dvd/<platform>/Product/components/packages/` directory on the DVD/Install Media.
2. Install the `SUNWxvmoc-sdk.pkg` package.
3. Locate the example you want. [Table F-1](#) lists the examples in the package.
4. Compile the example.
5. Execute the example, using the appropriate shell script.

Table F-1 Examples Included in the `SUNWxvmoc-sdk.pkg` Package.

Example	Description	Code
Basic client example using JMX	<p>This example demonstrates the JMX Messaging Protocol (JMXMP) connector with Java/JMX technology. This example does not require any external JARs except Java and JDMK. This example performs the following functions:</p> <ul style="list-style-type: none"> ■ Configures the connection. ■ Performs security settings. ■ Opens the connection (locally or remotely). ■ Uses the connection with JMX basic requests. ■ Closes the connection. 	BasicClient.java

Table F–1 (Cont.) Examples Included in the SUNWxvmoc-sdk.pkg Package.

Example	Description	Code
Domain model client example using WS-MAN	<p>This example demonstrates direct access to the read-only copy of the domain model. This example performs the following functions:</p> <ul style="list-style-type: none">■ Configures the connection.■ Performs security settings.■ Opens the connection (locally or remotely).■ Queries the domain model for all the server objects and displays the value of their model attributes.■ Queries the domain model for all the OperatingSystems objects and displays the value of their description attributes.■ Queries the domain model for all the ZoneController objects and displays the value of their ZoneControllerType attributes.■ Queries the domain model for all the ISOImage objects and displays the value of their DiskImagePath attributes.■ Queries the domain model for all the VirtServerImage objects and displays the value of their DomainFilePath attributes.■ Closes the connection.	DomainModelClient.java
Basic client example using WS-MAN	<p>This example performs the same functions as the basic client with JMX example. However, this example uses the WS-MAN connector with Java/JMX technology.</p>	SimpleClient.java
Server services access example using WS-MAN	<p>This example demonstrates the server public services. This example performs the following functions:</p> <ul style="list-style-type: none">■ Configures the connection.■ Performs the security settings.■ Opens the connection (locally or remotely)■ Creates a proxy for each xVM server service to verify everything is working.■ Displays the attribute values.■ Lists the service tags.■ Lists the nodes and their names.■ Lists the xVM servers and displays some information on each one including the list of guests.■ Closes the connection.	ServerClient.java ServerClientServices.java

Table F–1 (Cont.) Examples Included in the SUNWxvmoc-sdk.pkg Package.

Example	Description	Code
Services access using WS-MAN	<p>This example demonstrates specific Oracle Enterprise Manager Ops Center public services. This example performs the following functions:</p> <ul style="list-style-type: none">■ Configures the connection.■ Performs security settings.■ Opens the connection (locally or remotely).■ Creates a proxy for each specific service to verify everything is working.■ Closes the connection.	<p>OpsCenterClient.java OpsCenterClientServices.java</p>
Discovery and agent provisioning example using WS-MAN	<p>This example demonstrates some Oracle Enterprise Manager Ops Center public services. This example performs the following functions:</p> <ul style="list-style-type: none">■ Configures the connection.■ Performs security settings.■ Opens the connection (locally or remotely) with the satellite.■ Uses the discovery service to discover one machine.■ Registers the server asset for this machine.■ Provisions the agent on this machine.■ Unprovisions the agent on this machine.■ Deletes both the server and operating system asset created for this machine.■ Closes the connection.	<p>OCProvisionClient.java Utils.java</p>
Server example using WS-MAN	<p>This example demonstrates some Oracle Enterprise Manager Ops Center public services. This example performs the following functions:</p> <ul style="list-style-type: none">■ Configures the connection.■ Performs some security settings.■ Opens the connection (locally or remotely) with the satellite.■ Unattaches the xVM server if it is already attached.■ Deletes the asset associated with the xVM server.■ Uses the discovery service to discover one xVM server.■ Attaches the xVM server to the Oracle Enterprise Manager Ops Center.■ Closes the connection.	<p>OCXVMServer.java Utils.java</p>

List of Public APIs

[Table F–2](#) provides a brief description of the public APIs. For a complete list of packages, its classes, interfaces, and methods, refer to the Javadoc that is included in the SDK.

Table F–2 List of Public APIs

Package	Description
com.oracle.sysman.services.bios	Provides BIOS configuration services.
com.oracle.sysman.services.bootenvironment	Provides parameters for ABE related operations.
com.oracle.sysman.services.cloudmgmt	Provides virtual datacenter (vDC) administration and vDC management services.
com.oracle.sysman.services.common	Provides common constants and utility classes.
com.oracle.sysman.services.discovery	Provides services to perform discovery and manage discovery profiles.
com.oracle.sysman.services.exadatacell	Provides the service interface for exadata cell management.
com.oracle.sysman.services.network	Provides network domain services.
com.oracle.sysman.services.ocdoctor	Provides services for OCDoctor functionality in Ops Center.
com.oracle.sysman.services.ops	Provides Oracle pre-engineered system management services.
com.oracle.sysman.services.osanalytics	Provides services for operating systems and virtual server container analytics.
com.oracle.sysman.services.ovm	Provides services to perform operations on OVM deployments through the OVM manager.
com.oracle.sysman.services.rack	Provides rack management services.
com.oracle.sysman.services.reset	Provides the reset service interface to handle resetting and refreshing the service processor.
com.sun.xvm.services.agentprov	Provides service for managing the provisioning of agents. This package provides services for the installation, configuration, and removal of agent software to remote operating systems.
com.sun.xvm.services.agentupgrade	Provides the service interface for agent upgrades.
com.sun.xvm.services.alarm	Provides services for alarm management.
com.sun.xvm.services.authorization	Provides the operations for authorization management.
com.sun.xvm.services.bootenvironment	Provides the interface for boot environment service.
com.sun.xvm.services.cachemanager	Provides the service interface for managing the cached view of the domain model.
com.sun.xvm.services.cluster	Provides the service interface for cluster management. This service is responsible for viewing and managing Solaris cluster instances and their nodes.
com.sun.xvm.services.common	Provides the service interfaces that are common to different public API services.
com.sun.xvm.services.componentmgmt	Provides the service interface for hardware components management.
com.sun.xvm.services.datamodel	Provides the service for accessing low-level data models.
com.sun.xvm.services.deleteasset	Provides the service interface to delete assets such as servers, operating systems, service tags, and chassis. The virtualization assets such as hosts, guests, virtual image libraries, and networks can be removed by calling the appropriate methods on their respective services.

Table F–2 (Cont.) List of Public APIs

Package	Description
com.sun.xvm.services.discovery	Provides services to discover assets. The interface for discovery service provides capability to initiate a discovery operation for a set of IP addresses.
com.sun.xvm.services.event	Provides the interface for event service. The interface for event service allows the client to subscribe to notifications emitted from the domain model.
com.sun.xvm.services.fabricmgmt	Provides service interfaces for fabric management.
com.sun.xvm.services.fwdeployment	Provides services for firmware deployment management. This package manages deployments of firmware updates to servers and chassis.
com.sun.xvm.services.gear	Provides service tag and general gear management interfaces.
com.sun.xvm.services.grouping	Provides asset group management services.
com.sun.xvm.services.guest	Provides service interfaces for guest management.
com.sun.xvm.services.guest.provision	Provides service interfaces for guest provisioning.
com.sun.xvm.services.jobmanager	Provides job management services for performing job actions, controlling job status, and retrieving information on job actions and job history.
com.sun.xvm.services.migration	Provides the service interface for migration service. This package provides the service for live, warm and cold migration of guest domains.
com.sun.xvm.services.mos	Provides the My Oracle Support services.
com.sun.xvm.services.network	Provides the network management services. This package provides services for managing networks for guests.
com.sun.xvm.services.notification	Provides notifications management services.
com.sun.xvm.services.oemgc	Provides the service interface for OEM Grid Control repository integration.
com.sun.xvm.services.operation	Provides services for operation profile functionality.
com.sun.xvm.services.opscenter	Provides the service interface for XVM Update Job Manager.
com.sun.xvm.services.osdeployment	Provides services and types to manage deployment of operating systems to servers. This package provides services to access aggregated and detailed information on installed operating systems, monitor OS activity, and provision OS. This package is available only in Enterprise Controllers.
com.sun.xvm.services.osmonitoring	Provides the parameters for OS monitoring profile value object creation.
com.sun.xvm.services.pis	Provides the product and asset information service and types interfaces, used to classify assets according to their product name and hierarchy.
com.sun.xvm.services.productupdate	Provides the service interface to download update catalog files and bundles.
com.sun.xvm.services.profile	Provides the profile services to create, delete, and update profiles.
com.sun.xvm.services.provisioning	Provides services for managing server provisioning.
com.sun.xvm.services.proxyservice	Provides service for managing the provisioning of Proxy Controllers.
com.sun.xvm.services.report	Provides services and types for retrieving reports on utilization trend.

Table F–2 (Cont.) List of Public APIs

Package	Description
com.sun.xvm.services.reportsmanager	Provides services interfaces for the report manager.
com.sun.xvm.services.storagemgmt	Provides service interfaces for storage server management.
com.sun.xvm.services.svrmgmt	Provides service interfaces for server management.
com.sun.xvm.services.switchmgmt	Provides the physical and virtual switch management service interface.
com.sun.xvm.services.topology	Provides the topology service and types interfaces to get topology information of the managed resources. This package provides a tree-based navigation view of the Ops Center assets.
com.sun.xvm.services.user	Provides the services to manage user preferences.
com.sun.xvm.services.virtimagelib	Provides the services for virtual images library and virtual image management.
com.sun.xvm.services.virtpool	Provides the service for managing the life cycle of virtual pools sharing the same libraries or networks, and have the same processor architecture, integrating virtual pools with storage, networks, and virtualization hosts.
com.sun.xvm.services.virtservercontainer	Provides the service interface for managing virtualized server containers.
com.sun.xvm.services.win	Provides filters supported for Windows services.
com.sun.xvm.services.winpm	Provides service interfaces for Windows compliance reports and Windows software updates deployment.
com.sun.xvm.services.wme	Provides the service interface for Windows management extension administration.
com.sun.xvm.services.xvmserver	Provides the XVM server management service for managing the system level aspects of the XVM server.

Public API Access

- [Public API Access Through JMX Clients](#)
- [Public API Access Through WS-Management](#)

JMX

Java Management Extensions (JMX) technology provides the tools for building distributed, modular, and dynamic solutions for managing and monitoring devices, applications, and networks. The JMX API defines the notion of MBeans, or manageable objects, which expose attributes and operations in a way that allows remote management applications to access them. An important new feature of the JMX API in the Java Platform, Standard Edition 7 (Java SE 7) is its ability to create MXBeans, a substantial improvement over Standard MBeans. For more information on JMX, visit the Java Management Extensions (JMX) Technology website at <http://www.oracle.com/technetwork/java/javase/tech/javamanagement-140525.html>.

WS-Management

Web Services for Management (WS-MAN) is a specification for managing servers, devices, and applications using web services standards, providing a common way for systems to access and exchange management information across the entire IT infrastructure. It is commonly used in the Windows management stacks, and is

supported by Microsoft management tools, including the `winrm` command-line interface and COM objects that can be used from Visual Basic or from C#.

Public API Access Through JMX Clients

The simplest and preferred way of interacting with the public APIs is using a Java client. This Java client communicates with the public APIs using a secure remote JMX client. By using a remote JMX client, it is possible to directly interact with the Service MBeans in Java using the strongly-typed Java MBeans interface class library. The JMX MBean interfaces are documented in the javadoc that is included in the SDK. The following remote JMX management protocols are supported: JMXMP and WS-MAN.

Public API Access Through WS-Management

The public API MBeans are exposed over WS-Management through JSR-262. This specification defines a way to use Web Services to access JMX instrumentation remotely. Clients are not required to be Java applications. This permits a generic WS-Management client to interact with the public APIs. The WS-Management mappings for each service tier API is documented in the Javadoc that is included in the SDK.

Virtualization Management Stack

The virtualization management stack is a three-tier management paradigm. The stack consists of presentation tier, service tier, and an underlying domain model.

Presentation Tier

The presentation tier provides the user interface, displays the information provided by the service tier, and takes actions through the service tier.

Service Tier

The service tier contains all the core management logic. This tier exposes the management functionality through a set of public APIs that are used by the presentation tier and external systems management solutions. These public APIs are the ideal integration points for third-party software to control or to monitor the Enterprise Manager Ops Center software.

The service tier exposes its management APIs as a series of Java Management Extensions (JMX) MBean interfaces. These interfaces are made accessible as public APIs through several remote protocols, including JMXRemoting (JSR 160) and WS-Management (JSR 262).

The services are identified by MBean ObjectName, and these are of the form:

```
com.sun.xvm.services:type=<service name>
```

The main entry point to the public APIs is the MBean interface defined in `com.sun.xvm.services`.

Domain Model

The domain model provides the representation of underlying managed resources and associations between them. The domain model exposes a normalized data model of

these managed resources as a collection of JMX MXBeans. Each managed resource has one MXBean and each association between these resources as one MXBean.

Expression Query Language

Oracle Enterprise Manager Ops Center expression query language is an advanced feature that enables you to build query object constraints and create customized expression monitoring rules. See [Adding Monitoring Rules](#) for more information about monitoring rules. See the *Oracle® Enterprise Manager Ops Center Tuning Monitoring Rules and Policies Guide* for examples of how to tune the default rules that are provided with the software and how to add new rules.

The expression query language is mainly used to define monitoring rules and to check if an asset attribute verifies a given condition. The query language is closely modeled on the WHERE clause of SQL SELECT statements.

This section contains the following information:

- [Monitoring Attributes and Classes](#)
- [Monitoring Attributes and Classes](#)
- [Lexical Elements](#)
- [Grammar](#)
- [Semantics](#)
- [Informal Examples](#)

Monitoring Attributes and Classes

Asset monitoring attributes vary by asset type. The asset types and corresponding attributes are available in the API Javadocs that are part of the SDK package. The package installs the Javadocs in the `/opt/sun/xvm/sdk/xvm_oc/doc/javadoc/` directory. Open the Javadoc file to view all attributes and classes.

