

Oracle® Enterprise Manager

Cloud Control Host Lifecycle Management

Guide



13c Release 5
F37168-02
December 2021

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

F37168-02

Copyright © 2015, 2021, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

Contributors: Enterprise Manager Cloud Control Lifecycle Management Development Teams, Quality Assurance Teams, Customer Support Teams, and Product Management Teams.

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, then the following notice is applicable:

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	x
Documentation Accessibility	x
Related Resources	x
Conventions	x

1 Discovering Hosts and Software Deployments

Discovering Hosts Automatically and Adding Targets Manually	1-1
Discovering Hosts Manually and Adding Targets Manually	1-1

2 Provisioning Bare Metal Servers

Getting Started with Provisioning Bare Metal Servers	2-1
Overview Of Bare Metal Provisioning	2-2
Accessing Bare Metal Provisioning Page	2-3
Provisioning Environment for Bare Metals	2-4
Software Library and its Entities	2-4
Boot Server	2-4
Stage Server	2-4
Reference Host	2-4
RPM Repository	2-5
Bare Metal Provisioning Flow	2-5
Supported Releases of Linux	2-5
Setting Up Infrastructure for Bare Metal Provisioning	2-5
Setting Up Stage Server	2-6
Prerequisites to Setup a Stage Server	2-6
Setting up a Stage Server and Accessing the Management Agent files	2-6
Setting Up Boot Server and DHCP Server	2-9
Setting Up RPM Repository	2-10
Setting Up Oracle Linux RPM Repository	2-11
Exposing RPM Repository through HTTP or FTP	2-11
Configuring Stage Server	2-11

Configuring Boot Server	2-12
Configuring DHCP Server	2-12
Configuring RPM Repository	2-12
Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host	2-13
Configuring Software Library Components	2-13
Creating Operating System Component	2-13
Creating Disk Layout Component	2-15
Creating an Oracle Virtual Server Component	2-17
Prerequisites for Provisioning Operating Systems and Oracle VM Servers	2-18
Provisioning Operating Systems	2-18
Provisioning Oracle VM Servers	2-22
Viewing Saved Plans	2-25
Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers	2-26

3 Patching Linux Hosts

Overview of Patching Linux Hosts	3-1
About the Deployment Procedure for Patching Linux Hosts	3-2
Supported Linux Releases	3-2
Setting Up Infrastructure for Linux Patching	3-2
Prerequisites for Using the Linux Patching Feature	3-2
Setting Up the RPM Repository for Linux Patching	3-3
Prerequisites for Setting Up the RPM Repository	3-3
Setting Up the RPM Repository for Patching	3-4
Setting Up Linux Patching Group for Compliance Reporting	3-6
Prerequisites for Setting Up Linux Patching Group	3-6
Setting Up a Linux Patching Group	3-6
Patching Linux Hosts	3-8
Applying Patches on a Linux Patching Group Based on Compliance	3-8
Applying Ad Hoc or Emergency Patches on Linux Hosts	3-10
Managing Linux Configuration Files	3-12
Overview of Linux Configuration Files	3-13
Prerequisites for Managing Configuration Files	3-13
Creating a Linux Configuration File Channel	3-13
Uploading Linux Configuration Files to a Particular Channel	3-13
Prerequisites for Uploading Linux Configuration Files	3-13
Uploading Linux Configuration Files	3-13
Importing Linux Configuration Files from One Channel to Another	3-14
Prerequisites for Importing Linux Configuration Files	3-14
Importing Linux Configuration Files	3-14
Deploying Linux Configuration Files From a Particular Channel	3-14

Prerequisites for Deploying Linux Configuration Files	3-15
Deploying Linux Configuration Files	3-15
Deleting a Linux Configuration File Channel	3-15
Prerequisites for Deleting a Linux Configuration File Channel	3-15
Deleting Linux Configuration File Channels	3-15
Additional Linux Patching Tasks You Can Perform	3-16
Viewing Linux Patching Compliance History	3-16
Prerequisites for Viewing Linux Patching Compliance History	3-16
Viewing Linux Patching Compliance History	3-16
Patching Non-Compliant Linux Packages	3-17
Prerequisites for Patching Non-Compliant Linux Packages	3-17
Patching Non-Compliant Linux Packages	3-17
Rolling Back Linux Patch Update Sessions or Deinstalling Packages	3-17
Prerequisites for Rolling Back Linux Patch Update Sessions or Deinstalling Packages	3-18
Rolling Back Linux Patch Update Sessions or Deinstalling Packages	3-18
Registering a Custom Package Channel	3-18
Prerequisites for Registering a Custom Package Channel	3-18
Registering a Custom Package Channel	3-19
Cloning a Package Channel	3-19
Prerequisites for Cloning a Package Channel	3-19
Cloning a Package Channel	3-19
Copying Packages from One Channel to Another	3-20
Prerequisites for Copying Packages from One Channel to Another	3-20
Copying Packages from One Channel to Another	3-20
Adding Custom Packages to a Channel	3-21
Prerequisites for Adding Custom Packages to a Channel	3-21
Adding Custom Packages to a Channel	3-21
Deleting a Package Channel	3-22
Prerequisites for Deleting a Package Channel	3-22
Deleting a Package Channel	3-22

4 Monitoring and Managing Hosts

Host Consolidation Planner	4-1
Overview of Host Consolidation Planner	4-1
Creating a Host Consolidation Planner Project	4-2
Selecting the Host Consolidation Project Type	4-3
Specifying Host Source Candidates	4-3
Specifying Host Destination Candidates	4-3
Setting Data Collection Parameters	4-4
Choosing a Pre-configured Scenario	4-4

Reviewing and Saving the Host Consolidation Project	4-6
Creating a Host Consolidation Planner Scenario	4-6
Setting Up the Scenario	4-6
Defining Constraints for Host Consolidation	4-7
Planning the Destination for a Physical Server to Virtual Server Project	4-8
Planning the Destination for a Physical Server to Physical Server Project	4-8
Mapping Host Sources to Destinations	4-9
Reviewing and Saving the Host Consolidation Scenario	4-9
Other Host Consolidation Scenario Creation Options	4-10
Evaluating Host Consolidation Scenarios	4-10
Overview of Host Management	4-12
Host Statistics	4-12
Diagnosing Host Problems	4-12
Viewing Targets on the Host	4-13
Storage Statistics and History	4-13
Setting Up the Environment to Monitor Hosts	4-13
Required Installations	4-14
For Linux Hosts - Installing YAST	4-14
Setting Up Credentials	4-15
Setup Needed for Host Monitoring	4-15
Viewing Monitoring Configuration	4-16
Setting Up Monitoring Credentials	4-16
Target Setup Needed for Host Administration	4-16
Monitoring Hosts	4-17
Overall Monitoring	4-17
CPU Details	4-17
Memory Details	4-17
Disk Details	4-18
Program Resource Utilization	4-18
Log File Alerts	4-18
Metric Collection Errors	4-18
Storage Details	4-18
Storage Utilization	4-19
Overall Utilization	4-19
Provisioning Summary	4-19
Consumption Summary	4-20
ASM	4-20
Databases	4-20
Disks	4-20
File Systems	4-21
Volumes	4-22

Vendor Distribution	4-23
Storage History	4-23
Storage Layers	4-23
Storage Refresh	4-24
Ksplice for Oracle Linux	4-24
Customizing Your Host Monitoring Environment	4-24
Customizing the Host Home Page	4-25
Using Groups	4-25
Administering Hosts	4-26
Configuration Operations on Hosts	4-27
Configuring File and Directory Monitoring Criteria	4-27
Configuring Generic Log File Monitor Criteria	4-28
Configuring Program Resource Utilization Monitoring Criteria	4-29
Administration Tasks	4-30
Services	4-31
Default System Run Level	4-31
Network Card	4-32
Host Lookup Table	4-33
NFS Client	4-33
User and Group Administration (Users)	4-35
User and Group Administration (Groups)	4-36
Using Tools and Commands	4-37
Enabling Sudo and Power Broker	4-37
Executing the Host Command Using Sudo or PowerBroker	4-37
Using Remote File Editor	4-38
Adding Host Targets	4-39
Running Host Command	4-39
Accessing Host Command	4-39
Executing Host Command Using Sudo or Power Broker	4-39
Execute Host Command - Multiple Hosts	4-40
Execute Host Command - Group	4-41
Execute Host Command - Single Host	4-42
Load OS Script	4-43
Load From Job Library	4-43
Execution History	4-43
Execution Results	4-43
Miscellaneous Tasks	4-44
Enabling Collection of WBEM Fetchlet Based Metrics	4-44
Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts	4-44
Adding and Editing Host Configuration	4-45
Managing Oracle Linux Homes	4-46

Oracle Linux Home	4-46
Target Navigation Tree (TNT) of "Oracle Linux Home"	4-47
Oracle Linux Home Target	4-48
Oracle Linux Home page	4-48
General Region	4-48
Overview of Incidents and Problems	4-49
Host Flux	4-49
CPU	4-50
Memory	4-51
Oracle Linux Patching	4-52
Ksplice for Oracle Linux	4-55
Ksplice Metrics	4-55
Ksplice Patching	4-56
Ksplice Linux Hosts Page	4-57
Additional Setup for Real-time Monitoring	4-59
Overview of Real-Time Monitoring	4-59
Overview of Resource Consumption Considerations	4-59
OS File Monitoring Archiving	4-60
OS File Read Monitoring	4-60
Creating Facets That Have Very Broad Coverage	4-60
Cloud Control Repository Sizing	4-60
Configuring Monitoring Credentials	4-61
Preparing To Monitor Linux Hosts	4-62
OS File Monitoring	4-62
Debugging Kernel Module Or Other File Monitoring Issues	4-63
Preparing To Monitor Windows Hosts	4-65
Verifying Auditing Is Configured Properly	4-66
Subinacl External Requirements	4-66
Preparing To Monitor Solaris Hosts	4-67
Enabling BSM Auditing	4-67
Managing Audit Log Files	4-68
Preparing to Monitor AIX Hosts	4-69
Installation Prerequisite for AIX 5.3	4-69
Administering AIX Auditing	4-69
Verifying AIX System Log Files for the OS User Monitoring Module	4-70
Preparing To Monitor the Oracle Database	4-70
Setting Auditing User Privileges	4-71
Specifying Audit Options	4-71
Oracle Database Table Monitoring	4-72
Setting Up Change Request Management Integration	4-73
BMC Remedy Action Request System 7.1 Integration	4-73

Overview of the Repository Views Related to Real-time Monitoring Features	4-79
Modifying Data Retention Periods	4-84
Real-time Monitoring Supported Platforms	4-85
OS User Monitoring	4-85
OS Process Monitoring	4-87
OS File Monitoring	4-88
OS Windows Registry Monitoring	4-91
OS Windows Active Directory User Monitoring	4-91
OS Windows Active Directory Computer Monitoring	4-92
OS Windows Active Directory Group Monitoring	4-92

A Understanding PXE Booting and Kickstart Technology

About PXE Booting and Kickstart Technology	A-1
Subnet Provisioning Usecases	A-2

B Troubleshooting Issues

Troubleshooting Linux Provisioning Issues	B-1
Troubleshooting Linux Patching Issues	B-3
Frequently Asked Questions on Linux Provisioning	B-4

Index

Preface

The Cloud Control Host Lifecycle Management Guide introduces you to the suite of operating system and host lifecycle management solutions offered by Enterprise Manager Cloud Control, and describes in detail how you can use the discovery, provisioning, patching, and monitoring features to manage your data center.

Audience

This guide is primarily meant for administrators who want to use the discovery, provisioning, patching, and monitoring features offered by Cloud Control to meet their operating system and host lifecycle management challenges. As an administrator, you can be either a *Designer*, who performs the role of a system administrator and does critical operating system and host operations, or an *Operator*, who runs the default as well custom deployment procedures, patch plans, and patch templates to manage operating system and host configurations.

Documentation Accessibility

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

Access to Oracle Support

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

Related Resources

For more information, see the documents available in the Oracle Enterprise Manager documentation library: [Enterprise Manager Documentation](#).

Conventions

The following conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.

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

1

Discovering Hosts and Software Deployments

Discovery is the first step toward monitoring and managing the health of your software deployments. Discovery refers to the process of identifying unmanaged hosts and their software deployments, and adding them as manageable targets in Oracle Enterprise Manager Cloud Control (Cloud Control).

This chapter describes how you can discover the hosts and their software deployments, and add them to Cloud Control. In particular, this chapter describes the following:

- [Discovering Hosts Automatically and Adding Targets Manually](#)
- [Discovering Hosts Manually and Adding Targets Manually](#)

Discovering Hosts Automatically and Adding Targets Manually

Automatic discovery refers to the process of scanning hosts for Oracle software that can be managed and monitored by Cloud Control. By default, the automatic discovery runs every 24 hours to discover targets.

In automatic discovery, you enable a Management Agent running on the host to run an Enterprise Manager job that scans for unmanaged hosts. You then promote these unmanaged hosts to managed hosts by deploying Management Agents on these hosts, then you search for targets on these managed hosts, and finally you promote these targets to managed target status.

You can configure automatic discovery to set up a schedule for discovery, the target types to be discovered, and the hosts to scan for targets. You can then promote the discovered hosts to managed targets in Cloud Control. You can also regularly identify targets that have been newly added to the infrastructure, and add them to Cloud Control for monitoring.

Once automatic discovery has been configured, you can check the Auto Discovery Results page on a regular basis to see what targets have been discovered.

For information on automatically discovering and monitoring targets, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Discovering Hosts Manually and Adding Targets Manually

In addition to automatic discovery, Cloud Control enables you to manually add hosts as well as a wide variety of Oracle software and components as managed targets. When you add a target manually, you do not need to go through the process of discovery by adding the target directly. Discovering targets in this way eliminates the need to consume resources on the Oracle Management Agent to perform discovery when it is not needed.

For information on manually discovering and monitoring targets, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

2

Provisioning Bare Metal Servers

This chapter explains how you can provision Linux on bare metal servers using Oracle Enterprise Manager Cloud Control. In particular, this chapter covers the following:

- [Getting Started with Provisioning Bare Metal Servers](#)
- [Overview Of Bare Metal Provisioning](#)
- [Supported Releases of Linux](#)
- [Setting Up Infrastructure for Bare Metal Provisioning](#)
- [Provisioning Operating Systems](#)
- [Provisioning Oracle VM Servers](#)
- [Viewing Saved Plans](#)
- [Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers](#)

Tip:

Before you begin provisioning of Linux on bare metal boxes, it is advisable to set preferred credentials for the Stage Server. For more information, see [Setting Up Credentials](#) to set up preferred credentials. If you want to use a reference host, set credentials for the reference host also. You can also set preferred credentials when configuring the deployment procedure for provisioning Linux.

Note:

Before starting the provisioning Linux operations, ensure that you configure sudo privileges. For more information about configuring sudo privileges, see [Setting Up Credentials](#).

Getting Started with Provisioning Bare Metal Servers

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Linux operating system. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision Linux operating system. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 2-1 Getting Started with Provisioning Linux Operating System

Step	Description	Reference Links
Step 1	Knowing About The Supported Releases Know what releases of Linux are supported for provisioning.	<ul style="list-style-type: none">To learn about the releases supported for Linux Provisioning, see Supported Releases of Linux.
Step 2	Knowing the Use Case This chapter covers provisioning Linux. Understand the use case for Linux provisioning.	<ul style="list-style-type: none">To learn about provisioning bare metal boxes, see Provisioning Operating Systems.
Step 3	Setting Up Infrastructure Before you perform Linux provisioning, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, and setting up of Oracle Software Library.	<ul style="list-style-type: none">To learn about the prerequisites to be met for provisioning bare metal boxes, see Setting Up Infrastructure for Bare Metal Provisioning.
Step 4	Provisioning Linux Provision Linux on bare metal boxes.	<ul style="list-style-type: none">To provision Linux on bare metal boxes, follow the steps explained in Provisioning Operating Systems.

Overview Of Bare Metal Provisioning

Proliferation of low cost servers in our data centers has brought in a fresh set of management challenges. The well-acknowledged problems include the difficulty in managing consistency and compatibility across operating system and software deployments, server drifts and security vulnerabilities that lead to lack of compliance, difficulty in deploying software, difficulty in provisioning new servers with variety of configurations and applications, high cost of operation and difficulty in adapting to changes in workload of the environment. These lead to system administrators and DBAs spending significant amount of their time in software and server provisioning operations.

Oracle's answer to software and server management challenges is its Bare Metal Provisioning Application, an application built into Enterprise Manager Cloud Control. The application addresses all data center and server farm challenges by provisioning software and servers quickly and efficiently. The application uses standardized PXE (Pre Boot Execution environment) booting process for provisioning both bare-metal and live servers. It provides a role based User Interface, for easily creating gold images and initiating automated, unattended installs.

This section covers the following:

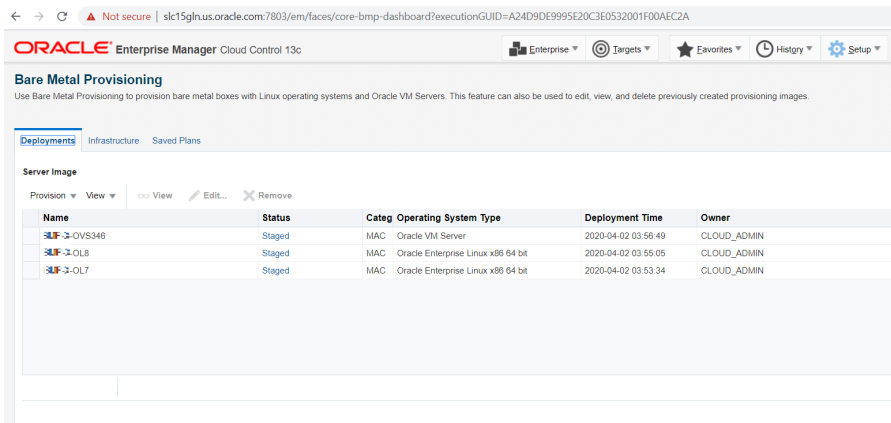
- [Accessing Bare Metal Provisioning Page](#)
- [Provisioning Environment for Bare Metals](#)
- [Bare Metal Provisioning Flow](#)

Accessing Bare Metal Provisioning Page

To access the Bare Metal Provisioning page, from **Enterprise** menu, select **Provisioning and Patching**, then click **Bare Metal Provisioning**. On the Bare Metal Provisioning home page, the following tabs are displayed:


Deployments tab allows you to provision Linux operating system or Oracle VM Server on bare metal boxes. All the servers that are provisioned are displayed on this page in the Server Image section.