Classes

The classes corresponding to the asset types are located in the `com.sun.hss.type` package and subpackages. The following are a few of the classes representing the asset types that you can use to create expression monitoring rules and some of the attributes available for monitoring:

- `com.sun.hss.type.os.OperatingSystem`
- `com.sun.hss.type.server.Server`
- `com.sun.hss.type.servercontainer.ServerContainer`
- `com.sun.hss.type.virtserver.VirtServer`
- `com.sun.hss.type.virtserver.VirtServerContainer`

For each class, you can deduce the attributes that the software can monitor by looking for the *getter* methods defined by the class. Those are the methods whose names start with *get* or *is* and which take no parameters. The name of the attribute corresponding to a method is built by removing the *get* or *is* prefix from the method name. For example, the `OperatingSystem` class defines a method named `getCpuUsage`, as follows:

```
@Description(value="CPU Usage information")
```

```
CpuUsage getCpuUsage()
```

The name of corresponding attribute name is `CpuUsage`.

AttributeType

The return type of the method describes the attribute type. To find the subfields of the attributes that the software can monitor, navigate to the definition of the corresponding class. The name of the subfields are also deduced from the *getter* methods defined on the attribute type class. Given the name of a *getter* method, you can deduce the subfield name by removing the *get* or *is* prefix and by setting the next letter to lower case. Following the previous example, when you navigate to the javadoc of the `CpuUsage` class, you will find the following *getter* method:

```
@ValueInfo(metricType="gauge",  
    thresholdRising=true,  
    units="%",  
    minValue="0",  
    maxValue="100")  
public float getUsagePercentage()
```

This means the subfield is named *usagePercentage*. The complete attribute / subfield name that you can use in a monitoring rule is `CpuUsage.usagePercentage`.

Lexical Elements

The following are some lexical elements:

- Attribute names are case sensitive.
- Keywords, such as `and`, `like`, and `between`, are not case sensitive.
- Use double quotes to access an attribute whose name, ignoring case, is the same as one of these keywords: `not`, `instanceof`, `like`, `true`, or `false`. For example, `"not"`. You can use double quotes to include non-identifier characters in the name of an attribute. For example, `"attribute-name-with-hyphens"`. To include the double quote character in the attribute name, write it twice. `"foo" "bar" "baz"` represents the attribute called `foo"bar"baz`.
- String constants are written with single quotes, for example `'this'`. A single quote within a string constant must be doubled, for example `'can' 't'`.
- Integer constants are written as a sequence of decimal digits, optionally preceded by a plus or minus sign. An integer constant must be a valid input to `Long.valueOf(String)`.
- Floating-point constants are written using the Java syntax. A floating-point constant must be a valid input to `Double.valueOf(String)`.
- A boolean constant is either `true` or `false`, ignoring case.

- Spaces cannot appear inside identifiers (unless written with double quotes), keywords, or multi-character tokens, such as <=. Spaces can appear anywhere else, but are not required except to separate tokens. For example, you can write the query $a < b$ and $5 = c$ as $a < b$ and $5 = c$, but you cannot remove any other spaces.

Grammar

query

andquery [OR query]

andquery

predicate [AND andquery]

predicate

(query) |

NOT predicate |

INSTANCEOF stringvalue |

LIKE objectnamepattern |

value predrhs

predrhs

compare value |

[NOT] BETWEEN value AND value |

[NOT] IN (value commavalues) |

[NOT] LIKE stringvalue

commavalues

[, value commavalues]

compare

= | < | > | <= | >= | <> | !=

value

factor [plusorminus value]

plusorminus

+ | -

factor

term [timesordivide factor]

timesordivide

* | /

term

attr | literal | (value)

attr

name [# name]

name:

identifier [.name]

identifier

Java-identifier | double-quoted-identifier

literal

booleanlit | longlit | doublelit | stringlit

booleanlit

FALSE | TRUE

stringvalue

stringlit

objectnamepattern

stringlit

Semantics

Table G–1 describes the grammar semantics and defines a function q that maps a string to a Java object, such as a QueryExp or a ValueExp.

Table G–1 Semantics

String s	$q(s)$
$query1$ OR $query2$	<code>Query.or(q(query1), q(query2))</code>
$query1$ AND $query2$	<code>Query.and(q(query1), q(query2))</code>
($queryOrValue$)	<code>q(queryOrValue)</code>
NOT $query$	<code>Query.not(q(query))</code>
INSTANCEOF $stringLiteral$	<code>Query.isInstanceOf(Query.value(q(stringLiteral)))</code>
LIKE $stringLiteral$	<code>new ObjectName(q(stringLiteral))</code>
$value1 = value2$	<code>Query.eq(q(value1), q(value2))</code>
$value1 < value2$	<code>Query.lt(q(value1), q(value2))</code>
$value1 > value2$	<code>Query.gt(q(value1), q(value2))</code>
$value1 \leq value2$	<code>Query.leq(q(value1), q(value2))</code>
$value1 \geq value2$	<code>Query.geq(q(value1), q(value2))</code>
$value1 <> value2$	<code>Query.not(Query.eq(q(value1), q(value2)))</code>
$value1 \neq value2$	<code>Query.not(Query.eq(q(value1), q(value2)))</code>
$value1$ BETWEEN $value2$ AND $value3$	<code>Query.between(q(value1), q(value2), q(value3))</code>

Table G-1 (Cont.) Semantics

String s	q(s)
<i>value1</i> IN (<i>value2</i> , <i>value3</i>)	Query.in(q(<i>value1</i>), new ValueExp[] { q(<i>value2</i>), q(<i>value3</i>)})
<i>value1</i> NOT IN (<i>value2</i> , <i>value3</i>)	Query.not(Query.in(q(<i>value1</i>), new ValueExp[] { q(<i>value2</i>), q(<i>value3</i>)}))
<i>value</i> LIKE <i>stringLiteral</i>	Query.match(q(<i>value</i>), translateWildcards(q(<i>stringLiteral</i>)))
<i>value</i> NOT LIKE <i>stringLiteral</i>	Query.not(Query.match(q(<i>value</i>), translateWildcards(q(<i>stringLiteral</i>))))
<i>value1</i> + <i>value2</i>	Query.plus(q(<i>value1</i>), q(<i>value2</i>))
<i>value1</i> - <i>value2</i>	Query.minus(q(<i>value1</i>), q(<i>value2</i>))
<i>value1</i> * <i>value2</i>	Query.times(q(<i>value1</i>), q(<i>value2</i>))
<i>value1</i> / <i>value2</i>	Query.div(q(<i>value1</i>), q(<i>value2</i>))
<i>name</i>	Query.attr(q(<i>name</i>))
<i>name1</i> # <i>name2</i>	Query.attr(q(<i>name1</i>), q(<i>name2</i>))
FALSE	Query.value(false)
TRUE	Query.value(true)
<i>decimalLiteral</i>	Query.value(Long.valueOf(<i>decimalLiteral</i>))
<i>floatingPointLiteral</i>	Query.value(Double.valueOf(<i>floatingPointLiteral</i>))

The `translateWildcards` function translates from the SQL notation for wildcards, using `%` and `_`, to the API notation using `*` and `?`. If the `LIKE` string already contains a `*` or `?`, these characters have their literal meanings, and are quoted in the call to `Query.match`.

Informal Examples

The formal specification of the language is described in [Lexical Elements](#). The attributes are located in the javadocs.

This section provides some informal examples.

```
Message = 'OK'
```

`Message = 'OK'` is verified if the `Message` attribute is the string `'OK'`.

```
Message like 'OK: %'
```

`Message` attribute whose value is a string beginning with `"OK: "`. The wildcard characters are the same as in SQL. In the query language, percent character means any sequence of characters and the underscore character means any single character.

```
FreeSpacePercent < 10
```

The `FreeSpacePercent` attribute whose value is a number less than 10.

```
FreeSpacePercent < 10 and WarningSent = false
```

Uses the same attribute as the previous example, but includes a boolean attribute `WarningSent` whose value is false.

```
SpaceUsed > TotalSpace * (2.0 / 3.0)
```

`SpaceUsed` and `TotalSpace` attributes where the first is more than two-thirds the second.

```
not (FreeSpacePercent between 10 and 90)
```

`FreeSpacePercent` attribute whose value is not between 10 and 90, inclusive.

```
FreeSpacePercent not between 10 and 90
```

Another way of writing the previous query.

```
Status in ('STOPPED', 'STARTING', 'STARTED')
```

`Status` attribute whose value is one of those three strings.

Example G-1 Root Disk and CPU Usage

[Example G-1](#) is an expression that will raise an alert when the root disk is above 80% and the CPU usage is above 10.

```
FileSystemUsages.name="/" .usedSpacePercentage >= 80) AND (CpuUsage.usagePercentage >= 10
```

Example G-2 Root Disk and System Load

[Example G-2](#) is an expression that will raise an alert when the root disk is above 80% and the system load is above 15.

```
FileSystemUsages.name="/" .usedSpacePercentage >= 70) AND (SystemLoad.average1Minute >= 15
```

Glossary

account

An account entitles designated cloud users the right to use computing, network, and storage resources of vDC. The account provides the required capabilities to manage these resources. Account defines the amount of vCPU, memory and storage resources that can be used from the available vDC resources.

actions pane

The Actions pane is used to start jobs based on the current selection in the Navigation pane. Selections in the Navigation pane or center pane change the display of operations in the Actions pane. The Actions pane is subdivided into four sections – Operate, Organize, Deploy, and Update.

active

Reflects the state of system and indicates whether monitoring is actually being performed. The active state is not editable. When a rule is not enabled, monitoring is not active. The status is displayed on the Alert Monitoring Rules page, which is accessed from the Monitoring tab. Text in the Active field indicates whether the parameter is active.

activate

Changes an inactive Oracle Solaris boot environment to the new default boot environment on reboot.

Agent Controller

The Agent Controller software communicates with the Enterprise Controller and is installed automatically when an asset is discovered to make the asset a managed asset. You can choose to manage resources remotely with proxy resources without putting an agent on the system. Some features of the product don't work without the agent, but discovery manages the assets without putting an agent on them.

agentless

A system that is managed with Oracle Enterprise Manager Ops Center without the Agent Controller software being installed.

alert monitor

Monitors the state of managed resources and their attributes and raise an alert when the state is outside the pre-defined thresholds.

Alternate Boot Environment

An alternate boot environment, or ABE, is an inactive Oracle Solaris boot environment.

annotations

Annotations are scripts or comments that you can associate with a incident. Annotations can be automated operations to solve a incident, a suggested action, or a comment. You can associate an annotation with a specific incident. Annotations can be added to the Incidents Knowledge Base.

assemblies

Assemblies are kind of infrastructure templates that contain a configuration of multiple virtual machines with their virtual disks and the inter connectivity between them. Assemblies can be created as a set of .ovf (Open Virtualization Format) and .img (disk image) files, or may all be contained in a single .ova (Open Virtualization Format Archive) file.

assets

Assets are physical or virtual piece of hardware, storage device, or operating system that you can manage with Oracle Enterprise Manager Ops Center.

Auto-Balancing Policy

An auto-balancing policy determines if, and how, a server pool is automatically load balanced. By default, automatic balancing is not selected. When you designate the server pool for automatic balancing, the software reviews the load on the virtualization hosts for the interval and day that you request. The software then migrates the guests, as needed, to balance the load. You can require administrator approval before the guests are moved. Also see placement policy and policy.

bandwidth flow

Bandwidth flow is the speed of a connection, or the amount of data that flows from a site's server out to the viewer at any given time.

Baseline

A dated collection of Oracle Solaris patches, patch metadata, and tools. Oracle releases Solaris baselines on a monthly basis. You can use the black lists and the white lists to modify a baseline and create a custom patch set.

baseline check

Baseline check is a feature of Oracle Enterprise Manager Ops Center Engineered Systems where the factory setup of eighth, quarter, half, and full rack configurations are considered as a normal or ideal setup. When the assets are discovered and associated with the rack, this setup is compared with the normal factory setup.

black list

A list of Oracle Solaris operating system patch IDs that you never want to apply to an asset. The black list is used when you are using a baseline to update an Oracle Solaris operating system.

block storage

A block storage library consists of LUNs (Logical Unit Number). Each LUN is a slice of a storage volume, which is storage space provided by a collection of disks.

Boolean Control Parameter

A monitoring rule that uses a true-false check.

Boot Environment

A collection of mandatory file systems (disk slices and mount points) that are critical to the operation of the Oracle Solaris operating system. These disk slices can be on the same disk or distributed across multiple disks.

branded zone

Zones that are capable of emulating user environments from operating systems other than Oracle Solaris 10. Zones supports different versions of Oracle Solaris operating system in the zones for running applications.

category

For Oracle Enterprise Manager Ops Center's Local Content, a category is the type of software that is uploaded to Oracle Enterprise Manager Ops Center for use at a site. The parent category is one of the types defined in Oracle Enterprise Manager Ops Center. The local category is a category defined for the site, for example a script for a quarterly inventory.

channel

An operating system distribution, such as Oracle Solaris 10 5/09 on x86 platform or Oracle Linux 5.5. A channel is also called a distribution.

Cloud

A cloud is a set of physical resources that can be divided and allocated to multiple users who can in turn create and use virtual resources as needed without impact to or awareness of the other users' resources. A cloud is implemented as a pool of servers sharing the same virtualization type, storage, networks and fabrics.

cluster heartbeat

Cluster heartbeat is used to verify if the Oracle VM Servers in a clustered server pool are up and running. The heartbeat function has a network component, where a TCP/IP communication channel is created with each Oracle VM Server. Each Oracle VM Server sends regular keep-active packets and these packets are used to determine if each Oracle VM Server is active.

connected mode

This is the default connection mode for Oracle Enterprise Manager Ops Center. With this mode, patch data is regularly downloaded from Knowledge Base through an Internet connection.

control domain

A domain that is created when Oracle VM Server for SPARC software is installed. The control domain contains the software packages for Oracle VM Server, including the domains manager application and the domains manager daemon (ldmd) process required for managing the logical domains. The interface to the hypervisor is through the domains manager. The control domain enables you to create, and manage logical domains and allocate virtual resources to the domains.

critical file system

File systems that are required by the Oracle Solaris operating system. When you use Solaris Live Upgrade, these file systems are separate mount points in the vfstab file of

the active and inactive boot environments. Example file systems are `root`, `/usr`, `/var`, and `/opt`. These file systems are copied from the source to the inactive boot environment.