Figure 2-1 Image showing Deployment tab for Bare Metal provisioning



Infrastructure allows you to setup the infrastructure required to provision bare metal machines. For information about the Stage Servers, Boot Servers, DHCP Servers, and RPM Repositories, see [Provisioning Environment for Bare Metals](#). For details on setting up and configuring each of these servers, see [Setting Up Infrastructure for Bare Metal Provisioning](#).

Figure 2-2 Infrastructure tab for Bare Metal Provisioning



Saved Plans tab allows you to *view* all the deployment procedures that were saved as a templates with all the essential attribute values for future runs. However, note that these plans can only be viewed from this tab, to run these saved plans see [Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers](#).

Provisioning Environment for Bare Metals

The deployment environment in the data center needs to be setup in a certain manner in order to support the provisioning application. Besides the Oracle Management Server (OMS) which hosts Cloud Control and Provisioning Application, the following need to be setup and configured before using the provisioning application.

Software Library and its Entities

For information about configuring Software Library and its entities, see [Setting Up Oracle Software Library](#).

Boot Server

One of the key requirements of application is the ability of the hardware server to boot up over the network (rather than from a local boot device). A boot server must be set up so that it is able to service the requests from the designated hardware servers in order for them to boot over the network. Boot server must be an Cloud Control target and should be able to receive the BOOTP and TFTP (Trivial File Transfer Protocol) requests over the network from the hardware server. See [Setting Up Boot Server and DHCP Server](#) for setting up a boot server with DHCP/TFTP combination. Also, see [Configuring Boot Server](#). It is also recommended that the users read about DHCP, PXE, and Redhat Kickstart technology before going through the boot server setup. See [Understanding PXE Booting and Kickstart Technology](#) for a detailed discussion on PXE.

Stage Server

During provisioning of an image on hardware servers, the required binaries and files are first transferred to a stage server. This is known as **Staging** phase and is responsible for preparing images to be installed over the network, and exposing installable or executable software elements over the network to the target hardware server being provisioned.

The Provisioning application requires at least one stage server on which all the activities related to staging can be performed. Stage server should again be an Cloud Control target. Refer to section [Setting Up Stage Server](#) for setting up a stage server. Also, see [Configuring Stage Server](#).

Reference Host

A Reference Host (also called a **gold machine**) is the machine that the Provisioning application uses as a reference to create the Linux operating system component. The Provisioning application picks up the list of RPMs (along with their versions) installed on the reference host, and fetches those RPMs from a RPM repository to create an Linux OS component that represents the operating system installed on the reference host. The reference host must be an Cloud Control target.

RPM Repository

The Provisioning application picks up the RPMs for the operating system from the RPM repository. At least one repository needs to be setup for use by the Provisioning application. From the networking perspective, you are advised to keep the RPM Repository as close to the target machines as possible. It will help in bringing down the installation time drastically by reducing the time taken to transfer RPMs from the RPM Repository to the hardware servers. If you have multiple hardware server groups residing at physically different locations, it would be better to have one RPM Repository for each of these locations. Refer to section [Setting Up RPM Repository](#) for setting up a RPM repository. Also refer to section [Configuring RPM Repository](#).

Bare Metal Provisioning Flow

The provisioning process consists of the following two high-level tasks:

1. Setting Up Provisioning Environment ([Setting Up Infrastructure for Bare Metal Provisioning](#)):
 - Setting up and configuring Boot/DHCP server and Stage server, setting up RPM repository and Software Library
 - Optionally, creating baremetal provisioning entities
2. Provisioning Linux using Bare Metal Provisioning Application ([Provisioning Operating Systems](#)):
 - Launching the Baremetal Provisioning wizard to configure the bare metal machines using MAC addresses, subnet, or re-imaging Cloud Control hosts.
 - Powering up the bare metal machine on the network to begin the PXE-based OS boot and install process. For information about PXE Booting and KickStart, see [Understanding PXE Booting and Kickstart Technology](#).

Supported Releases of Linux

Cloud Control supports bare metal provisioning of 32-bit and 64-bit variants of the following operating systems:

- Oracle Linux 6
- Oracle Linux 7
- Oracle VM Server 3.4
- Oracle Linux 8

Setting Up Infrastructure for Bare Metal Provisioning

This section describes how to set up the infrastructure required to provision bare metal machine. In particular, this section describes the following:

- [Setting Up Stage Server](#)
- [Setting Up Boot Server and DHCP Server](#)
- [Setting Up RPM Repository](#)
- [Configuring Stage Server](#)

- [Configuring Boot Server](#)
- [Configuring DHCP Server](#)
- [Configuring RPM Repository](#)
- [Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host](#)
- [Configuring Software Library Components](#)

Setting Up Stage Server

This section contains:

- [Prerequisites to Setup a Stage Server](#)
- [Setting up a Stage Server and Accessing the Management Agent files](#)

Prerequisites to Setup a Stage Server

Ensure that you meet the following prerequisites before setting up the stage server:

- The user or role used to create the top-level directory (stage directory) where you stage the Agent rpms should have Sudo access to *root*. To ensure that you have sudo access on the stage storage, log in to the Cloud Control console and set the sudo privileges.

 **Note:**

Oracle recommends that the stage server must have very limited access due to the criticality and sensitivity of the data it hosts. The super administrator can enforce this by creating one account on the stage server, and setting it as the preferred credential, to be used by all the provisioning users in Cloud Control. This preferred credential should also be a valid ORACLE_HOME credential (belonging to ORACLE_HOME owner's group).

- The user creating the top-level directory must have write permissions on it. To ensure that you have write access on the stage server, log in to the Cloud Control console and set the privileged preferred credentials for the stage server host.
- The minimum space requirement for the stage directory is 100 MB.

Setting up a Stage Server and Accessing the Management Agent files

To set up a stage server, and access the Management Agent RPM files, follow these steps:

1. Set up an NFS Stage Server or a HTTP Server.
 - To setup an NFS Stage Server, see [Setting up an NFS Stage Server](#).
 - To set up a HTTP Server, see [Setting up a HTTP Stage Server](#).
2. Log in to the stage server running on the Management Agent, and create a top-level directory to store all the Management Agent installation files. In this section, the variable `STAGE_TOP_LEVEL_DIRECTORY` is used to refer to the top level directory on the stage server.

For example:

```
User: aime
Stage Server: upsgc.example.com
Stage Directory: /scratch/stage
```

 **Note:**

In this case, the `aime` user should have `sudo` access to `root`, and should have write permissions on `/scratch/stage` directory

3. To create and copy the Management Agent Files to Stage location, run the following commands on the OMS:

```
For using the NFS Stage Server:
STAGE_TOP_LEVEL_DIRECTORY=/scratch/stage
```

```
For using the HTTP based Stage Server:
STAGE_TOP_LEVEL_DIRECTORY=/var/www/html/stage
```

```
emcli get_agentimage_rpm -destination="${STAGE_TOP_LEVEL_DIRECTORY:?}" -
platform="Linux x86-64"
```

```
[root@upgps1t12 stage]# pwd
/scratch/stage
[root@upgps1t12 stage]# ls
1 10 21 6 9 oracle-agt-13.4.0.0-1.0.x86_64.rpm
```

 **Note:**

- If NFS is used then the staging process will automatically discover the agent rpm and there's no requirement for you to provide a URL for the rpm.
- If HTTP is used then a URL will be required to reference the Agent rpm. The Agent URL is `http://host.example.com/agent_dir/oracle-agt-13.4.0.0-1.0.x86_64.rpm`. For more information on setting up HTTP Stage Server, see [Setting up a HTTP Stage Server](#).

Setting up an NFS Stage Server

During the installation, hardware servers mount the stage directory so that all the files required for installation appear as local files. In such a scenario, the stage server functions as the NFS server, and the hardware servers as its clients. IF the stage server is an NFS server then any files that it NFS exports must be available to its clients; for files on NAS storage it might be necessary to configure the NAS to allow this to happen.

Make sure that you perform the following steps on the stage server:

1. Run the following command to install an NFS service:

```
rpm --quiet -q nfs || yum -y install nfs
```

2. Run the following commands to configure NFS to export the stage server's top level directory (STAGE_TOP_LEVEL_DIRECTORY):

```
STAGE_TOP_LEVEL_DIRECTORY=/scratch/stage  
echo "${STAGE_TOP_LEVEL_DIRECTORY}*(ro, sync)" >>/etc/exports
```

3. To reflect these changes on the NFS daemons, run the following command:

- Oracle Linux 6:

```
service nfs restart
```

- Oracle Linux 7 and up:

```
systemctl start nfs
```

4. Ensure NFS starts up on reboot, and is working now:

- Oracle Linux 6:

```
chkconfig nfs on
```

- Oracle Linux 7 and up:

```
systemctl enable nfs
```

5. To install a Management Agent see: Oracle Enterprise Manager Cloud Control Basic Installation Guide .

Setting up a HTTP Stage Server

To setup a HTTP Stage Server, follow these steps:

1. Run the following commands to install a stage server and start it:

- For Oracle Linux 6:

```
rpm --quiet -q httpd || yum -y install httpd  
service httpd restart  
chkconfig httpd on
```

- For Oracle Linux 7 and above:

```
rpm --quiet -q httpd || yum -y install httpd  
systemctl start httpd  
systemctl enable httpd
```

2. Create a HTTP stage directory as follows:

```
mkdir /var/www/html/stage
```

3. The URL to access the HTTP stage server is:

```
http://host.example.com/stage
```

4. To install a Management Agent see: Oracle Enterprise Manager Cloud Control Basic Installation Guide .

 **Note:**

`/var/www/html/stage` is the stage directory, and `http://host.example.com/stage` is the base URL.

Setting Up Boot Server and DHCP Server

 **Note:**

Ensure that you have 2 GB RAM available for boot server, stage server, and RPM repository server.

If you have the required boot server, stage server, and RPM repository already created, then set up the preferred credentials.

Complete the following steps to setup a machine as the boot server:

1. Install DHCP and TFTP Servers if not already installed.

The two servers could be running either on the same machine, or on different machines. Oracle recommends running the TFTP server on the same host machine as the DHCP server. In case the two servers are installed and configured on different machines, the machine running the TFTP server will be referred to as the boot server.

2. Configure the TFTP server:

Ensure that the *pxelinux* boot loader (**pxelinux.0**) exists in the directory that is configured for your TFTP server (`/var/lib/tftpboot` in the given examples).

3. Configure DHCP Server:

Edit the `/etc/dhcpd/dhcpd.conf` file. A sample **dhcpd.conf** file for PXE setup is shown below:

```
allow booting;
allow bootp;

option domain-name <domain_name>;
option domain-name-servers dns_servers;
option routers <default_router>;

subnet <subnet-number> netmask <netmask> {
    [ parameters ]
    [ declarations ]
}
# Group the PXE bootable hosts together
```

```
group {  
  
# PXE-specific configuration directives...  
  
next-server <TFTP_server_IP_address>;  
  
filename "linux-install/pxelinux.0";  
  
host <hostname> {  
hardware ethernet <MAC address>;  
fixed-address <IP address>;  
}  
}
```

The *next-server* option in the DHCP configuration file specifies the host name or IP Address of the machine hosting the TFTP server. Oracle recommends running the TFTP Server on the same host machine as the DHCP Server. Therefore, this address should be the IP Address or host name for the local machine. The *filename* option specifies the boot loader location on the TFTP server. The location of the file is relative to the main TFTP directory.

Any standard DHCP configuration file is supported. The sample file format above shows one entry (line 12-15) for each target host. The DHCP service must be restarted every time you modify the configuration file.

4. Enable the *TFTP* service. Edit the `/etc/xinetd.d/tftp` file to change the `disable` flag as `no` (default=`no`).
5. Restart the following services:

- For Oracle Linux 6:

```
service dhcpd restart  
service xinetd restart
```

- For Oracle Linux 7 and above:

```
systemctl restart dhcpd  
systemctl restart xinetd
```

6. Install Oracle Management Agent. This step is not necessary if the DHCP and Boot servers are installed on the Cloud Control server. For more information see: Oracle Enterprise Manager Cloud Control Basic Installation Guide.

Setting Up RPM Repository



Note:

It is recommended that you use RAM of 2 GB.

Setting Up Oracle Linux RPM Repository

You can set up Oracle Linux Repository by using the Oracle Linux installation media as follows:

1. Download Oracle Linux from <http://edelivery.oracle.com/linux>.
2. Copy all the contents of the first CD to a directory, for example: `Root Directory`.
3. Add custom RPM to the repository as follows:
 - a. If there are custom RPM installed on the reference host that need to be provisioned on the bare metal machine, make sure to copy them to the directory containing the RPMS packages.
 - b. Run the `createrepo <Root Directory>` command on this directory.
4. Create a symbolic link in `/var/www/html` to `<Directory>` directory.

The repository should now be available through HTTP if an Apache server is running.

Exposing RPM Repository through HTTP or FTP

To expose RPM Repository through HTTP, follow these steps:

1. Ensure that Apache Web Server is installed and HTTP service is running.
2. Create a symbolic link in document root to RPM Repository directory. For example, `/var/www/html` to `<RPM_REPOS>` directory.

To expose RPM Repository through FTP, ensure that FTP server is running.

Configuring Stage Server

During provisioning of an image on hardware servers, the required binaries and files are first transferred to a stage server. This is known as Staging phase and is responsible for preparing images to be installed over the network, and exposing installable or executable software elements over the network to the target hardware server being provisioned.

The Provisioning application requires at least one stage server on which all the activities related to staging can be performed. From the networking perspective, you are advised to keep the stage server as close to the target machines as possible. It will help in bringing down the installation time drastically, by reducing the time taken to transfer image data from the stage server to the hardware servers. If you have multiple hardware server groups residing at physically different locations, it would be better to have one stage server for each of these locations. Stage server should again be an Cloud Control target.

Follow these steps:

1. Log in to Cloud Control as an administrator.
2. From the **Enterprise** menu, select **Provisioning and Patching** and then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the Stage Servers section, click **Add Server**.
4. In the Add Staging Server dialog, select a **Stage Server**, specify a **Stage Directory**, for example, `/scratch/stage`, and **Base URL**, for example, `file://stgserver.example.com/scratch/stage`. Click **OK**.

Configuring Boot Server

One of the key requirements of application is the ability of the hardware server to boot up over the network (rather than from a local boot device). A boot server must be set up so that it is able to service the requests from the designated hardware servers in order for them to boot over the network. Boot server must be an Cloud Control target and should be able to receive the BOOTP and TFTP (Trivial File Transfer Protocol) requests over the network from the hardware server. Refer to Setting Up Boot Server for setting up a boot server with DHCP/TFTP combination.

Follow these steps:

1. Make sure that you have administrator privileges.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the Boot Servers section, click **Add**.
4. In the Add Boot Server dialog, select a **Boot Server** and specify a **TFTP Boot Directory**. For example: `/var/lib/tftpboot`. Click **OK**.

Configuring DHCP Server

Follow these steps:

1. Ensure that you have administrator privileges.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the DHCP Servers section, click **Add**.
4. In the Add DHCP Server dialog, select a **DHCP Server** and specify a **DHCP Configuration File**, for example, `/etc/dhcpd.conf` that has been modified to support your target hosts. Click **OK**.

Configuring RPM Repository

The Provisioning application picks up the RPM for the operating system from the RPM repository. At least one repository needs to be setup for use by the Provisioning application.

To configure the RPM repository follow these steps:

1. Ensure that you have administrator privileges.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the RPM Repositories section, click **Add**.
4. In the Add RPM Repository Server dialog, specify a **Repository Name** and **URL**. For RPM repository either accessible by HTTP or on a local server, specify the URL in the HTTP format, for example, `http://example.com/OL7/`. For NFS location, specify the URL as `file://example/OL7/`.

Click **OK**.

Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host

Ensure that the following criteria are met before provisioning:

Table 2-2 Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host

Resource Name	Checklist
Boot Server	<ul style="list-style-type: none"> • DHCP server is up and running. • The <code>next_server</code> entry in <code>/etc/dhcp/dhcpd.conf</code> file points to this boot server. • TFTP is up and running. • Boot Server is present in the same Subnet where the target machines to be provisioned are present or will be added. • Management Agent is installed. • Boot server machine is visible as a managed target in Cloud Control. • A brand new PXE-bootable box actually detects the boot server and starts to boot it, even if no image is installed yet.
Stage Server	<ul style="list-style-type: none"> • Large storage, High Memory and Sufficient Memory. • If NAS server is used for storage then it should have NFS support. • Management Agent is installed. • Boot server machine is visible as a managed target in Cloud Control. • The required agent rpm is staged for installing agents on targets. • Preferred Credentials are set. • Stage server is reachable from the box to be provisioned, or the same Subnet.
RPM Repository	<ul style="list-style-type: none"> • RPM Repository is as close as possible to the target servers. • Install tree structure is as indicated in Configure RPM repository section. • RPM repository is available via HTTP. • Provide the exact URL and test the RPM repository access over HTTP.
Reference Host	<ul style="list-style-type: none"> • Agent is installed on local disk and not on NFS mounted directory. • Preferred Credentials are set.
Software Library	<ul style="list-style-type: none"> • Shared storage used for Software Library is accessible through NFS mount points to all OMS servers.

Configuring Software Library Components

To set up and configure the Software Library, see *Setting Up Oracle Software Library in Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

You can create the following Bare Metal provisioning entities and store them in Software Library:

- [Creating Operating System Component](#)
- [Creating Disk Layout Component](#)
- [Creating an Oracle Virtual Server Component](#)

Creating Operating System Component

Follow these steps to create an operating system component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the Software Library Home, from the **Actions** menu, select **Create Folder**.
3. In the Create Folder pop up, specify a **Name** and **Description** for the folder and select the folder location. For example, create a folder `BMP-OL77` to represent the components you will use to provision a bare metal server of Oracle Linux 7 Update 7 Click **Save**.
4. From the **Actions** menu, select **Create Entity** and then **Bare Metal Provisioning Components**.
5. In the Create Entity: Bare Metal Provisioning Components dialog box, select **Operating System Component** and click **Continue**.
6. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

 **Note:**

The component name must be unique to the parent folder that it resides in. Sometime even when you enter a unique name, it may report a conflict, this is because there could be an entity with the same name in the folder that is not visible to you, as you do not have view privilege on it.