Dashboard

Displays a high-level overview of an asset or a group of assets on the user interface. The information of the selected asset or group is displayed in the Center Pane.

Deployment Plans

Defines the sequence of steps that must be carried out on an asset to deploy. Deployment plans also include the specification or profile that each step should apply, and the resources that are required to apply it such as network addresses, host names and so on. Customized deployment plan enables you to perform hardware, firmware and operating system provisioning activities in a repeatable fashion.

disconnected mode

This is the alternate connection mode for Oracle Enterprise Manager Ops Center. Instead of relying on an Internet connection for updates, patch data is acquired using the harvester script and moved to the Enterprise Controller.

discovery

This is the method for adding assets to Oracle Enterprise Manager Ops Center. Assets can be discovered using a variety of protocols, by their service tags, or by declaring hardware so that it can be configured and provisioned with an operating system.

distribution

For an operating system, a distribution is a specialized version of the operating system.

distribution groups

Distribution groups are defined especially for Oracle Exalogic system based vDCs. Place your vServers in a distribution group and ensure that no two vServers run on the same Oracle VM Servers.

Domain Name Service (DNS)

DNS is a network protocol that issues IP addresses within a specified range to devices on the network.

Dynamic System Domains

In M-Series servers, you can partition the available hardware resources into smaller logical systems called as dynamic system domains. Dynamic System Domains run their own copies of the operating system and offer a very high level of isolation from other domains in the system because the partitioning occurs at the hardware level.

dynamic storage library

When the block storage library uses LUNs constructed from a storage array that is a managed asset, the block storage library is dynamic. You can add storage capacity as needed by adding LUNs supplied by the storage array.

When the block storage library relies on a storage array that is not a managed asset, the block storage library is static. Because Oracle Enterprise Manager Ops Center has less information about the storage array, you cannot increase the number of LUNs in the storage library.

enabled

A monitoring rule that is enabled is actively monitoring a parameter. By default, all rules are enabled. Users can disable and enable parameters on a per asset or group basis. The status is displayed on the Alert Monitoring Rules page, which is accessed from the Monitoring tab. Text in the Enabled field indicates whether the parameter is enabled.

Enterprise Controller

This is the central server for Oracle Enterprise Manager Ops Center software. The Enterprise Controller hosts the user interface and communicates with the Knowledge Base. Enterprise Controller stores management information, such as firmware and operating system images, plans, profiles, and policies and also stores the asset data and site customizations. All operations, or jobs, are initiated from the Enterprise Controller.

Enumerated Control Parameter

A monitoring rule that uses a series of values.

Exclusive IP Mode

A dedicated network interface is allocated to the zone. You can choose the network interface when you assign the network to a zone.

Expression Parameter

A monitoring rule that uses an instruction to execute something that returns a value.

/etc Directory

The directory that contains critical system configuration files and maintenance commands.

/etc/netboot Directory

The directory on a WANboot server that contains the client configuration information and security data that are required for a WANboot installation.

/export File System

A file system on an operating system server that is shared with other systems on a network. For example, the /export file system can contain the root (/) file system and swap space for diskless clients and the home directories for users on the network. Diskless clients rely on the /export file system on an operating system server to boot and run.

fabrics

Fabrics are network topologies where network nodes connect with each other through one or more network switches. A true fabric provides a direct connection between any two ports, and supports single step/lookup-based processing. Regardless of its various components, a fabric appears on the outside as a single, logical device with a single, consistent state.

The term is popular in telecommunication, Fibre Channel storage area networks, and other high-speed networks, including InfiniBand.

filesystem storage

A software or storage library that relies on a file system on the Enterprise Controller's system or a shared file system on an NFS server that the Enterprise Controller mounts.

global zone

In Oracle Solaris Zones, the global zone is both the default zone for the system and the zone used for system-wide administrative control. The global zone is the only zone from which a non-global zone can be configured, installed, managed, or uninstalled. Administration of the system infrastructure, such as physical devices, routing, or dynamic reconfiguration (DR), is only possible in the global zone. Appropriately privileged processes running in the global zone can access objects associated with other zones.

group

A group is a user-defined set of assets. Assets can be added to groups based on asset attributes such as type or location. A group can include other groups. Assets can be manually added in addition to the rules based addition using attributes. Any type of asset that can be in a group can be added manually to any user-defined asset group.

guest

Guest refers to a virtual machine that is configured and installed in a virtualization host. For example, the logical domains in an Oracle VM Server host are referred to as guests in a server pool.

GUID

Globally Unique Identifier. A pseudo-random 128-bit number that is computed by Windows to identify any component in the computer that requires a unique number. In Oracle Enterprise Manager Ops Center, GUIDs are used to identify LUNs.

Hardware Virtualization (HVM)

Hardware virtualization is a technology that is used to create multiple virtual systems on a single piece of physical hardware. When you create a hardware virtualized (HVM) guest, you must supply an ISO file in a repository to create the virtual machine.

Hardware Virtualized with Paravirtualized Drivers (PVHVM)

PVHVM is identical to HVM, but has additional paravirtualized drivers for improved performance of the virtual machine. PVHVM improves the performance level of Microsoft Windows running in guests.

host name

The name by which a system is known to other systems on a network. This name must be unique among all the systems within a particular domain (usually, this means within any single organization). A host name can be any combination of letters, numbers, and dashes (-), but it cannot begin or end with a dash.

hypervisor

A hypervisor is the software that enables multiple virtual machines to be multiplexed on a single physical machine. The hypervisor code runs at a higher privilege level than the supervisor code of its guest operating systems to manage use of the underlying hardware resources by multiple supervisor kernels.

Image Packaging System (IPS)

Image Packaging System is an Oracle Solaris 11 package that contains operating system components and a manifest that provides basic metadata.

incident

An event that triggers an alert when a monitored attribute does not meet the monitoring parameters. A new incident is displayed in the Unassigned Incidents queue in the Message Center. From the Message Center you can view and act on incidents.

Incident Knowledge Base

A custom database of annotations that are associated with known incidents.

InfiniBand

InfiniBand is a switched fabric communications link primarily used in high-performance computing. Its features include quality of service and failover, and it is designed to be scalable. The InfiniBand architecture specification defines a connection between processor nodes and high performance I/O nodes such as storage devices.

InfiniBand transmission rates begin at 2.5 GBps.

IPMP

IPMP (IP network multipathing) provides physical interface failure detection and transparent network access failover. You can configure one or more physical interfaces into an IP multipathing group, or IPMP group. After configuring IPMP, the system automatically monitors the interfaces in the IPMP group for failure.

JET Templates

JumpStart Enterprise Toolkit provides a framework to simplify and extend the JumpStart functionality provided within the Oracle Solaris operating system. You can use JET to install Oracle Solaris on the SPARC and x86/64 platforms. You create JET templates to customize the operating system configuration options as required.

JMX

Java Management Extensions (JMX) technology provides the tools for building distributed, modular, and dynamic solutions for managing and monitoring devices, applications, and networks. The JMX API defines the notion of MBeans, or manageable objects, which expose attributes and operations in a way that enables remote management applications to access them. The public API in Oracle Enterprise Manager Ops Center can be accessed through JMX-Remoting.

Knowledge Base

The Knowledge Base is the repository for metadata about Oracle Solaris and Linux operating system components. Knowledge base stores information about patch dependencies, patch compatibilities, withdrawn patches, downloads, and deployment rules and also stores URL of operating system vendor download sites and downloads the components at set intervals. The Enterprise Controller must have Internet connection to connect to the Knowledge Base.

least allocated

Least allocated is a parameter in the server pool placement policy. The lowest allocated CPU and memory is the total static resource allocation across all guests on the virtualization host. The other placement policy parameter is relative load.

libraries

A collection of virtual machine images and disk images that are located under the same file system. When a server pool is created, one or more libraries are assigned to the server pool. Server pools can share the same libraries.

link aggregation

Link aggregation is a standard defined in IEEE802.3ad. An aggregated link consists of several interfaces on a system configured as a single, logical unit. Link aggregation increases the speed and high availability of a connection between a server and a switch.

LUN

LUN stands for Logical Unit Number. In storage, a LUN is the number assigned to a SCSI protocol entity, that handles (I/O) operations. A SCSI target provides a LUN for each storage volume.

management

An asset is managed when Oracle Enterprise Manager Ops Center can monitor it and target it with jobs. Operating systems can be managed with or without an Agent Controller, but operating system update functions are only available with an Agent Controller.

manifest

Each Oracle Solaris 11 package has an associated manifest that describes how the package is put together. The package manifest provides basic metadata about the package (such as name, description, version, and category), what files and directories are included, and the package dependencies.

membership graph

Shows a graphical relationship between assets and status of the connection. A blue line shows the working connection and a red line represents the faulted or disconnected status. The membership graph is displayed in the Center Pane.

message center

Displays all incidents, alerts, and notifications. Message Center helps you to view and manage incidents, notifications, and service request, and display warranty information.

MTU

MTU stands for Maximum Transmission Unit. MTU is the largest packet size, in bytes, that can be sent over a network.

monitoring policy

A set of monitoring rules that defines alert conditions. Policies are either system-defined, user-defined, or generic. Each monitoring policy contains one or more alert monitors for a specific type of resource. An alert is raised when the state is outside the pre-defined condition.

monitoring rule

Contains monitoring parameters that state the values and boundaries for an asset's activity. The set of rules is called a monitoring policy.

MPxIO

MPxIO provides a multipathing solution for storage devices accessible through multiple physical paths. MPxIO is included as a part of the distribution in Solaris 10 onwards.

NAT

NAT stands for Network Address Translation. NAT is a protocol that enables a network to use many internal-only IP addresses and a few Internet-facing IP addresses.

navigation pane

Navigation pane is an important part of the user interface of Oracle Enterprise Manager Ops Center. navigation pane contains Message Center, Assets, Plan Management, Networks, Libraries, Reports, vDC Management, and Administration. The Assets section of the Navigation pane lists all the asset that are managed by Oracle Enterprise Manager Ops Center, grouped by its type and the required criteria.

network

A network enables guests to communicate with each other or with the external world (that is, the Internet). When a server pool is created, one or more networks is assigned to the server pool. Server pools can share the same networks.

network bonding

Network bonding refers to the combination of network interfaces on one host for redundancy and/or increased throughput. Redundancy is the key factor you use to protect your virtualized environment from loss of service due to failure of a single physical link. This network bonding equals as the Linux network bonding. Using network bonding in Oracle VM might require some switch configuration.

network domain

A system of centralized network administration, in which the permissions that grant access to resources in the network are maintained in one or more servers. Network Domains use a hierarchical structure that enables you to assign permissions to collaborate with different departments in an organization.

A large network may have several domains based on the needs of each set of users.

NIS

NIS stands for Network Information System. NIS is a network naming and administration system for smaller networks. NIS is similar to the Internet's domain name system (DNS) but designed for a smaller network.

non-global zone

A virtualized operating system environment created within a single instance of the Oracle Solaris operating system. One or more applications can run in a non-global zone without interacting with the rest of the system. Non-global zones are also called zones.

non-sparse copy

A clone of the type "non-sparse copy" is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

notifications

An email, pager, or user interface message that is automatically sent by Oracle Enterprise Manager Ops Center when specified conditions are met. You can configure separate notification profiles for different assets and different users. You can configure the software to send notification for specific incidents, or when a critical or warning incident is detected.

Opaque Data

An opaque data is a data type that is incompletely defined in an interface, so that its values can only be manipulated by calling subroutines that have access to the missing information.

/opt

A file system that contains the mount points for third-party and unbundled software.

Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control is a single, integrated solution for managing all aspects of the Oracle Cloud and the applications running on it. Oracle Enterprise Manager Cloud Control couples a potent, top-down monitoring approach to delivering the highest quality of service for applications with a cost-effective automated configuration management, provisioning, and administration solution.

Oracle Engineered System

Oracle Engineered Systems are hardware and software integrated systems that are designed for a specific enterprise purpose. Oracle Engineered System helps in reducing the cost and complexity of the IT infrastructures, and increases the productivity and performance.

Oracle Services

Provides integrated methods of maintaining and displaying current contracts, warranty information, contract dates, and service requests in Oracle Enterprise Manager Ops Center.

Oracle Solaris Clusters

Oracle Solaris Clusters is a high availability software product for Solaris operating system. Oracle Solaris Clusters are used to improve the availability of software services such as databases, file sharing on a network, electronic commerce websites, or other applications. You can now manage Oracle Solaris Clusters as any other asset using Oracle Enterprise Manager Ops Center.

Oracle Solaris Zones

Oracle Solaris Zones is a software partitioning technology used to virtualize operating system services, and provide an isolated and secure environment for running applications. When you create a non-global zone, you produce an application execution environment in which processes are isolated from all other zones. This isolation prevents processes that run in a zone from monitoring or affecting processes that run in any other zones. See also global zone and non-global zone.

Oracle Solaris 11 Software Update Library

Oracle Solaris 11 Software Update Library repository is located on the Enterprise Controller. This contains the Oracle Solaris 11 packages that you need to install, provision, and update your Oracle Solaris 11 operating system.

Oracle VM Server for SPARC

Oracle VM Server is a virtualization technology that enables the creation of multiple virtual systems by a hypervisor in the firmware layer, interposed between the operating system and the hardware platform. This is designed to abstract the hardware and can expose or hide various system resources, allowing for the creation of resource partitions that can operate as discrete systems, complete with virtual CPU, memory and I/O devices.