Click **+Add** to attach files that describe the entity better such as `readme`, `collateral`, `licensing`, and so on. Ensure that the file size is less than 2 MB.

In the **Notes** field, include information related to the entity such as changes being made to the entity or modification history that you want to track.

7. In the Basic Operating System page, select a **Time Zone** and specify the **Root Password**.

In the Operating System Users List, add the users for the operating system by specifying the **User Name**, **Password**, **Primary Group**, and **Additional Groups**. Specify if you want to **Enable Sudo Access** for the user.

In the Fetch Configuration properties from Reference Enterprise Manager Host target section, select **Fetch Properties** to apply the host properties. Select the reference host and select the **Configurations** you want to fetch.

Click **Next**.

8. In the Advanced Configuration page, specify the agent properties, boot configuration, and other configuration as explained in the tables.

The Configure Package Selection section displays the packages from the operating component or reference host you specified in the previous screen. You can retain or remove these packages from the component.

Click **Next**.

9. In the Review page, verify the information and click **Finish**.

The operating system component will be saved in Software Library with the status Ready.

Table 2-3 Agent Settings

Element	Description
Install User	User name for installing the agent.
Install Group	Install group for agent.
Agent Registration Password	Specify the password to be used to register the agent with Oracle Management Server.
RPM URL	Location where agent RPM are stored.

Table 2-4 Additional OS Configuration

Element	Description
Require TTY	Select this option if you want <code>sudo</code> user to Log in to a separate terminal.
SELinux	You can choose to enable or disable SELinux.
Mount Point Settings	Specify entries for the <code>/etc/fstab</code> file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
NIS Settings	Specify entries for the <code>/etc/yp.conf</code> file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
NTP Settings	Specify entries for the <code>/etc/ntp.conf</code> file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for <code>/etc/inittab</code> file. All processes are started as part <code>init</code> operation in boot process. <code>init</code> operation decides the processes that will start on booting of a machine or when runlevel changes.
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 2-5 Boot Configuration and Configuration Scripts

Element	Description
Advanced Configuration & Power Interface	Specify settings for boot time parameter for kernel (<code>acpi</code>) in the <code>/boot/grub/grub.conf</code> file.
Use Para-Virtualized kernel	Select if you are using para-virtualized kernel.
Post Install Script	Specify any set of commands that need to be executed on the newly provisioned machine. These commands will be appended to the post section of the kickstart file.
First Boot Script	Specify any set of commands that need to be executed on the newly provisioned machine when it starts up for the first time.

Creating Disk Layout Component

Follow these steps to create a disk layout component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.

2. From the Software Library Home, from the **Actions** menu, select **Create Folder**.
3. In the Create Folder popup, specify a **Name** and **Description** for the folder and select the folder location. Click **Save**.
4. From the **Actions** menu, select **Create Entity** and then **Bare Metal Provisioning Components**.
5. In the Create Entity: Bare Metal Provisioning Components dialog box, select **Disk Layout Component** and click **Continue**.
6. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

 **Note:**

The component name must be unique to the parent folder that it resides in. On some occasions when you enter a unique name, it may report a conflict. This is because there may be an entity with the same name in the folder that is not visible to you, as you do not have view privilege on it.

- Click **+Add** to attach files that describe the entity better like readme, collateral, licensing, and so on. Ensure that the file size is less than 2 MB.
- In the **Notes** field, include information related to the entity like changes being made to the entity or modification history that you want to track.

7. In the Configure page, specify the hard disk, RAID, partition, and logical configurations.
 - To specify the hard disk profile, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**.
 - To specify the Partition Configuration, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.
 - To specify RAID configuration, click **Add**. Specify the **Device Name** and **Capacity**.
 - To specify the Logical Volume Group Configuration, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.
 - To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

 **Note:**

Valid characters for **Logical Volume Name** and **Logical Group Name** names are: a-z, A-Z, 0-9, and these special characters: + _ . - . See **lvms(8)** for valid names.

Click **Next**.

8. In the Review page, verify the information and click **Finish**.

The disk layout component will be saved in Software Library with the status Ready.

Creating an Oracle Virtual Server Component

Follow these steps to create an operating system component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the Software Library Home, from the **Actions** menu, select **Create Folder**.
3. In the Create Folder popup, specify a **Name** and **Description** for the folder and select the folder location.
4. From the **Actions** menu, select **Create Entity** and then **Bare Metal Provisioning Components**.
5. In the Create Entity: Bare Metal Provisioning Components dialog box, select **Oracle Virtual Server Component** and click **Continue**.
6. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

Note: The component name must be unique to the parent folder that it resides in. Sometime even when you enter a unique name, it may report a conflict, this is because there could be an entity with the same name in the folder that is not visible to you, as you do not have view privilege on it.

Click **+Add** to attach files that describe the entity better like readme, collateral, licensing, and so on. Ensure that the file size is less than 2 MB.

In the **Notes** field, include information related to the entity like changes being made to the entity or modification history that you want to track.

7. In the Basic Operating System page, select a **Time Zone** and specify the **Root Password** and the **OVM Agent Password**.

In the Operating System Users List, add the users for the operating system by specifying the **User Name**, **Password**, **Primary Group**, and **Additional Groups**. Specify if you want to **Enable Sudo Access** for the user.

Click **Next**.

8. In the Advanced Configuration page, specify the Dom0 Configuration, Boot Configurations, and Additional OS Details as explained in the tables.

Click **Next**.

9. In the Review page, verify the information and click **Finish**.

The oracle virtual server component will be saved in the Software Library with the status Ready.

Table 2-6 Additional OS Details

Element	Description
Mount Point Settings	Specify entries for the /etc/fstab file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.

Table 2-6 (Cont.) Additional OS Details

Element	Description
NIS Settings	Specify entries for the /etc/yp.conf file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
NTP Settings	Specify entries for the /etc/ntp.conf file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for/etc/inittab file. All processes are started as part of init operation in boot process. Init operation decides the processes that will start on booting of a machine or when runlevel changes.
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 2-7 Boot Configuration and Configuration Scripts

Element	Description
Post Install Script	Specify any set of commands that need to be executed on the newly provisioned machine. These commands will be appended to the post section of the kickstart file.
First Boot Script	Specify any set of commands that need to be executed on the newly provisioned machine when it starts up for the first time.

Prerequisites for Provisioning Operating Systems and Oracle VM Servers

- Ensure that you meet the prerequisites described in *Setting Up Oracle Software Library* in *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.
- Ensure that you set up the bare metal provisioning infrastructure described in [Setting Up Infrastructure for Bare Metal Provisioning](#).
- Ensure that you have Cloud Control administrator privileges.

Provisioning Operating Systems

The following sections explain how to provision Linux on bare metal boxes:

Follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In the Server Image section, from the **Provision** menu, select **Operating System**.

3. In the General/Target Selection page, in the General section, specify the **Deployment Name**. Select the **Operating System** you want to provision and provide a description. Select the **Patching Groups** and **Monitoring Templates** you want to associate with the system.

In the Target Selection section, select the Provisioning Category as one of the following:

- **MAC Addresses** if you want to provision the bare metal systems by specifying MAC addresses. Click **Add** to specify the list of MAC Address. In the Add MAC dialog box, specify the MAC addresses. Click **OK**.
Optionally, click **Add from File** to add the MAC address from a file. In the Add from File dialog box, click **Browse** and select the file from the location where you have stored it.
- **Subnet** to specify the subnet for the bare metal provisioning. In the Subnet to be Provisioned section, specify the **Subnet IP, Netmask, Number of Network Interfaces**, and **Bootable Network Interface**.
- **Re-image EM Host Targets** to re-provision an existing Cloud Control host target. In the Enterprise Manager Hosts to be Provisioned section, click **Add** to search and select the host target. Click **OK**. Select the **Bootable Network Interface**.

Optionally, you can click **Save As Plan** and save the configuration details you have specified. Specify a name and description and click **OK** to save the plan. You can later use the saved plan to provision bare metal boxes. You can save as plan on any page of the wizard or configure the wizard completely and save the plan on the last page of the wizard.

Click **Next**.

4. In the Deployment page, in the Infrastructure section, specify:
 - a. **Stage Server** and select the **Storage**. Select **Run Stage Server Pre-requisite checks** to check if the stage server is configured properly.
 - b. **Boot Server** and select **Run Boot Server Pre-requisite checks** to check if the Boot server is configured properly.
 - c. **DHCP Server** and select **Run DHCP Server Pre-requisite checks** to check if the DHCP server is configured properly.
 - d. **Local RPM Repository**.

In the Fetch Configuration Properties from Pre-Created Components section, select the **Operating System Component, Disk Layout Component**, and **Provisioning Directive** from the Software Library. Otherwise, you can specify the operating system, disk layout, and other properties in the respective pages.

Click **Next**.

5. In the Basic OS Details page, set the **Time Zone** and **OS Root Password**. In the Add Operating System Users list section, click **Add**. Specify the **User Name, Password, Primary Group**, and **Additional Groups** to add the operating system users. Enable or Disable sudo access. Click **OK**.

If you have a reference host from which you want to provision your bare metal servers, then in the Fetch Properties from Reference Enterprise Manager Host Target section, select **Fetch Properties** to select reference host properties. Select the reference host and the configurations you want to fetch. Specify reference host credentials. The credentials you specify must have root access or you must have sudo privileges set up for the target.

You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password

and select the Run Privilege. Choose to **Save Credentials** to use these credentials in future deployments.

Click **Next**.

6. In the Additional OS Details page, specify agent settings, configuration scripts, package selection, and additional operating system configuration, and boot configuration as explained in [Table 2-8](#), [Table 2-9](#), and [Table 2-10](#). The Configure Package Selection section displays the packages from the operating component or reference host you specified in the previous screen. You can retain or remove these packages for your provisioning operation.

If you have selected an OS component in step 4, these settings will be displayed here. You can edit or retain these values.

Click **Next**.

7. In the Disk Layout page, specify hard disk profile, partition configuration, RAID configuration, Logical Volume Group configuration, and Logical Volume configuration.
 - To specify the hard disk profile, click **Add**. Specify the **Device Name** and **Capacity**.
 - To specify the Partition Configuration, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.
 - To specify RAID Configuration, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**. To configure RAID, ensure that your hard disk has two partitions at the minimum.
 - To specify the Logical Volume Group Configuration, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.
 - To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

 **Note:**

Valid characters for **Logical Volume Name** and **Logical Group Name** names are: a-z, A-Z, 0-9, and these special characters: + _ . -. See **lvm(8)** for valid names.

 **Note:**

If you selected a Disk Layout component in step 4, these settings will be displayed here. You can edit, remove, or retain these values.

Click **Next**.

8. In the Network page, the network properties for the MAC Address or Subnet as specified during target selection, is displayed. Click **Add** to configure the network interfaces. In the Input Network Interface Properties dialog box, specify the Interface name. Select the **Configuration Type** as:

- **Static** if you want to specify the IP addresses
- **DHCP** if you want the DHCP server to assign a network address
- **Network Profile** if you want to assign network addresses from a network profile.

Select the **Interface Type** as bond master, slave, or non-bonding.

Click **Next**.

9. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Specify the Stage Server and Boot Server credentials. You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password and select the Run Privilege. Choose to **Save Credentials** to use these credentials in future deployments. Click **Next**.
10. In the Review page, verify that the details you have selected are correctly displayed and submit the job for the deployment. If you want to modify the details, click **Back** repeatedly to reach the page where you want to make the changes. Click **Save As Plan** to save the configuration details you have specified. Specify a name and description and click **OK** to save the plan. You can later use the saved plan to provision bare metal boxes. For more information, see [Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers](#). Click **Submit**.
11. The Deployment Procedure is displayed in the Bare Metal Provisioning page with Status Running. Click on the Status message.
12. In the Procedure Activity page, view the job steps and verify that Status is Success. If the status is Failed, view the steps that have failed, and fix them and resubmit the job.
13. After bare metal systems have been provisioned, verify that they appear in the All Targets page.

Table 2-8 Agent Settings

Element	Description
Install User	User name for installing the agent.
Install Group	Install group for agent.
Agent Registration Password	Specify the password to be used to register the agent with Oracle Management Server.
Agent RPM URL	Agent RPM location.

Table 2-9 Additional OS Configuration

Element	Description
Require TTY	Select this option if you want <code>sudo</code> user to Log in to a separate terminal.
SELinux	You can choose to enable or disable SELinux.
Mount Point Settings	Specify entries for the <code>/etc/fstab</code> file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
NIS Settings	Specify entries for the <code>/etc/yp.conf</code> file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.

Table 2-9 (Cont.) Additional OS Configuration

Element	Description
NTP Settings	Specify entries for the <code>/etc/ntp.conf</code> file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for the <code>/etc/inittab</code> file. All processes are started as part init operation in boot process. Init operation decides the processes that will start on booting of a machine or when runlevel changes.
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 2-10 Boot Configuration and Configuration Scripts

Element	Description
Advanced Configuration & Power Interface	Specify settings for boot time parameter for kernel (acpi) in the <code>/boot/grub/grub.conf</code> file.
Use Para-Virtualized kernel	Select if you are using para-virtualized kernel.
Post Install Script	Specify any set of commands that need to be executed on the newly provisioned machine. These commands will be appended to the post section of the kickstart file.
First Boot Script	Specify any set of commands that need to be executed on the newly provisioned machine when it starts up for the first time.

 **Note:**

After Linux is provisioned on the bare metal system, out-of-box Deployment Procedures can be used to provision Database and other Oracle products on the server.

Provisioning Oracle VM Servers

To provision an Oracle VM server on a bare metal box, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In the Server Image section, from the **Provision** menu, select **Oracle VM Server**.
3. In the General/Target Selection page, in the **General** section, specify a unique **Deployment Name**. In the **Target Selection** section, select one of the following **Provisioning Category**:
 - **MAC Addresses:** If you want to provision the bare metal systems by specifying MAC addresses. Click **Add** to specify the list of MAC Address. Alternately, to add the addresses from a file, click **Add from File**. In the Add from File dialog box, select the file that contains the addresses and click **OK**.

- **Subnet:** To specify the subnet for the bare metal provisioning. In the Subnet to be Provisioned section, specify the **Subnet IP, Netmask, Number of Network Interfaces,** and **Bootable Network Interface.**

The Oracle VM Registration section allows you to select an OVM Manager registered in cloud to manage the Oracle VM servers you are provisioning. To do so, click the search icon. From the Select Target dialog box, select a target VM machine, and click **Select**. Optionally, you can click **Save As Plan** and save the configuration details you have specified. Specify a name and description and click **OK** to save the plan. You can later use the saved plan to provision future bare metal boxes. You can **Save As plan** on any page of the wizard or configure the wizard completely and save the plan on the last page.

Click **Next**.

4. In the Deployment page, in the Infrastructure section:
 - a. Select **Stage Server**, and a location on the stage server for preparing images to be installed over the network. Select **Run Stage Server Pre-requisite checks** to check if the stage server is configured properly.
 - b. Select **Boot Server**, and select **Run Boot Server Pre-requisite checks** to check if the Boot server is configured properly.
 - c. Select **DHCP Server**, and select **Run DHCP Server Pre-requisite checks** to check if the DHCP server is configured properly
 - d. Select **Local RPM Repository** from the available list.

In the Fetch Configuration Properties from Pre-Created Components section, select an existing **Operating System Component, Disk Layout Component,** and **Provisioning Directive** from the Software Library home page. By doing so, the property values of the selected entities are fetched from Software Library, and are populated accordingly. Doing so allows you can skip updating the remaining pages in the wizard and directly go to the scheduling page. However, if you do not have the required entities on Software Library, then you can specify the operating system, disk layout, and other properties in the subsequent pages.

Click **Next**.

5. In the Basic OS Details page, set the **Time Zone, OS Root Password,** and the Oracle VM Agent password. In the Operating System Users list section, click **Add**. Specify the **User Name, Password, Primary Group,** and **Additional Groups** to add the operating system users. Enable or Disable sudo access. Click **OK**.
Click **Next**.
6. In the Additional OS Details page, do the following:
 - a. In the Dom0 Configuration section, you can provide the memory and power requirements for the target provisioned.
 - b. In the Additional OS details, you can click the configure icon to add certain other configuration details such as: **Mount Point Settings, NIS Settings, NTP Settings, Kernel Parameter Settings, Inittab Settings, Firewall Settings.**
 - c. In the Boot Configuration section, click **First Boot** to add commands/scripts that must be run on the system when it boots for the first time after installation. Click **Post Install** to provide commands to run on the system once the installation is complete.

Click **Next**.

7. In the Disk Layout page, specify **Hard Disk Profile, RAID configuration,** and **Logical Configuration.**

- a. To specify the **Hard Disk Profile**, click **Add**. Specify the **Device Name** and **Capacity**.
- b. To specify the **Partition Configuration**, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.
- c. To specify **RAID Configuration**, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**. To configure RAID, ensure that your hard disk has two partitions at the minimum.
- d. To specify the **Logical Volume Group Configuration**, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.
- e. To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

 **Note:**

Valid characters for **Logical Volume Name** and **Logical Group Name** names are: a-z, A-Z, 0-9, and these special characters: + _ .
-. See **lvm(8)** for valid names.

 **Note:**

If you selected a Disk Layout component in step 4, these settings will be displayed here. You can edit, remove, or retain these values.

Click **Next**.

8. In the Network page, the network properties for the MAC Address or Subnet as specified during target selection, is displayed.
Click **Add** to configure the network interfaces. In the **Add Network Interface** dialog box, specify the Interface name. Select the **Configuration Type** as:

- **Static** if you want to specify the IP addresses.
- **DHCP** if you want the DHCP server to assign a network address.
- **Network Profile** if you want to assign network addresses from a network profile.

Select the **Interface Type** as bond master, slave, or non-bonding.
Click **Next**.

9. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Specify the Stage Server and Boot Server credentials. You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password and select the Run Privilege. Choose **Save Credentials** to use these credentials in future deployments.
Click **Next**.
10. In the Review page, verify that the details you have selected are correctly displayed and submit the job for the deployment. If you want to modify the details, click **Back** repeatedly to reach the page where you want to make the changes.
Click **Submit**.

Click **Save As Plan** to save the configuration details you have specified. Specify a name and description and click **OK** to save the plan. You can later use the saved plan to provision future bare metal boxes. For more information, see: [Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers](#) .

11. The **Deployment Procedure** is displayed in the Bare Metal Provisioning page with Status *Running*. Click the Confirmation message.
12. In the **Procedure Activity** page, view the job steps and verify that **Status** is *Success*. If the status is *Failed*, review the failed steps, fix and resubmit the job.
13. After bare metal systems have been provisioned, verify that they appear in the All Targets page.

Table 2-11 Additional OS Configuration

Element	Description
Mount Point Settings	Specify entries for the <code>/etc/fstab</code> file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
NIS Settings	Specify entries for the <code>/etc/yp.conf</code> file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
NTP Settings	Specify entries for the <code>/etc/ntp.conf</code> file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for <code>/etc/inittab</code> file. All processes are started as part <code>init</code> operation in boot process. Init operation decides the processes that will start on booting of a machine or when <code>runlevel</code> changes.
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 2-12 Boot Configuration and Configuration Scripts

Element	Description
Post Install Script	Specify any set of commands that need to be executed on the newly provisioned machine. These commands will be appended to the post section of the kickstart file.
First Boot Script	Specify any set of commands that need to be executed on the newly provisioned machine when it starts up for the first time.

Viewing Saved Plans

To view saved plans, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. On the Bare Metal Provisioning page, click **Saved Plans**.

**Note:**

To edit the saved plans, see [Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers](#).

Using Saved Plans for Provisioning Linux Operating System and Oracle VM Server on Bare Metal Servers

To edit the saved plans, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In Server Image section, from Provision menu, select Using Saved Plan.
3. From the Saved Plans dialog box, select any template to pre-populate the provisioning wizard with the saved values, and click **Continue**.
4. Update the Deployment Name, the Provisioning Category information in the General/Target selection page.
5. Follow steps 4 to step 10 listed in the section [Provisioning Operating Systems](#).
6. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Also, update the Stage Server and Boot Server credentials.
7. In the Review page, verify all the details you have selected, and click **Submit**.

3

Patching Linux Hosts

This chapter explains how you can patch Linux hosts using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Overview of Patching Linux Hosts](#)
- [About the Deployment Procedure for Patching Linux Hosts](#)
- [Supported Linux Releases](#)
- [Setting Up Infrastructure for Linux Patching](#)
- [Patching Linux Hosts](#)
- [Additional Linux Patching Tasks You Can Perform](#)
- [Managing Linux Configuration Files](#)

Overview of Patching Linux Hosts

Linux Host Patching is a feature in Cloud Control that keeps the hosts in an enterprise updated with security fixes and critical bug fixes, especially in a data centre or a server farm. This feature in Cloud Control enables you to:

- Set up Linux RPM Repository based on Unbreakable Linux Network (ULN) channels
- Download Advisories (Erratas) from ULN
- Set up a Linux Patching group to update a group of Linux hosts and collect compliance information
- Allow non-compliant packages to be patched
- Rollback/uninstall packages from a host
- Manage RPM repositories and channels (clone channels, copy packages from one channel into another, delete channels)
- Add RPMs to custom channels
- Manage configuration file channels (create/delete channels, upload files, copy files from one channel into another)

The following are concepts related to Linux patching:

Type	Description
Linux Host	A host target in Cloud Control that is running the Linux operating system.
Linux Patching Group	A set of managed Linux hosts that are associated with a common list of RPM repositories. Every group is configured with an update schedule, according to which a recurring job is triggered, that will update the hosts of the group with the associated RPM repositories.
Unbreakable Linux Network (ULN)	Unbreakable Linux Network (ULN) is a Web site hosted by Oracle to provide updates for Oracle Linux.

Type	Description
ULN Channel	A channel is a group of RPM packages on ULN. For example, the <code>ol6_latest</code> channel contains all the packages for Oracle Linux 6.
RPM Repository	RPM repository is a directory that contains RPM packages and their metadata (extracted by running <code>yum-arch</code> and <code>createrepo</code>). The RPM repository is accessible via http or ftp. An RPM repository can be organized to contain packages from multiple channels. For example, <code>/var/www/html/yum/OracleLinux/OL6/latest</code> might contain packages from the <code>ol6_latest</code> channel on ULN.
Custom Channel	A channel that is created by the user to store a set of custom RPM packages. Custom channels can be added to the RPM repository.
Configuration Channel	A channel that is created by the user to store a set of Linux configuration files. Configuration channels can be used in the Linux patching application GUI to deploy configuration files on Linux hosts.

About the Deployment Procedure for Patching Linux Hosts

Cloud Control provides the following deployment procedures for Linux patching:

- *Patch Linux Hosts*
This deployment procedure enables you to patch Linux hosts.
- *Linux RPM Repository server setup*
This deployment procedure enables you to set up a Linux RPM repository server. To set up the Linux RPM repository server, see [Setting Up the RPM Repository for Patching](#).

Supported Linux Releases

The following releases are supported for Linux patching:

- Oracle Linux 6
- Oracle Linux 7
- Oracle Linux 8

Setting Up Infrastructure for Linux Patching

This section describes the setup requirements for Linux patching. In particular, this section describes the following:

- [Prerequisites for Using the Linux Patching Feature](#)
- [Setting Up the RPM Repository for Linux Patching](#)
- [Setting Up Linux Patching Group for Compliance Reporting](#)

Prerequisites for Using the Linux Patching Feature

To use the Linux Patching feature, meet the following prerequisites:

1. Meet the basic prerequisites described in Setting Up Your Infrastructure
2. Install yum on all your Oracle Linux 6 and above target hosts.
3. Enable the following commands through SUDO:
 - /bin/cp
 - /bin/rm
 - /bin/chmod
 - /sbin/chkconfig
 - yum
 - sed
 - rpm

Setting Up the RPM Repository for Linux Patching

This section describes how you can set up the RPM repository. In particular, this section describes the following:

- [Prerequisites for Setting Up the RPM Repository](#)
- [Setting Up the RPM Repository for Patching](#)



Note:

The RPM repository can be set up in a shared location. This configuration is supported. The same EM repository is shared by using the `symlink` (symbolic link) in the folder `/var/www/html` to a shared file system. In case the host target goes down then the RPM repository also is unavailable.

The RPM repository can exist on the OMS or on a non-OMS designated host target.

Prerequisites for Setting Up the RPM Repository

Before setting up the RPM repository, meet the following prerequisites:

- Identify a Redhat or Oracle Linux host, install a Management Agent, and point to the OMS. This host must have the `sudo` package installed.
- Obtain a valid Customer Support Identifier (CSI) number from your Oracle sales representative.
After obtaining a valid CSI number, ensure that you create a ULN account. To create a ULN account, access the following URL:
<https://linux.oracle.com/register>
- Ensure that the `/var/www/html/` directory on the host on which you plan to set up the RPM repository has at least 60 GB of free disk space per channel.
- Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.
For example: `http://h1.example.com`. If this works, then it is confirmed that Apache is installed and listening on port 80.

- Ensure that the `createrepo` package is installed on the RPM Repository host. To obtain this package, subscribe to the `el*_addon` or the `ol*_addon` channel.
- Ensure that the `uln-yum-mirror` and `yum-utils` packages are installed on the RPM Repository host. To obtain the `yum-arch` and the `uln-yum-proxy/uln-yum-mirrors` packages, subscribe to the add-ons channel. To obtain the `yum-utils` package, subscribe to the latest channel.
- If the RPM Repository host is not running on the same operating system as the channels, you must import the public key manually to the channel belonging to the release. To do so, follow these steps:
 1. Download the OL 6 or above key from: <http://public-yum.oracle.com/RPM-GPG-KEY-oracle-ol6>
 2. Store it under the following directory on your host: `/usr/share/rhn`
 3. Run the following command:`rpm --import /usr/share/rhn/RPM-GPG-KEY-oracle-ol6`

 **Note:**

Modify the path accordingly to the channel belonging to the release.

- Ensure that the Enterprise Manager user has the `EM_LINUX_PATCHING_ADMIN` role. If the Enterprise Manager user does not have this role, ensure that the super user grants them.
- Ensure that the Oracle GPG keys are installed on the host on which you plan to set up the RPM Repository.
To install the Oracle GPG keys on a host running on Oracle Linux 6 or above, run the following command:`rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY.`

Setting Up the RPM Repository for Patching

Log in with super user privileges to set up an RPM Repository that downloads latest RPM packages and advisories from ULN. Follow these steps:

To set up an RPM Repository that downloads the latest RPM packages and advisories from ULN, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, click **Setup RPM Repository**.
3. On the Setup RPM Repository page, in the RPM Repository Server section, select the RPM Repository server by clicking the search icon. Select the host assigned for subscribing to ULN.
4. In the Credentials section, ensure that the **Normal Host Credential** user has write access to the stage location, and the **Privileged Host Credential** user can sudo with root privilege. Click **Apply**.
5. In the Deployment Procedure submission confirmation, click **Linux RPM Repository Server Setup**. The deployment procedure starts a job to download latest RPM packages and Advisories from the subscribed ULN channels.