Oracle VM Server for SPARC was previously known as Logical Domains, it is a virtualization technology designed to run on CMT based servers.

Oracle VM Server for x86

Oracle VM Server for x86 is a managed virtualization environment or part of such an environment, that is designed to provide a lightweight, secure, server-based platform for running virtual machines. Oracle VM Server for x86 is based upon an updated version of the underlying Xen hypervisor technology, and includes Oracle VM Agent.

Oracle Solaris ZFS

An Oracle Solaris operating system file system that uses storage pools to manage physical storage.

Paravirtualization

Paravirtualization enables you to select a location for the mounted ISO file from which you create the virtual machine. Before you create the virtual machine using the paravirtualized method, you must mount the ISO file on an NFS share, or HTTP or FTP server.

parent repositories

Any hosted Oracle repository that Oracle Solaris 11 Software Update Library can use to upload, or sync, content.

photorealistic view

Photorealistic view displays the front and rear views of the rack. All slots and the respective assets are displayed. Positions within the rack are displayed in a 2-dimensional view. All assets in the rack have a specific image. The health status of assets such as OK, Warning, and Critical are displayed in the form of colored buttons.

placement policy

Determines whether the guest is placed on a virtualization host with the lowest relative load or the least allocated. By default, new guests are placed on the server with the lowest load and are automatically started. The placement policy is defined when a server pool is created. Server pools can have different placement policies.

policy

Defines how a job is performed and sets the automation level of the job. A policy file is similar to a response file. If there is a conflict between a profile and policy, the profile overrides the policy.

Private vNet

vNet that is unique to a given account is called Private vNet.

profile

Defines the configuration of components for a specific type of system. By using a profile, you can define what is enabled, and not enabled, to be installed on a system. If there is a conflict between a profile and policy, the profile overrides the policy.

Proxy Controller

Proxy Controllers link the managed assets to the Enterprise Controller and act as proxies for operations that must be located close to the managed assets, such as operating system provisioning. Proxy Controllers distribute the network load and provide for fan-out capabilities to minimize network load. Proxy Controllers perform management operations on assets and report the results to the Enterprise Controller. An Oracle Enterprise Manager Ops Center installation must have at least one functioning Proxy Controller.

relative load

Relative load is a parameter in the server pool placement policy. Lowest relative load is based on the lowest memory and CPU utilization for the virtualization host over the past three weeks. The other placement policy parameter is least allocated.

repository

A repository is a central place that stores an aggregation of data in an organized way, usually in a computer storage. Depending on how the term is used, a repository may be directly accessible to users or may be a place from which specific databases, files, or documents are obtained for further relocation or distribution in a network.

root

The top level of a hierarchy of items. `root` is the one item from which all other items are descended. See `root` directory or `root (/)` file system.

root directory

The top-level directory from which all other directories stem.

root file system

The top-level file system from which all other file systems stem. The `root (/)` file system is the base on which all other file systems are mounted, and is never unmounted. The `root (/)` file system contains the directories and files critical for system operation, such as the kernel, device drivers, and the programs that are used to boot a system.

RPM

A package manager used by many versions of the Linux operating system.

rule parameters

Define the monitoring parameters. The following types of rule parameters are available: Threshold, Boolean Control, Enumerated Control, and Expression. Some parameters are editable. All active parameters can be disabled.

SAN Storage Library

Storage Attached Network (SAN) storage which is used for providing storage spaces for managed assets in Oracle Enterprise Manager Ops Center. The SAN storage library consists of groups of LUNs.

script

A command file that is associated with one of Oracle Enterprise Manager Ops Center's actions, either before the action occurs (pre-action script), or after the action completes (post-action script).

security group

The organization of users and other domain objects into groups for easy administration of access permissions is known as a security group. A Security Group enables you to specify certain security settings on an instance specific basis. You have the ability to filter traffic based on IP's (a specific address or a subnet), packet types (TCP, UDP or ICMP), and ports (or a range of ports). You can also grant access to an entire security group. This enables your trusted computers to access each other without having to open ports to the public.

server management

Server management is used to manage the physical Oracle VM Servers in a server pool, for example, to update the Oracle VM Agent on the different Oracle VM Servers.

server pool

A server pool is a resource pool of virtualization hosts that share compatible chip architecture, which facilitates actions such as moving guests between virtualization host instances. Members of the server pool have access to the same network and storage library resources. Guests can access the images contained in the server pool's library. Several server pools can share the same network and library storage resources.

server templates

Server templates provide pre-built images for creating vServers. They can be uploaded individually or as part of an Assembly. Server templates can be created from an existing vServer.

service tag

Service tags are XML files that identify assets uniquely. Assets with service tags can be discovered using the Find Assets wizard.

Shared IP Mode

The global zone shares its network interface with one or more zone. You must define the network interface when you assign the network to the global zone.

snapshot

Snapshot, a point in time image of a volume is a non-bootable copy of a boot environment that uses much less disk space than a boot environment. You can create a boot environment from a snapshot.

software libraries

A software library can be a local file system on the Enterprise Controller or a mount point on an NFS server. The software library is used to store the operating system images for provisioning, branded images, flars, firmwares, profiles, operating system updates, custom programs and scripts.

sparse Ccopy

A clone of the type "sparse copy" is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size.

static route

Specifies the route taken by the network for external access. You define a default gateway for the network; however, this default gateway may not be reachable to a given subnet. In this case, you must add a static route for this specific subnet.

status pane

The Status pane in the Jobs section describes about the state of the incidents like jobs in progress, jobs failed, jobs partially successful, jobs stopped, jobs schedules, jobs successful and so on.

SCCM

Microsoft System Center Configuration Manager (SCCM), is used to update Windows operating systems.

syncing

Syncing is the process of reconfiguring or updating the Oracle Solaris 11 Software Update Library with the Oracle Solaris 11 Image Packaging System (IPS).

synchronizing

Updates an inactive boot environment to match an active boot environment.

system groups

Default asset groups that automatically organize your assets by type in the user interface.

System-defined Rules

Attribute specific monitoring rules that are hard-coded into drivers. You can disable a system-defined rule, but cannot edit, move, or reconfigure these types of rules.

Thin Clone

A thin clone is a clone of a physical disk that takes up only the amount of disk space actually in use; not the full specified disk size.

threshold parameters

A monitoring rule that uses a numeric value above or below a defined level.

time server

The network device that provides accurate time for synchronizing network activity.

unmanaged storage

Unmanaged storage is the storage resource that is unknown to Oracle Enterprise Manager Ops Center. When you add storage to zones using the native CLI or manage existing zone environments, the zone's storage is not identified and termed as unmanaged.

User-defined Network Domain

A network domain provides custom network resources from an Ethernet or InfiniBand fabric to virtualization hosts, server pools, or virtual datacenters so that new networks can be created as needed. A user-defined network domain supplements the Default Network Domain that is always available and cannot be deleted.

User-defined Rules

Monitoring rules that are associated with, and determined by, the type of managed resource. You can apply a user-defined rule to many different attributes.

/usr File System

A file system on a standalone system or server that contains many of the standard UNIX programs.

Sharing the large /usr file system with a server rather than maintaining a local copy minimizes the overall disk space that is required to install and run the Solaris software on a system.

/var File System

A file system or directory (on standalone systems) that contains system files that are likely to change or grow over the life of the system. These files include system logs, vi files, mail files, and UUCP files.

vDC

vDC is a collection of physical servers and storage that are placed on a common network. These physical resources are organized into a pool that are accessed by self-service users. This offers an access point through which you can allocate and control the resources inside. This is created during the set up phase. An example of a vDC is a single Exalogic system that might consist of multiple racks.

vNets

vNets are managed networks and their associated logical (L2) fabrics that can be associated with a vDC and its Accounts.

vServer

vServer is an entity that provides the outward interface of a standalone operating system. This may be a Virtual Machine (VM) or a Solaris Container or a similar construct. This consumes CPU and memory resources. This can be a member of one or multiple vNets.

vServer Type

vServer type is a profile for vServer creation that defines size of memory, size of disk and number of vCPUs to be used when creating a new vServer instance, that is used in combination with a Server Template.

VID

VLAN Identifier. Part of the VLAN tag inserted into Ethernet frame that specifies its VLAN.

white list

A list of Oracle Solaris operating system patch IDs that you always want to be applied to a host. The white list is used when you are using a baseline to update an Oracle Solaris operating system.

virtual disk image

A virtual disk image is a representation of a virtual storage device that is associated with a virtual machine. Such storage can represent a virtual hard disk or a virtual CD/DVD.

virtualization host

Oracle VM Server that are managed by Oracle Enterprise Manager Ops Center is referred to as virtualization host. The virtualization host contains a hypervisor and its local resources and network connections.

virtual machine

A virtual machine is a software implementation of a computing environment in which an operating system or program is installed and run.

A virtual machine typically emulates a physical computing environment, requests for CPU, memory, hard disk, network, and other hardware resources that are managed by a virtualization layer which translates these requests to the underlying physical hardware.

virtual machine template

A Virtual Machine Template provides a standardized group of hardware, and software settings that is used repeatedly to create virtual machines configured with those settings.

virtual server image

A virtual server image is the persisted specification and state of a virtual machine. A virtual server is created when you create a guest. The virtual server image contains the general specification of the guest such as CPU, network, memory, and the type of physical storage that is backing the guest. A virtual server image is also referred to as a guest image.

Virtual Local Area Network (VLAN)

VLAN is a group of network resources connected to different network segments that behave as if they were connected to a single network segment. All transmissions from the VLAN are identified by a unique VLAN tag.

volume

A volume is an identifiable unit of data storage that is sometimes physically removable from the computer or storage system. In tape storage systems, a volume may be a tape cartridge. In mainframe storage systems, a volume may be a removable hard disk. Each volume has a system-unique name or number that enables it to be specified by a user.

WINS

WINS stands for Windows Internet Naming Service. The WINS server converts NetBIOS names to IP addresses.

WS-Man

Web Services for Management (WS-MAN) is a specification for managing servers, devices, and applications using web services standards. WS - Man provides a common way for systems to access and exchange management information across the entire IT infrastructure. The public API in Oracle Enterprise Manager Ops Center can be accessed through WS-Management.

World Wide Name (WWN)

WWN is a unique identifier in a Fibre Channel or Serial Attached SCSI storage network. Each WWN is an 8-byte number derived from an IEEE OUI and vendor information.

zone

Also called non-global zones, are a virtualized operating system environment created within a single instance of the Oracle Solaris operating system. One or more applications can run in a non-global zone without interacting with the rest of the system.

A

- ABE See Boot environments, 12-36
- Access points, 2-18
- Accounts
 - private networks, 18-7, 18-14
 - quotas, 18-14
 - vDC, 18-3, 18-12, 18-13, 18-14, 18-15, 18-16, 18-21
- Acknowledge Incident, 9-8, 9-10, 9-16
- Acting on an incident, 9-18
- activate, 12-56
- Add Annotation to Incident, 9-8, 9-10
- Add Assets, 2-4, 2-6
- Add Assets to Group, 2-32
- Add Assets Using Profile, 16-5
- Add Content, 5-8
- Add File System, 14-29
- Add LUNs, 6-11
- Add monitoring rules, 4-7
- Add Storage, 6-12, 14-27
 - logical domains, 15-25
 - Migrate Zone, 14-12
- Add Virtualization Host to Server Pool, 17-21
- Address Allocation Method, 7-16
 - Automatic, 18-23
 - static, 18-23
- Agent Controllers, 2-3
 - access point, 2-18
 - installing, 2-9, 2-13
 - log file, C-2
 - monitoring, 4-1
 - operating systems, 12-9
 - Oracle VM Server for x86, 16-7
 - uninstalling, 2-19
- agentadm
 - requirements, 2-9
- Agentless assets, 2-3
- Agentless-managed operating systems, 12-9, 12-12
 - changing mode, 12-10
- Agent-managed assets, 2-3
- Agent-managed operating systems, 12-9, 12-12
 - changing mode, 12-10
 - updating, 13-3
- Aggregated link, 7-13
- Alert Monitoring Rules, 4-4
- Alert monitors, 4-8
- Alerts, 9-1
 - clearing, 9-5
- All Unassigned Incidents, 9-7
- Analytics
 - operating systems, 12-12, 12-13, 12-14, 12-15, 12-16, 12-17, 12-18, 12-19, 12-23
 - Oracle VM Server for SPARC, 15-17
- Annotations, 9-14
 - Automated Operation, 9-12
 - Comment, 9-12
 - comment, 9-12
 - deleting, 9-12
 - incident, 9-10, 9-17
 - Suggested Action, 9-12
 - viewing, 9-11, 9-13, 9-18
- API for vDC, 18-19
- Asset attributes, 2-18, 2-29
 - adding, 2-29
 - deleting, 2-29
 - viewing, 2-29
- Asset management
 - access points, 2-18
 - adding assets to a group, 2-32
 - Agent Controller, 2-3
 - agentless, 2-3
 - agent-managed, 2-3
 - asset attributes, 2-18
 - changing groups, 2-33
 - credentials, 2-3, 2-16
 - copying, 2-18
 - creating, 2-16
 - deleting, 2-18
 - editing, 2-18
 - updatng, 2-16
 - upgrading, 2-16
 - deleting groups, 2-33
 - global zones, 14-11, 14-12
 - grouping, 2-29, 2-30, 2-31, 2-32
 - incidents, 9-11
 - jobs, 3-6
 - logical domains, 15-29, 15-30
 - Move Storage, 14-28
 - moving groups, 2-33
 - Oracle Services, 9-22
 - Oracle Solaris 11, 2-20
 - Oracle Solaris Clusters, 2-20

- removing assets, 2-19
- removing assets from a group, 2-33
- roles, 2-1
- SPARC SuperCluster, 20-28
- Sun SPARC Enterprise M-Series Servers, 2-20, 2-21
- Sun ZFS Storage Appliances, 2-20, 2-21
- Windows, 2-20
- Assign Incident, 9-8, 9-10, 9-15
- Assign Ownership, 16-11
- Associate Libraries, 17-24
- Associating network domains, 17-22
- Attach Networks, 7-8
 - Oracle VM Server for SPARC, 15-15
 - Oracle VM Server for x86, 16-14
 - server pools, 17-23
- Automated Operation, 9-12
- Automatic Load Balancing Policy, 17-3, 17-4
- Automatic method, 7-16
- Automatic Recovery, 17-4, 17-12
- Available LUNs, 6-10

B

- Badges
 - default library, 5-3
 - incident severity, 9-5
- Balance resources
 - server pools, 17-26
- Bandwidth flow, 7-9
- Bandwidth management, 7-9
- Baseline Analysis Report, 10-9, 13-7
 - black list, 10-10
 - white list, 10-10
- Baseline Check, 20-6
- beadm, 12-37
- Black list, 10-10, 10-13
- Block storage, 6-9
 - dynamic, 6-9
 - LUNs
 - selecting, 6-10
 - specifying, 6-11
 - static, 6-10
- Bonded interface, 16-13, 17-21
- Boolean Control, 4-4
- Boot environments, 12-36
 - activating, 12-39, 12-40
 - active, 12-36
 - alternate, 12-36, 12-50
 - deleting, 12-40
 - dual, 12-36
 - icons, E-3
 - incidents, 12-38
 - Live Upgrade, 12-49
 - monitoring, 12-37
 - Oracle Solaris 10, 12-37, 12-47, 12-55, 12-56
 - Oracle Solaris 10-8, 12-45, 12-47, 12-48, 12-49, 12-50, 12-51, 12-52, 12-53, 12-54
 - Oracle Solaris 11, 12-37, 12-41, 12-42, 12-43, 12-44, 12-45

- policy, 12-49
- profiles, 12-50
- requirements, 12-49
- update profile, 13-16
- updating operating systems, 13-13, 13-16
- viewing, 12-38
- Boot interface, 12-30
- Boot Zone, 14-24
- Branded root zones, 14-2
 - packages, 14-14
 - requirements, 14-14, 14-15

C

- Change History Report, 10-8
- Charts
 - creating, 12-22
 - operating systems, 12-17, 12-21
- Chassis, 4-11
- Cisco switches, 4-11
- CLI
 - access, 19-3
 - command scripts, 19-10
 - general commands, 19-5
 - modes, 19-10
- CLI for vDC, 18-19
- Close Incident, 9-8, 9-10
- Closing an incident, 9-20
- Cloud Admin, 20-8
 - prerequisites, 20-8
- Cloud User, 20-8
 - prerequisites, 20-8
- Cloud user, 18-18
- Cluster file system, 16-17, 17-15
- Cluster Heartbeat, 16-13
- Cluster profiles, B-1, B-2, B-3, B-4
- Command Line Interface
 - See CLI, 19-3
- Comment, 9-12, 9-18
- Common Vulnerability and Exposure
 - See CVE Report, 10-21
- Compare Catalogs, 13-6
- Compute Nodes, 20-18, 20-29, 20-30
- Compute nodes, 20-28
 - SPARC SuperCluster, 20-19
- Configuration
 - jobs, 3-9
- Configure Parent Repositories, 5-7
- Configure YUM Repository, 16-16
 - Oracle VM Manager, 16-16
- Connectivity
 - networks, 7-17
- Connectivity check interval
 - user interface preferences, 1-2
- Console timeout
 - user interface preferences, 1-2
- Consoles
 - logical domains, 15-29
 - Oracle VM Server for x86, 16-6
- Contracts, 9-21, 9-22, 9-23

- Control domains, 15-11, 15-12, 15-13, 15-14
 - networks, 15-14
 - Oracle VM Server for SPARC, 15-2
 - reboot, 15-17
 - storage libraries, 15-14
- Copy Policy, 4-14
- Create Logical Domains, 17-24
- Create Network, 7-7
- Create New NAS Library, 6-17
- Create Policy, 4-14
- Create Private Network, 7-7
- Create Profile
 - Discovery, 16-4
- Create Update Library, 5-8
- Create Virtual Machines, 17-24
- Create Zones, 17-24
- Credentials
 - access points, 2-18
 - asset management, 2-3, 2-16
 - copying, 2-18
 - creating, 2-16
 - deleting, 2-18
 - editing, 2-18
 - upgrading, 2-16
 - ssh, 16-6
- Critical, 4-3, 9-14
- CSV reports, 10-1
- Custom content, 5-14
- CVE Report, 10-21, 13-7

D

- Data link, 7-9
- Database
 - log file, C-2
 - Oracle Exalogic Elastic Cloud, 20-7
- Datacenter Enterprise Controller, 20-5
- Declaring servers, 2-4
 - operating systems, 2-4
 - service processors, 2-5
- Default network domain, 7-5
 - server pools, 17-5
- Default Policy icon, 4-14
- Delete Access Point, 2-18
- Delete Asset, 2-19
- Delete Group, 2-33
- Delete Local Component File, 5-20
- Delete Logical Domain, 15-29
- Delete Selected Guest, 16-26
- Delete Server Pool, 17-27
- Delete Zone, 14-25
- Deleting annotations, 9-12
- Deployment plans, 12-32, 12-33
 - boot environments, 12-52
 - Linux, 12-33
 - logical domains, 15-21, 15-22, 15-23
 - non-global zones, 14-19, 14-20
 - operating systems, 12-33, 12-34
 - Oracle Solaris 11, 12-34
 - Oracle VM Server for SPARC, 15-10

- Oracle VM Server for x86, 16-9
 - provisioning operating systems, 12-35
 - updating operating systems, 13-13
 - updating Oracle Solaris 11, 13-11, 13-12
 - virtual machine, 16-24
- DHCP, 12-25
- Discovering assets, 2-7
- Discovery
 - assets, 2-6
 - logical domains, 15-29
 - Oracle Solaris Clusters, 2-28
 - Oracle VM Manager, 16-3, 16-4
 - Oracle VM Server, 16-5, 16-6
 - Oracle VM Server for x86, 16-6, 16-7
 - requirements, 16-6
 - ssh, 16-6
 - servers, 2-4, 2-5, 2-6
 - SPARC SuperCluster, 20-28
 - storage servers, 16-11, 16-12
 - Sun SPARC Enterprise M-Series Servers, 2-21
 - Sun ZFS Storage Appliances, 2-22, 2-23, 2-25, 2-26, 2-27
- Discovery file, 2-4
- Discovery profile, 2-6
 - creating a profile, 2-7
- Disk image, 5-9
- Distribution Update Report, 10-24
- DLMP, 7-13
- Dynamic block storage, 6-9
- Dynamic Host Configuration Protocol (DHCP), 12-33
- Dynamic private networks, 7-4, 18-6, 18-7
- Dynamic System Domains, 2-21, 4-11

E

- EC Library, 5-3
 - changing, 5-3
 - Initial EC Library, 5-3
 - updating, 5-3
- Edit Attributes, 16-15
- Edit Group, 2-32
- Edit IPMI Configuration, 16-15
- Edit Local Component File, 5-20
- Editing monitoring rules, 4-5, 4-6
- Engineered Systems
 - See Oracle Engineered Systems, 20-1
- Enterprise Controller
 - Oracle Engineered Systems, 20-5
- Enumerated Control, 4-4
- Ethernet network, 7-1
- Exadata Storage Servers, 20-27, 20-30
 - SPARC SuperCluster, 6-20, 20-19
- Exalogic, 20-6, 20-7
 - Cloud, 20-2
- Exalogic Control, 20-6, 20-7
 - Controls tab, 20-12
 - database, 20-7
 - Incidents tab, 20-13
 - Networks tab, 20-12, 20-13

- Exalogic Rack View, 20-14, 20-15
 - actions, 20-16
- Exalogic Report, 10-28
- Exalogic System Admin, 20-8
 - prerequisites, 20-8
- Exalogic System Report, 20-17
 - Baseline Check, 20-2, 20-18
 - Compute Nodes, 20-18
 - Power distribution units, 20-18
 - Report Parameters, 20-17
 - Storage Appliances, 20-18
 - Summary Table, 20-17
 - Switches, 20-18
 - System Control Software, 20-18
 - Validation, 20-2, 20-18
- Exalogic System View, 20-9, 20-10, 20-11, 20-12, 20-13
 - actions, 20-14
- Exclusive IP mode, 7-9, 14-4, 17-11
- Exporting log files, 3-6
- Expression, 4-4
- Expression Monitoring Rule Query Language, G-1
- eXtended System Control Facility
 - See XSCF, 2-21
- Extract a Monitor Policy, 4-13

F

- Fabrics
 - fully-managed, 7-4, 18-6
 - host-managed, 7-4
 - physical, 7-3
 - unmanaged, 7-4, 18-6
- Failure policy, 3-2
- Fibre Channel SAN, 6-18
- File servers, 4-11, 6-18
- File system libraries, 5-4
- File systems
 - boot environments, 12-43, 12-48
- Find Assets, 2-6
- Firewall, 7-26
- Firmware
 - power distribution units, 5-13
 - Profile for PDU, 5-14
- Firmware Compliance Report, 10-5, 11-23
- Firmware images, 5-10
 - metadata, 5-10
 - PDU, 5-13
- Formatting reports, 10-1
- Fully-managed fabrics, 7-4, 18-6

G

- Generate Report, 10-3
- Global zones, 14-2
 - Attach Network, 14-6
 - deleting, 14-12
 - discovering, 14-11
 - IPMP groups, 14-7
 - Link aggregation, 14-7

- monitoring policies, 4-11
- networks, 14-4, 14-10, 14-11
- requirements, 14-3
- storage libraries, 14-4, 14-12, 14-27
- updating operating systems, 14-41
- GnuPG key, 16-16
- Groups
 - adding assets, 2-32
 - attributes, A-1
 - changing groups, 2-33
 - changing position, 2-33
 - creating groups, 2-31
 - deleting, 2-33
 - editing, 2-32
 - icons, E-5
 - removing assets, 2-33
 - system groups, 2-30
 - tags, 2-29
 - user-defined, 2-30
 - viewing data, 2-31
- Guest metadata, 5-9
- Guests
 - See virtual hosts, 17-1
- GUID
 - LUN, 6-12
 - WWN, 6-12

H

- Halt Zone, 14-25
- Hardware
 - icons, E-6
- High Availability
 - storage, 6-20
- Historical data, 4-3, 4-6
- Host Compliance Report, 10-20, 13-7
- Host-managed fabrics, 7-4

I

- IaaS, 18-2
- Icons
 - Acknowledge Incident, 9-8, 9-10
 - Add Annotation to Incident, 9-8, 9-10
 - Add Assets Using Profile, 16-5
 - Add Content, 5-8
 - Administration, E-1
 - annotation, 9-10
 - Assign Incident, 9-8, 9-10
 - Boot Environment, E-3
 - Close Incident, 9-8, 9-10
 - Connectivity, E-4
 - Default Policy, 4-14
 - Edit View, 10-3
 - Generate Report, 10-3
 - Group, E-5
 - Hardware, E-6
 - Incident, E-8
 - Jobs, E-9
 - Kernal, E-10

- Mark Incident as Repaired, 9-8, 9-10
- Message Center, E-11
- Monitoring, E-11
- Network, E-13
- Open Service Request, 9-10
- Operating System, E-13
- Processes, E-15
- Report, E-15
- Software Library, E-10
- SuperCluster, E-17
- Switch, E-18
- Take Action on Incident, 9-8, 9-10
- Unbind Network, 14-11
- Universal, E-19
- Update Policy, E-20
- Update Profiles, E-20
- URL, 9-7
- View Alerts, 9-8, 9-9
- View Annotations, 9-8, 9-9
- View Comments, 9-8, 9-10
- View Possible Impacts and Causes, 9-8, 9-9
- View Suggested Actions, 9-8, 9-10
- Virtualization, E-14
- wrench, 12-21
- Image Packaging System, 5-5
- Images
 - disk image, 5-9
 - HVM, 16-22
 - import, 5-11
 - ISO, 5-9
 - Oracle VM Server for x86, 16-7
 - OS, 5-10
 - PVM, 16-22
 - upload, 5-11
 - virtual machines, 16-22
- Import images, 5-11
- Incident Compliance Report, 10-17, 10-18, 13-7
- Incident Detail Report, 10-26
- Incident Severity Badges, 9-5
 - Asset icon, 9-5
- Incident Summary Report, 10-25
- Incident tab for vDC, 18-11
- Incidents, 9-1
 - Acknowledge Incident, 9-8
 - acknowledging, 9-16
 - acting, 9-18
 - Add Annotation to Incident, 9-8
 - alerts, 9-4
 - annotating, 9-17
 - annotations, 9-10
 - Assign Incident, 9-8
 - Assigned to Others, 9-6
 - assigning, 9-15
 - boot environments, 12-38
 - Close Incident, 9-8
 - closing, 9-20, 20-6
 - comments, 9-18
 - Critical, 9-14
 - details, 9-15
 - icons, 9-6, E-8
 - ID, 9-7
 - Informational, 9-14
 - Mark Incident as Repaired, 9-8
 - Message Center, 9-6
 - monitoring rules, 9-13, 9-14
 - My Incidents, 9-6
 - My Service Requests, 9-6
 - new, 9-7
 - Notifications, 9-6
 - Open Service Requests, 9-6
 - Oracle Engineered Systems, 20-6
 - Oracle Exalogic Elastic Cloud, 20-13
 - Relayed, 9-6, 20-6
 - Relayed Service Requests, 9-7
 - repaired, 9-20
 - roles, 9-2
 - Service Requests Opened by Others, 9-6
 - SPARC SuperCluster, 20-24
 - Take Action on Incident, 9-8
 - unassigned, 9-6, 9-7
 - unresolved, 9-15
 - View Alerts, 9-8
 - View Annotations, 9-8
 - View Comments, 9-8
 - View Possible Impacts and Causes, 9-8
 - View Suggested Actions, 9-8
 - viewing, 20-6
 - viewing annotations, 9-18
 - Warning, 9-14
- Incidents Assigned to Others, 9-6
- Incidents Knowledge Base, 9-11, 9-12, 9-14
 - annotations, 9-11, 9-12
 - operational plans, 9-11, 9-12
- Incidents tab, 9-14
- InfiniBand network, 7-1
- Info, 4-3
- Informational, 9-14
- Infrastructure-as-a-Service, 18-2
- Initial EC Library, 5-3
- Initialize Oracle Solaris 11 Software Update Library, 5-6
- Installing operating system
 - See Provisioning, 12-24
- Intelligent Platform Management Interface
 - See IPMI, 16-15
- Interactive reports, 10-1
- Internet Connection Firewall., 2-20
- IP Multipathing, 7-10
- IPMI, 16-15
 - Wake-on-LAN, 16-15
- IPMP groups, 7-10, 14-5
 - creating, 7-11, 7-12, 14-7, 14-8
 - global zones, 14-7
 - Oracle VM Server for SPARC, 15-15
- IPS, 5-5, 5-8
- IPv4, 7-15
- IPv6, 7-15
- iSCSI storage arrays, 4-11
- ISO image, 5-9

J

JET

See JumpStart Enterprise Toolkit, 12-31

JetFLASH, 12-31

Job status popup duration

user interface preferences, 1-2

Jobs

actions, 3-8

Actual Run, 3-2

by asset, 3-6

configuring, 3-9

copying, 3-9

debugging, 3-9

deleting, 3-9

details, 3-5

event logs, 3-6

failure policy, 3-2

icons, E-9

order of tasks, 3-2

repeating, 3-8

roles, 3-3

Run ID, 3-2

searching, 3-7

Simulated, 3-2

status, 3-1

stopping, 3-8

targets, 3-5

tasks, 3-6

Update Microsoft Windows, 13-19

updating operating systems, 13-14

viewing, 3-4, 3-5, 3-8

Jobs pane, 3-1

JumpStart Enterprise Toolkit, 12-31, 12-32, D-1

location, D-1, G-5

parameters, D-1

templates, 12-32

K

Kernal

icons, E-10

Knowledge Base, 5-3

L

Launch Web Interface, 20-6

LDAP, 20-7

Libraries

images, 5-1

local content, 5-1

software, 5-1

storage libraries, 6-1

Link aggregation, 7-13

Configure Bonding, 17-21

creating, 7-14, 14-9, 14-10

global zones, 14-7

Networks

Link aggregation, 14-6

Oracle VM Server for SPARC, 15-9, 15-15

Linux, 13-1, 13-13

update policies, 13-8

update profiles, 13-9

updates, 13-3

Linux and Oracle Solaris 8-10 Software Update Library, 5-8

Linux provisioning parameters, 12-29

Linux See operating systems, 13-1

Linux SuSE provisioning parameters, 12-29

Linux, Solaris 8-10 Software Update Library, 5-8

Live Migrate, 16-13

Live Upgrade, 12-45, 12-47, 12-49, 12-50, 12-51, 12-55, 12-56

updated operating systems, 12-54

updating operating systems, 13-13

Local actions, 5-15

Local Categories, 5-15

Local Component File, 5-20

Local Configuration File, 5-17

Local content, 5-14

actions, 5-15

categories, 5-15

changing, 5-20

configuration files, 5-17

editing, 5-20

packages, 5-16

Local libraries, 5-4

monitoring policies, 4-11

Local Packages, 5-16

Local storage

File system storage, 6-8, 6-9

Logging events, 3-6

Logical domains

console, 15-29

deleting, 15-29

deployment plans, 15-21, 15-22, 15-23

discovering, 15-29

migrating, 15-26, 15-27

modifying, 15-24

monitoring policies, 4-11

networks, 15-14, 15-25

operating systems, 15-22

Oracle Engineered Systems, 20-19

Oracle VM Server for SPARC, 15-2, 15-18, 15-19,

15-20, 15-21, 15-22, 15-23, 15-24, 15-25, 15-26,

15-27, 15-29, 15-30

profiles, 15-19, 15-20

server pools, 17-24, 17-26

shutdown, 15-24

start, 15-24

storage libraries, 15-14, 15-25

Logical Domains Manager, 15-11

Logical Units

See LUNs, 6-18

Logs

cacao, C-3

events, C-1

installation, C-1

upgrades, C-3

Lowest allocated CPU and memory resources, 17-15

Lowest relative load, 17-14

lucreate, 12-37, 12-49

LUNs, 6-3, 6-9

 Add LUNs, 6-11

 selecting, 6-10

 specifying, 6-11

 volume groups, 6-18

M

Macros, 5-15

Maintenance mode

 Incidents

 disabling, 9-19

 migrating virtual hosts, 17-4

Managed networks

 Subnets, 7-6

Manual Net Boot, 12-30, 15-13

Manually Adding LUNs, 6-11

Mark an Incident as Repaired

 Repaired incidents, 9-20

Mark Incident as Repaired, 9-8, 9-10

Maximum Transmission Unit, 7-16

Membership graph preferences

 user preferences, 1-2

Message Center, 9-2, 9-14

 annotations, 9-11

 icons, E-11

 incidents, 9-6, 20-6

 My Incidents, 9-8

 Notifications, 9-10

 relayed service requests, 20-6

Metadata, 6-3

 firmware, 5-10

 guest, 5-9

Microsoft System Center Configuration

 Manager, 13-19

Microsoft Windows

 registry, 13-19

 update job, 13-20

 updates, 13-19, 13-20

 updating, 13-19

Migrate Logical Domain, 15-26, 15-27

Migrate Zones, 14-32

Migrating logical domains

 server pool, 17-26

Migrating virtual machines

 server pool, 17-26

Migrating zones, 14-30, 14-31, 14-34, 14-35, 14-36, 14-37

 server pool, 17-25

Minimize Power Consumption Policy, 17-3, 17-15

Modify Physical Connectivity, 14-10

 Networks

 Oracle VM Server for SPARC, 15-16

Monitored attributes, 4-7, A-1

 collection, 4-7

 simple, 4-7

 struct-like, 4-7

 structure, 4-7

Monitoring

 boot environments, 12-37

 icons, E-11

 operating systems, 12-11, 12-16, 12-21

 Oracle VM Server for SPARC, 15-17

 roles, 4-1

 thresholds, 4-3

Monitoring policies, 4-1, 4-3, 4-8, 12-11

 chassis, 4-11

 Cisco switches, 4-11

 copying, 4-14

 creating, 4-8, 4-13, 4-14

 default, 4-9, 4-14

 deleting, 4-14

 details, 4-10

 Dynamic System Domains, 4-11

 extracting, 4-9, 4-13

 file servers, 4-11

 global zones, 4-11

 groups of assets, 4-12

 iSCSI storage arrays, 4-11

 local libraries, 4-11

 logical domains, 4-11

 modifying, 4-12

 M-Series servers, 4-11

 NAS libraries, 4-11

 non-global zones, 4-11

 operating systems, 4-11

 Oracle VM Server for SPARC, 4-11

 Oracle VM Server for x86, 4-11

 power distribution units, 4-11

 Remote Oracle Engineered System, 4-11

 rule set, 4-8

 SAN libraries, 4-11

 SAN storage arrays, 4-11

 server pools, 4-12

 Servers, 4-12

 Solaris Cluster, 4-12

 Solaris Cluster Node, 4-12

 Solaris Cluster Zone Cluster Group, 4-12

 Solaris Cluster Zone Cluster Node, 4-12

 Storage, 4-12

 switches, 4-12

 system-defined, 4-9, 4-11

 user-defined, 4-9

 viewing assets, 4-14

 Virtual machines, 4-12

Monitoring rules, 4-1, 4-3, 12-11

 active, 4-4, 4-5

 adding, 4-7

 Critical, 4-3

 disabled, 4-4

 editing, 4-5, 4-6

 enabled, 4-4

 inactive, 4-4, 4-5

 incidents, 9-13, 9-14

 Info, 4-3

 state, 4-4

 system-defined, 4-3

 thresholds, 4-6

 user-defined, 4-3, 4-4

- Using Expression monitoring rules, 4-7
- Warning, 4-3
- Monitoring tab, 4-3
 - details, 4-10
- Move Asset to Group, 2-33
- Move Group, 2-33
- Move Storage, 14-12, 14-27
- MPxIO, 6-18
- M-Series servers, 2-21
 - monitoring policies, 4-11
- Multipathing, 6-18
- My Incidents, 9-6, 9-8
 - Acknowledge Incident, 9-10
 - Add Annotation to Incident, 9-10
 - Assign Incident, 9-10
 - Close Incident, 9-10
 - Mark Incident as Repaired, 9-10
 - new, 9-8
 - Open Service Request, 9-10
 - Take Action on Incident, 9-10
 - unassigned, 9-8
 - View Alerts, 9-9
 - View Annotations, 9-9
 - View Comments, 9-10
 - View Possible Impacts and Causes, 9-9
 - View Suggested Actions, 9-10
- My Oracle Support, 9-2
- My Service Requests, 9-6

N

- NAS libraries, 17-24
- NAS software libraries, 5-4
- NAS storage
 - monitoring policies, 4-11
 - NFS, 6-4
- NAS Storage Appliances, 6-16
- NAS storage libraries, 6-4
 - creating, 6-17
- Netboot
 - virtual machines, 16-22
- Network domains, 7-4
 - default, 7-5
 - server pools, 17-21, 17-22, 17-23
 - user-defined, 7-5
 - vDC, 18-7
 - virtual datacenter, 18-6
 - virtual machines, 16-23
- Network File Service, 6-4
- Network Interface Card, 7-18
- Network-attached storage
 - See NAS, 6-4
- Networks, 7-6
 - Connectivity, 7-17
 - exclusive IP mode, 14-4
 - firewall, 7-26
 - global zones, 14-4, 14-10, 14-11
 - icons, E-13
 - IPMP groups, 14-5
 - logical domains, 15-25
 - MTU, 7-16
 - non-global zones, 14-30
 - Oracle Enterprise Manager Ops Center
 - operations, 7-19
 - Oracle Solaris Zones, 7-8
 - Oracle VM Server for SPARC, 7-9, 15-14, 15-15
 - Oracle VM Server for x86, 7-9, 16-13, 16-14
 - roles, 16-13
 - ports, 7-26
 - profiles, 7-19
 - properties, 7-14
 - protocols, 7-26
 - requirements, 7-7
 - roles, 7-2
 - routing mode, 7-15
 - server pools, 7-8, 17-5, 17-11, 17-23
 - shared IP mode, 14-4
 - static route, 7-15
 - switches, 7-18
 - utilization, 7-17
 - vDC, 7-9, 18-10, 18-22, 18-23
 - virtual datacenter, 18-4, 18-5
 - VLAN, 7-17
- New Update OS Job, 13-14
- NFS Client, 6-16
- NFS Server, 6-16
- NIC See also PNIC, 7-18
- Non-global zones, 14-2
 - Add File System, 14-29
 - Add Storage, 14-27
 - Boot Zone, 14-24
 - branded root zones, 14-2
 - clones, 14-25
 - creating, 14-13
 - Delete Zone, 14-25
 - deleting, 14-12
 - deployment plans, 14-19, 14-20
 - file systems, 14-21
 - Halt Zone, 14-25
 - migrating, 14-30, 14-31, 14-32, 14-34, 14-35, 14-36, 14-37
 - modifying, 14-23
 - monitoring policies, 4-11
 - networks, 14-30
 - Oracle Solaris 11 Software Update Library, 14-15
 - profiles, 14-13, 14-15, 14-16
 - Reboot Zone, 14-25
 - recovering, 14-38, 14-39
 - Replicate Zone, 14-25
 - requirements, 14-13, 14-14
 - server pools, 14-40, 17-24, 17-25
 - Shutdown Zone, 14-25
 - sparse root zones, 14-2
 - storage, 14-29
 - targets, 14-17
 - updating operating systems, 14-40, 14-42
 - updating operating systems in parallel, 14-43
 - whole root zones, 14-2
 - zpool, 14-21
- Notifications, 9-6, 9-10

O

OCDDoctor

- debugging a job, 3-9

Open Service Requests, 9-6, 9-10

Operating system updates

- roles, 13-4

Operating systems

- agentless-managed, 12-9, 12-12

 - changing mode, 12-10

- agent-managed, 12-9, 12-12

 - changing mode, 12-10

- Analytics, 12-12, 12-13

 - charts, 12-17

 - custom, 12-15, 12-23

 - Summary, 12-14

- boot environments, 12-36, 12-37, 12-38, 12-39, 12-40, 12-41, 12-42, 12-43, 12-44, 12-45, 12-47, 12-48, 12-49, 12-50, 12-51, 12-52, 12-53, 12-54

- charts, 12-21

- CPU Utilization, 12-14

- deployment plans, 12-33, 12-34, 12-35

- History, 12-18

- icons, E-13

- logical domains, 15-22

- Metrics, 12-19

- Microsoft Windows, 13-19, 13-20

- monitor thresholds, 12-16

- monitoring, 12-11

- monitoring policies, 4-11

- Oracle VM Server for x86, 16-7, 16-9

- parameters, 12-29

- Processes, 12-15

- profiles, 12-8, 12-29, 12-30, 12-31, 12-35

- provisioning, 12-24, 12-32

- roles, 12-3

- Services, 12-16

- system catalog, 13-6

- threshold, 12-21

- update job, 13-14

- update policies, 13-7, 13-8

- update profiles, 13-9, 13-10

- updates, 13-3, 13-4, 13-11, 13-12, 13-13, 13-14, 13-15, 13-16, 13-19

 - requirements, 13-3

- vDC, 18-23

- Virtualization Analytics, 12-23

- zones, 14-40, 14-41, 14-42, 14-43

operating systems

- status, 12-7

Operational plans, 4-3, 9-14

- boot environments, 12-52

- incidents, 9-11, 9-12

Oracle 8,9,10, 12-33

Oracle Engineered Systems, 20-1

- Baseline Check, 20-6

- incidents, 20-6

- Launch Web Interface, 20-6

- logical domains, 20-19

- monitoring policies, 4-11

- Oracle Exalogic Elastic Cloud, 20-2, 20-6, 20-7,

 - 20-8, 20-9, 20-10, 20-14, 20-15, 20-16, 20-17, 20-18

- rack, 20-5

- service requests, 20-6

- SPARC SuperCluster, 20-19, 20-20

 - actions, 20-25

 - add assets, 20-28

 - Baseline Check, 20-30

 - compute nodes, 20-28, 20-29, 20-30

 - Dashboard, 20-21, 20-22

 - Details, 20-23

 - discovery, 20-28

 - Exadata Storage Server, 20-27

 - Incidents, 20-24

 - Infrastructure Networks, 20-24

 - Network Connectivity, 20-24

 - Networks, 20-24

 - power distribution units, 20-30

 - reports, 20-29

 - storage, 20-27, 20-28, 20-30

 - switches, 20-28, 20-30

 - validation, 20-30

 - viewing, 20-20

 - viewing racks, 20-25, 20-26

- SPARC SuperCluster T4-4, 20-19

- status, 20-5

Oracle Enterprise Manager Ops Center Software

- Developer's Kit (SDK), F-1

- SUNWxvmoc-sdk.pkg, 4-7

Oracle Exalogic Elastic Cloud, 20-2, 20-6, 20-7, 20-8

- actions, 20-14

- Baseline Check, 20-2, 20-18

- compute nodes, 20-18

- Controls, 20-12

- Dashboard, 20-10

- Details, 20-11

- Incidents, 20-11, 20-13

- Infrastructure Networks, 20-13

- Membership Graph, 20-10

- Network Connectivity, 20-13

- Networks, 20-12

- power distribution units, 20-18

- reports, 20-17

- software, 20-18

- Status, 20-11

- storage appliances, 20-18

- Summary, 20-10

- switches, 20-18

- validation, 20-2, 20-18

- viewing, 20-9

- viewing racks, 20-14, 20-15, 20-16

Oracle Knowledge Base, 5-3

Oracle Services, 9-21

- contracts, 9-21, 9-22, 9-23

- icons, E-4

- requirements, 9-22

- service requests, 9-21, 9-23, 9-24

- warranty, 9-21, 9-22

Oracle Solaris

- update policies, 13-8

- update profiles, 13-9
- updates, 13-3
- Oracle Solaris 10
 - boot environments, 12-37, 12-47, 12-55, 12-56
 - branded root zones, 14-14
 - non-global zones, 14-14
 - sparse root zones, 14-14
 - updates, 5-8
 - whole root zones, 14-14
- Oracle Solaris 10-8, 13-13
 - boot environments, 12-45, 12-47, 12-48
 - profiles, 12-49
- Oracle Solaris 11, 13-11
 - boot environment, 13-16
 - boot environments, 12-37, 12-41, 12-42, 13-16
 - file systems, 12-43
 - profiles, 12-44
 - snapshots, 12-43, 12-45
 - zones, 12-43
 - branded root zones, 14-15
 - non-global zones, 14-14
 - updates, 13-11, 13-12
- Oracle Solaris 11 deployment plan, 12-35
- Oracle Solaris 11 Package Repository, 5-3
- Oracle Solaris 11 provisioning profiles, 12-35
- Oracle Solaris 11 Software Update Library, 5-6, 13-11
 - adding content, 5-8
 - alternate repository, 5-8
 - deleting, 5-8
 - initializing, 5-6
 - non-global zones, 14-15
 - parent repositories, 5-6
 - status, 5-6
 - synchronizing, 5-7
 - viewing, 5-7
- Oracle Solaris 8-10 Software Update Library, 5-4
- Oracle Solaris Cluster
 - monitoring policy, 4-12
 - upgrading, B-1
- Oracle Solaris Cluster node
 - monitoring policy, 4-12
- Oracle Solaris Cluster Zone Cluster Groups
 - monitoring policy, 4-12
- Oracle Solaris Cluster Zone Cluster Node
 - monitoring policy, 4-12
- Oracle Solaris Clusters
 - discovery, 2-28
 - importing, B-4
 - uploading, B-3
- Oracle Solaris Container
 - See Oracle Solaris Zones, 14-1
- Oracle Solaris OS, 13-1
- Oracle Solaris See operating systems, 13-1
- Oracle Solaris Update Compliance Report, 10-16, 13-7, 13-15
- Oracle Solaris Zones, 14-1
 - boot environments, 12-43
 - Live Upgrade, 12-47, 12-55, 12-56
 - networks, 7-8
- roles, 14-2
- server pools
 - adding virtualization hosts, 17-20
 - creating, 17-12
 - modifying, 17-20
 - networks, 17-11
 - policies, 17-12
 - requirements, 17-11
 - storage libraries, 17-24
- Oracle VM Manager, 16-1
 - details, 16-10
 - discovering, 16-3
 - discovery profile, 16-3, 16-4, 16-5
 - file system, 16-11, 16-12
 - iSCSI storage, 16-11, 16-12
 - Oracle Engineered Systems, 20-7
 - Oracle VM Server, 16-5
 - ownership, 16-11
 - password, 16-7
 - SAN storage, 16-11, 16-12
 - server pools, 16-17
 - storage, 16-11, 16-12
 - YUM repository
 - GPG key, 16-16
- Oracle VM Server, 16-6
 - block storage, 16-12, 16-13
 - file system, 16-12
 - networks, 7-9
 - storage libraries, 16-12, 16-13
- Oracle VM Server for SPARC, 15-1, 15-13
 - Analytics, 15-17
 - Attach Networks, 15-15
 - control domains, 15-2, 15-11, 15-12, 15-13, 15-14
 - deployment plans, 15-10
 - details, 15-17
 - installing, 15-11, 15-12
 - IPMP groups, 15-15
 - Link aggregation, 15-9, 15-15
 - logical domains, 15-2, 15-18, 15-19, 15-20, 15-21, 15-22, 15-23, 15-24, 15-25, 15-26, 15-27, 15-29, 15-30
 - Modify Physical Connectivity, 15-16
 - modifying, 15-14
 - monitoring, 15-17
 - monitoring policies, 4-11
 - networks, 15-14, 15-15
 - operating systems, 12-31, 15-22
 - profile, 15-5, 15-6, 15-13
 - reboot, 15-17
 - requirements, 15-3, 15-4, 15-5, 15-12, 15-13
 - roles, 15-2
 - server pools, 15-10, 15-18
 - adding virtualization hosts, 17-20
 - creating, 17-8
 - networks, 17-6
 - policies, 17-6
 - requirements, 17-5
 - storage, 17-6
 - service processor, 15-13
 - storage libraries, 15-14, 17-24

- Unbind Network, 15-16
- Oracle VM Server for x86, 16-1, 17-1
 - Agent Controller, 16-7
 - deployment plans, 16-9
 - discovery profile, 16-6, 16-7
 - editing, 16-15
 - image, 16-7
 - IPMI, 16-15
 - monitoring policies, 4-11
 - networks, 16-13, 16-14
 - management network, 16-13
 - ownership, 16-11
 - Place in Maintenance Mode, 16-15
 - profile, 16-7
 - roles, 16-2
 - server pools, 16-11, 16-17, 17-15
 - adding virtualization hosts, 17-20
 - cluster file system, 16-17, 17-15
 - Create Server Pool, 16-18
 - creating, 17-17
 - policies, 16-17, 17-16
 - requirements, 16-17
 - Server Pool Master, 16-17, 17-16
 - storage libraries, 17-24
 - updating, 16-16
 - virtual machines, 16-20
 - creating, 16-21
 - creating from a virtual machine, 16-22
 - creating from image, 16-22
 - creating from network, 16-22
 - creating from templates, 16-21
 - profiles, 16-21
- Oracle VM Server Host, 15-11
- Oracle VM Storage Repositories, 6-6, 16-12, 16-23, 17-24
- OS image, 5-10
- OS Update, 13-7

P

- Package Compliance Report, 10-25, 13-7
- Parent repositories, 5-3, 5-4
 - Oracle Solaris 11 Software Update Library, 5-6
 - synchronizing, 5-7
- Partition
 - partition keys, 7-3
- PDF reports, 10-1
- PDU
 - firmware, 5-13
 - profile, 5-14
- Physical fabric, 7-3
- Physical network interface card, 7-6
- Place in Maintenance Mode
 - Maintenance mode, 16-15
- Placement Policy, 17-2
 - Lowest allocated CPU and memory resources, 17-15
 - Lowest relative load, 17-14
 - Minimize power consumption, 17-15
- Plan Management

- Incident Knowledge Base, 9-11
- Monitoring policies, 4-9
- Plans
 - provisioning, 12-32
- PNIC, 7-6
- Policies
 - Automatic Load Balancing Policy, 17-3
 - thresholds, 17-4
 - Minimize Power Consumption, 17-3
 - monitoring, 4-3, 4-8, 4-12
 - Oracle Solaris Zones, 17-12
 - Oracle VM Server for x86, 17-16
 - Placement Policy, 17-2
 - server pools, 16-17, 17-2, 17-3, 17-4, 17-6, 17-12, 17-16
 - update, 13-7, 13-8
- Policy
 - boot environments, 12-49
- Ports, 7-26
 - Ethernet switch
 - InfiniBand switch, 18-6
- Power Distribution Units, 20-30
- Power distribution units
 - monitoring policies, 4-11
- Private networks, 7-4, 7-7, 18-14, 18-26, 18-27
 - accounts, 18-7
 - dynamic, 7-4, 18-6, 18-7
 - Physical fabric, 18-6
 - requirements, 18-7
 - static, 7-4, 18-7
- Private vNets, 18-7, 18-26
 - allocating vIP, 18-27
 - creating, 18-27
 - managing, 18-27
- Processes
 - icons, E-15
- Profile Analysis Report, 10-12, 13-7
 - black list, 10-13
 - white list, 10-13
- Profiles, 15-13
 - boot environments, 12-44, 12-49, 12-50, 12-53, 13-16
 - clusters, B-1, B-2, B-3, B-4
 - discovery, 2-6, 2-7, 16-12
 - logical domains, 15-19, 15-20
 - network, 7-19
 - non-global zones, 14-15, 14-16
 - operating systems, 12-8
 - Oracle VM Manager, 16-3, 16-4
 - Oracle VM Server for SPARC, 15-5, 15-6
 - Oracle VM Server for x86, 16-6, 16-7
 - provisioning operating systems, 12-29, 12-35
 - boot interface, 12-30
 - JumpStart Enterprise Toolkit, 12-31
 - Linux parameters, 12-29
 - Linux SuSE parameters, 12-29
 - Manual Net Boot, 12-30
 - Oracle Solaris profiles, 12-30
 - Oracle VM Server for SPARC, 12-31
 - storage, 6-20

- system catalog, 13-16
- update, 13-9, 13-10
- updating operating systems, 13-14
- virtual machines, 16-21
- zones, 14-15, 14-16
- Protocols, 7-26
- Provisioning
 - operating systems, 12-24
- Provisioning operating systems
 - hardware virtualization, 16-21
 - hardware virtualization with paravirtual
 - drivers, 16-21
 - HVM, 16-21
 - HVM with PV, 16-21
 - Oracle VM Server for x86, 16-7, 16-9
 - paravirtualization, 16-21
 - PVHVM, 16-21
 - PVM, 16-21
 - virtual machines, 16-25
- Provisioning plans, 12-32
- Public API
 - API, F-4
 - JMX, F-6, F-7
 - WS-Management, F-6, F-7
- Public external networks, 18-7
- Public networks, 7-4, 7-7, 18-26
- Publisher, 5-3

R

- Racks
 - Oracle Engineered Systems, 20-5
 - Oracle Exalogic Elastic Cloud, 20-14, 20-15, 20-16
 - SPARC SuperCluster, 20-19, 20-25, 20-26
- RAID Controller, 6-15
- Reboot Guests, 16-26
- Reboot Zone, 14-25
- Recommended Software Configuration
 - Report, 10-15, 13-7
- Recovering zones, 14-38, 14-39
- Relayed incidents, 9-6, 20-6
 - closing, 20-6
 - viewing, 20-6
- Relayed Service Requests
 - viewing, 20-6
- Relayed service requests, 9-7, 20-6
- Release Ownership, 16-11
- Remove Asset from Group, 2-33
- Replicate Zone, 14-25
- Report Result
 - updating operating systems, 13-15
- Reports
 - Baseline Analysis Report, 10-9, 13-7
 - Change History Report, 10-8
 - CVE Report, 10-21, 13-7
 - Distribution Update Report, 10-24
 - Exalogic Report, 10-28
 - Exalogic System Report, 20-2, 20-17, 20-18
 - Firmware Compliance Report, 10-5, 11-23
 - formatting, 10-1
 - Generate Report, 10-3
 - Host Compliance Report, 13-7
 - Host Compliance Report for Microsoft
 - Windows, 10-20
 - Host Compliance Report for Oracle Solaris or
 - Linux, 10-20
 - icons, E-15
 - Incident Compliance Report, 10-17, 13-7
 - Incident Compliance Report for Microsoft
 - Windows, 10-18
 - Incident Compliance Report for Oracle Solaris or
 - Linux, 10-17
 - Incident Detail Report, 10-26
 - Incident Summary Report, 10-25
 - Oracle Solaris Update Compliance Report, 10-16,
 - 13-7, 13-15
 - OS Update, 13-7
 - Package Compliance Report, 10-25, 13-7
 - Profile Analysis Report, 10-12, 13-7
 - Recommended Software Configuration
 - Report, 10-15, 13-7
 - Results, 10-3
 - roles, 10-2
 - scheduling, 10-1
 - Service Pack Compliance Report, 10-24
 - SPARC SuperCluster System Report, 20-29, 20-30
 - System Catalog, 13-6
 - System Catalog Report, 10-23
 - System Information Report, 10-27
 - templates, 10-2, 10-3
- Repositories
 - Oracle Solaris 11 Package Repository, 5-3
 - Oracle VM Server, 6-6
 - parent, 5-3, 5-4
 - publisher, 5-3
 - YUM, 16-16
- requirements, 13-3
- Roles
 - asset management, 2-1
 - Cloud Admin, 20-8
 - Cloud User, 20-8
 - Exalogic System Admin, 20-8
 - incidents, 9-2
 - jobs, 3-3
 - monitoring, 4-1
 - networks, 7-2
 - operating system updates, 13-4
 - operating systems, 12-3
 - Oracle Solaris Zones, 14-2
 - Oracle VM Server for SPARC, 15-2
 - Oracle VM Server for x86, 16-2
 - Reports, 10-2
 - server pools, 17-2
 - Software libraries, 5-1
 - storage libraries, 6-1
 - SuperCluster System Admin, 20-20
 - vDC Cloud, 18-19
 - virtual datacenters, 18-3
- roles, 13-4
- Routing mode, 7-15

Rule set, 4-8
Rules for monitoring, 4-3
Run ID, 3-2

S

SAN libraries, 17-24
SAN storage, 6-3
 monitoring policies, 4-11
SAN storage arrays
 monitoring policies, 4-11
SCCM See Microsoft System Center Configuration Manager, 13-19
Scheduling reports, 10-1
Scripts, 5-15
See NFS, 6-4
Server Management, 16-13
Server Pool Master, 16-17, 17-16
Server pools, 17-1
 adding virtualization hosts, 17-21
 Automatic Load Balancing Policy, 17-3, 17-4
 Automatic Recovery, 17-4
 balance resources, 17-26
 deleting, 17-27
 logical domains, 17-24, 17-26
 modifying, 17-19
 monitoring policies, 4-12
 network domains, 17-21, 17-22, 17-23
 networks, 7-8, 17-5, 17-23
 non-global zones, 17-24, 17-25
 Oracle Solaris Zones, 17-11, 17-12, 17-20
 Oracle VM Manager, 16-17
 Oracle VM Server for SPARC, 15-10, 15-18, 17-5, 17-6, 17-8, 17-20
 Oracle VM Server for x86, 16-11, 16-17, 16-18, 17-15, 17-16, 17-17, 17-20
 Placement Policy, 17-2
 policies, 17-2, 17-3, 17-4, 17-6
 roles, 17-2
 software libraries, 17-5
 storage libraries, 6-5, 17-5, 17-24
 vDC, 18-10, 18-11
 virtual datacenter, 18-4
 virtual hosts, 17-2
 virtual machines, 17-24, 17-26
 virtualization hosts, 17-2
 zones, 14-40
Servers
 monitoring policies, 4-12
Service Pack Compliance Report, 10-24
Service Requests
 Oracle Engineered Systems, 20-6
 Relayed, 20-6
Service requests, 9-21, 9-23, 9-24, 20-6
 automated, 9-24
 creating, 9-24
Service Requests Opened by Others, 9-6
Service tags, 2-6
Session timeout
 user interface preferences, 1-2
Shared IP mode, 7-9, 14-4, 17-11
Shut Down Guest, 16-26
Shutdown a logical domain, 15-24
Shutdown Zone, 14-25
Snapshot
 volume, 18-31
Snapshots
 assembly templates, 18-32
 boot environments, 12-43, 12-45
 volumes, 18-32, 18-33
Software libraries, 5-1
 creating, 5-5
 EC Library, 5-3
 file system, 5-4
 Firmware images, 5-1
 icons, E-10
 Linux and Oracle Solaris 8-10 Software Update Library, 5-8
 local, 5-4
 NAS, 5-4
 OS images, 5-1
 roles, 5-1
 server pools, 17-5
 specifying EC Library, 5-3
 uploading, 5-18
 viewing, 5-4, 5-19
Solaris See Oracle Solaris, 13-1
SPARC SuperCluster
 Dashboard, 20-21
 Incidents tab, 20-24
 Networks tab, 20-24
 Rack View, 20-25
 System View, 20-24
SPARC SuperCluster Rack View, 20-26
SPARC SuperCluster System Report
 Baseline Check, 20-30
 Compute Nodes, 20-29
 Exadata Storage Server, 20-30
 power distribution units, 20-30
 Storage Appliances, 20-30
 Switches, 20-30
 validation, 20-30
SPARC SuperCluster System View
 actions, 20-25
SPARC SuperCluster T4-4, 20-2
 See Oracle Engineered Systems, 20-19
SPARC SuperCluster View, 20-20, 20-21, 20-23, 20-24
 Membership Graph, 20-22
 Status, 20-22
 Summary, 20-22
Sparse root zones, 14-2
 requirements, 14-14
ssh, 16-6
Start Logical Domain, 15-24
Start page preferences
 user preferences, 1-2
Start Selected Guest, 16-26
Static block storage, 6-10
Static IP, 7-16
Static private networks, 7-4, 18-7

- Static route, 7-15
- Status
 - jobs, 3-1
- Storage
 - Exadata Storage Server, 20-27
 - monitoring policies, 4-12
 - non-global zones, 14-29
 - SPARC SuperCluster, 20-27
 - storage libraries, 6-1
 - vDC, 18-10, 18-22, 18-30, 18-31, 18-32
 - virtual datacenter, 18-4, 18-5
- Storage array
 - iSCSI SAN, 6-18
- Storage Connect plug-ins
 - Oracle VM Manager, 6-15
- Storage libraries
 - block storage, 6-9
 - dynamic, 6-9
 - static, 6-10
 - file system, 6-3, 6-8, 6-9
 - global zones, 14-4
 - high availability, 6-20
 - icons, E-10
 - local, 6-3, 6-8
 - creating, 6-8
 - deleting, 6-9
 - editing, 6-8
 - viewing, 6-8
 - logical domains, 15-25
 - metadata, 6-3
 - NAS, 6-4, 6-17, 17-24
 - non-global zones, 14-27
 - Oracle Solaris Zones, 17-24
 - Oracle VM Server, 16-12, 16-13
 - Oracle VM Server for SPARC, 15-14, 17-24
 - Oracle VM Server for x86, 17-24
 - repositories, 6-6
 - roles, 6-1
 - SAN, 6-3, 17-24
 - server pools, 6-5, 17-5, 17-24
 - virtual machines, 16-23
- Storage role, 16-13
- Storage servers, 16-12
 - Exadata, 6-20
- Suggested Action, 9-12
- Sun ZFS Storage Appliances, 2-21
 - account, 2-23
 - discovery, 2-22
 - discovery profile, 2-26, 2-27
 - LUNs, 2-23, 2-25
 - ocuser, 2-23
 - setting up, 2-22
 - SPARC SuperCluster, 20-19
- SUNWjet, 12-31
- SuperCluster, 20-2
 - icons, E-17
 - See Oracle Engineered Systems, 20-19
- SuperCluster System Admin, 20-20
 - prerequisites, 20-20
- SuperCluster System Report, 20-29

- Report Parameters, 20-29
- Summary Table, 20-29
- Suspend Selected Guest, 16-26
- Switch Management Access, 12-10
- Switches, 7-18, 20-30
 - icons, E-18
 - monitoring policies, 4-12
 - Oracle Engineered Systems, 20-28
 - SPARC SuperCluster, 20-19
- synchronize, 12-55
- System Catalog, 13-6
 - creating profiles, 13-16
 - updating operating systems, 13-14, 13-15
 - viewing, 13-6
- System Catalog Report, 10-23
- System Information Report, 10-27
- System-defined monitoring policies, 4-9, 4-11
- System-defined rules, 4-3

T

- Table refresh frequency
 - user interface preferences, 1-2
- Tags
 - asset attributes, 2-18, 2-29
 - service, 2-6
- Take Action on Incident, 9-8, 9-10
- Take Ownership, 16-11
- Task Execution Order, 3-2
- Tasks, 3-6
- Templates
 - assembly, 18-32
 - deleting report template, 10-3
 - edit report template, 10-3
 - JumpStart Enterprise Toolkit, 12-32
 - Reports, 10-2
 - virtual machines, 16-21
 - vServer, 18-22, 18-28, 18-30
- Thresholds, 4-3, 4-4, 4-6
- Time interval
 - user interface preferences, 1-2
- Time zone attributes, A-10

U

- Unassigned incidents, 9-6
- Unbind Network, 14-11, 15-16
- Unbreakable Linux Network, 16-16
- Unconfigured assets, 2-5
- Universal icons, E-19
- Unmanaged fabrics, 7-4, 18-6
- Update, 13-1
- Update job
 - Microsoft Windows, 13-20
- Update Oracle VM Server Host, 16-16
- Update policies, 13-7
 - custom, 13-8
- Update Policy
 - icons, E-20
- Update Profile, 12-53

- Update Profiles, 13-16
 - icons, E-20
- Update profiles, 13-9, 13-10
- Update Virtual Datacenter, 18-10
- Updates See operating system updates, 13-1
- updating, 2-16
- Upload, 5-18
 - results, 5-19
- Upload images, 5-11
- User Accounts
 - See *Oracle Enterprise Manager Ops Center Administration Guide*, 20-8
- User interface preferences
 - connectivity check interval, 1-2
 - console timeout, 1-2
 - job status popup duration, 1-2
 - session timeout, 1-2
 - table refresh frequency, 1-2
 - time interval, 1-2
- User preferences
 - membership graph preferences, 1-2
 - user preferences summary, 1-3
- User preferences summary
 - user preferences, 1-3
- User-defined monitoring policies, 4-9, 4-14
- User-defined rules, 4-3
 - Boolean Control, 4-4
 - Enumerated Control, 4-4
 - Expression, 4-4
 - parameters, 4-4
 - rule types, 4-4
 - thresholds, 4-4

V

vDC

- accounts, 18-3, 18-14
 - adding, 18-12
 - deleting, 18-16
 - quotas, 18-21
 - removing, 18-13
 - updating, 18-15
- API, 18-19
- CLI, 18-19
- cloud user, 18-18
 - creating, 18-8
 - deleting, 18-12
 - managing, 18-10
- network, 18-6
- network domain, 18-7
- Networks, 18-10
- networks, 7-9, 18-4, 18-5, 18-22, 18-23
- operating systems, 18-23
- requirements, 18-7
- roles, 18-3
- server pools, 18-4, 18-10, 18-11
- storage, 18-4, 18-5, 18-10, 18-22
- storage volumes, 18-30
- updating, 18-10
- user resources, 18-22

- volumes, 18-5
 - creating, 18-31
 - deleting, 18-32
 - importing, 18-31
 - snapshot, 18-31
 - updating, 18-32
- vServer type, 18-16
 - creating, 18-17
 - deleting, 18-18
 - editing, 18-18
 - system-defined, 18-16
- vServers, 18-22
 - cloud user, 18-18
 - creating, 18-23, 18-24
 - managing, 18-25
 - network, 18-23
 - server templates, 18-22, 18-28, 18-30
- vDC Cloud
 - roles, 18-19
- View Alerts, 9-8, 9-9
- View Annotations, 9-8, 9-9
- View Associated Assets, 4-14
- View Comments, 9-8, 9-10
- View Interactive, 10-1
- View Possible Impacts and Causes, 9-8, 9-9
- View Suggested Actions, 9-8, 9-10
- Viewing annotations, 9-13, 9-18
- Viewing comments, 9-18
- Viewing incident details, 9-15
- Viewing unresolved incidents, 9-15
- Virtual datacenter See vDC, 18-2
- Virtual disk, 6-3
- Virtual hosts
 - logical domain, 17-1
 - non-global zones, 17-1
 - virtual machines, 17-1
- Virtual image, 5-9
- Virtual machines, 15-2, 18-7, 18-22
 - creating, 16-21, 16-22
 - deleting, 16-26
 - deployment plans, 16-24
 - monitoring policies, 4-12
 - network, 16-13
 - network domain, 16-23
 - operating systems, 16-25
 - Oracle VM Server for x86, 16-20
 - profiles, 16-21
 - reboot, 16-26
 - server pools, 17-24, 17-26
 - shutdown, 16-26
 - starting, 16-26
 - storage libraries, 16-23
 - suspending, 16-26
 - type, 18-16
- Virtual resources, 18-22
- Virtual servers, 18-3
- Virtual switches, 7-9, 15-14, 15-15, 17-6
- Virtualization
 - icons, E-14
- Virtualization hosts

- adding to server pool, 17-20, 17-21
- control domain, 17-1
- global zone, 17-1
- Oracle Solaris Zones, 17-1
- Oracle VM Server for SPARC, 17-1
- Oracle VM Server for x86, 17-1
- Virtualization stack, F-7
- Virtualization types, 16-21, 16-22
 - HVM, 16-21
 - HVM with PV, 16-21
 - PVHVM, 16-21
 - PVM, 16-21
- VLAN, 7-17
- VM Type
 - See Virtualization types, 16-22
- Volumes
 - snapshots, 18-32
 - vDC, 18-5, 18-30, 18-31, 18-32
- vServer type, 18-16, 18-17, 18-18, 18-23, 18-24
- vServers, 18-3, 18-12, 18-18
 - create server templates, 18-28
 - delete server templates, 18-30
 - register server templates, 18-30
 - update server templates, 18-30
 - upload server templates, 18-28
 - vDC, 18-22, 18-23, 18-24, 18-25, 18-28

W

- Wake on LAN, 16-18
- Wake-on-LAN, 16-15
- WAN boot, 12-25
 - disable and enable, 12-27
 - requirements, 12-26
 - setup, 12-27
- Warning, 4-3, 9-14
- Warranty, 9-21, 9-22
- Web Interface, 20-6
- White list, 10-10, 10-13
- Whole root zones, 14-2
 - requirements, 14-14
- Windows
 - discovering, 2-20
 - updating, 13-19
 - Windows Firewall, 2-20
 - WMI, 2-20
- Windows Firewall, 2-20
- Windows Management Instrumentation
 - See WMI, 2-20
- Windows Update, 13-19
- WMI, 2-20, 13-19

Y

- YUM Repository, 16-16

Z

- zone path, 14-18
- Zpools, 12-37, 12-38, 12-42, 14-21, 14-27