6. (Optional) If you want to change the refresh mode to 30 seconds, then from the **View Data** list, select **Real Time: 30 Second Refresh**.
7. Click the status of the manual step **Register with ULN** to verify if your host has been registered to ULN.
If you have registered your host to ULN, then select the target and click **Confirm**, and then click **Done** to go to the main flow.

If you have not registered your host to ULN, then perform the following steps on your Linux host:
 - a. Log in to the RPM Repository server machine.
 - b. Check if your host can connect to ULN. If your host cannot connect to the ULN directly, you can configure **uln_register** to use a proxy server. To configure access to ULN using a proxy server, follow these instructions: https://linux.oracle.com/uln_faq.html#9
 - c. Register the host to ULN by following the steps at: https://linux.oracle.com/uln_faq.html#2

 **Note:**

While registering, you can choose the user name and password. This credential will be used to log in to <http://linux.oracle.com>

8. Click the status of the step **Subscribe to ULN channels**. When you register a Linux server to ULN, it will be subscribed to a channel that has the latest Oracle Linux packages for the appropriate architecture. If no additional channels are needed to be subscribed to your host, then select the target and click **Confirm**, and then click **Done** to go to the main flow. If some additional channels are needed to be subscribed to your host, then perform the following steps:
 - a. Log in to ULN:<http://linux.oracle.com/>
 - b. Click on the **Systems** tab to manage subscriptions for each subscribed server.
 - c. Subscribe to all the additional channels you need.

 **Note:**

- Make sure to install the `createrepo` package from the `ol*_latest` channel.
- Make sure that the `uln-yum-mirror` and `yum-utils` packages are installed on your Linux host. To obtain the `uln-yum-mirror` package, subscribe to the Add On channel. To obtain the `yum-utils` package, subscribe to the latest channel.

- d. Verify the list of subscribed channels on ULN.
9. Once the deployment procedure ends successfully, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
10. On the Patching Setup page, click **Manage RPM Repository** to verify if the ULN channels are displayed in the Cloud Control console.

11. On the Manage RPM Repository page, check if all the subscribed channels are listed and if all the packages are downloaded.

Setting Up Linux Patching Group for Compliance Reporting

This section describes how you can set up a Linux Patching group for compliance reporting by associating the group with the RPM Repository (each subscribed ULN channel is a repository) created in [Setting Up the RPM Repository for Linux Patching](#).

In particular, this section describes the following:

- [Prerequisites for Setting Up Linux Patching Group](#)
- [Setting Up a Linux Patching Group](#)

Prerequisites for Setting Up Linux Patching Group

Before setting up the Linux Patching Group, meet the following prerequisites:

- Set up RPM Repository server or set a custom RPM Repository as a channel in Cloud Control.
- Install yum on all your Oracle Linux 6 and above target hosts.
- Install `Sudo` on the target hosts.
- Ensure that the Enterprise Manager user logs in to the OMS with super user privileges.
- Ensure that the Enterprise Manager user has the `EM_LINUX_PATCHING_ADMIN` role. If the Enterprise Manager user does not have these, make sure that the Super User grants them.

Setting Up a Linux Patching Group

Log in with Super User privileges to set up a Linux patching group. Follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching Setup page, click **Setup Patching Groups**.
3. On the Setup Patching Groups page, click **Create**.
4. On the Create Group: Properties page, enter a unique name for the group. Select the maturity level, Linux distribution, and Linux hosts to be added to the group. Click **Next**.
5. On the Create Group: Package Repositories page, select the RPM Repositories that must be associated with the patching group (click the search icon to select repository).

In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum performs a GPG signature check on the packages obtained from the specified repositories. Sometimes, yum may require a public GPG key to verify the packages obtained from the repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.

In the Stage Location section, specify the location where you want the Linux patching configuration and log files to be created.

In the Update Hosts section, select **Automatically Update Hosts** if you want to auto-update the host, that is, to schedule an update job (schedule specified as one of the subsequent step) to update all non-compliant packages from the selected package repository.

In the Excluded Packages section, for **Excluded Packages**, specify the list of packages that you do not want to update while patching the Linux hosts. If the list of packages that you do not want to update during the patching process is present in a file, click **Import From File** to specify the location of the file. The wizard obtains the required packages from the specified file.

In the Rollback Last Update Session section, select **Enable 'Rollback Last Update Session'** to enable the Rollback Last Update Session feature for the group in the Undo Patching wizard. If this feature is not enabled here, it is not visible in the Undo Patching wizard for the group.

In the Package Compliance section, you can choose whether to include *Rogue* packages in compliance reporting or not.

In the Packages Updated on Reboot section, for **Packages updated on Reboot**, specify the list of packages that must be updated only when the host is rebooted.

6. Click **Next**.
7. On the Create Group: Credentials page, enter the host credentials or choose to use preferred credentials. Click **Next**.
8. On the Create Group: Patching Script page, enter any pre/post patching operations to be done. This is not a mandatory step. Click **Next**.

 **Note:**

Steps 8 and 9 will be skipped if **Automatically Update Hosts** was not selected.

9. On the Schedule page, set the schedule for the update job. Click **Next**.
10. On the Review page, validate all the parameters. Click **Finish**.
11. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**. Verify the compliance report generated. The group created will have at least one out-of-date package.

[Table 3-1](#) describes the jobs that are submitted for setting up a Linux patching group.

Table 3-1 Jobs Submitted for Setting Up Linux Patching Group

Job	Description
Patching Configuration	This job configures all the hosts for patching. It creates configuration files to be used by the yum and up2date tools on each host. This job is executed just once on all the hosts contained in the Linux Patching group immediately.
Compliance Collection	Compares the versions of the packages already installed in each machine contained in the Linux Patching group with the package versions in the selected RPM Repositories, and generates Compliance Reports for indicating which packages are outdated. This job is executed once every 24 hours (after the group is set up) on all the hosts contained in the Linux Patching group.

Table 3-1 (Cont.) Jobs Submitted for Setting Up Linux Patching Group

Job	Description
Package Information	Collects the metadata information of each package contained in the selected RPM Repositories. This job is executed daily.
Packages Update	Updates non-compliant packages. This job will update the packages installed on the hosts in the group to ensure that they are up-to-date with respect to the package repositories for that group. This job will be submitted only when the option "Update Hosts" is selected in the step "Package Repositories" of the Linux Patching group wizard, and its schedule can be customized in the step "Schedule"

Patching Linux Hosts

This section describes how to patch your Linux hosts. It consists of the following:

- [Applying Patches on a Linux Patching Group Based on Compliance](#)
- [Applying Ad Hoc or Emergency Patches on Linux Hosts](#)



Note:

Before patching your Linux hosts, ensure that the Enterprise Manager user has the `EM_PATCH_DESIGNER` role and the `OPERATOR_ANY_TARGET` privilege. If the Enterprise Manager user does not have these, ensure that the super user grants them.

Applying Patches on a Linux Patching Group Based on Compliance

If the Linux Patching Compliance Home page reports that a particular Linux patching group is not compliant, you can choose to patch the group. To apply patches on this Linux patching group, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, in the Compliance Report section, select the Linux patching group that you want to patch, then click **Schedule Patching**.
3. On the Package Repository page, in the LINUX Distribution section, select the tool that you want to use to update the RPM.

 **Note:**

If the Linux host to be patched is running on Oracle Linux 6 (OL6) or later, then you must use the yum tool for patching. The up2date patching tool is not supported for this Linux version. If you do not use the yum tool in this scenario, the patching process fails on the *Configure Host For Patching* step with the following error: You are not selecting 'yum' as the tool to update the RPMs in this system. 'yum' is the only supported tool for updating RPMs in Oracle Linux 6 operating system

- a. If you have selected yum as the patching tool, make sure that you select the patching mode that you want to use. Select **Package update and new package installation** if you plan to update the existing packages, as well as install new packages. Select **Package update only** if you plan to only update the existing packages, and not install any new packages.
- b. In the Stage Location section, specify the location where you want the Linux patching configuration and log files to be created.
- c. In the Package Repository section, select the RPM repositories that you want to use.
- d. In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum or up2date performs a GPG signature check on the packages obtained from the specified repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.
- e. In the Advanced Options section, by default, the **Hide obsolete updates** option is selected. Selecting this option hides the obsolete packages on the Select Updates page. If you want to view these packages on the Select Updates page, ensure that you deselect this option.
- f. If you have selected yum as the patching tool, in the Advanced Options section, select one of the following patch application modes:
 - **Most suitable architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, for the suitable RPM architectures that are installed on the Linux hosts that you are patching.
If you select this option, Cloud Control runs the following yum command: `yum install|update packagename`
 - **Specific architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, on only those Linux hosts that have the RPM architecture of the selected package.
If you select this option, Cloud Control runs the following yum command: `yum install|update packagename.arch`
 - **Specific version and architecture**, if you want yum to install only the specific version of the package selected on the Select Updates page, or update the existing version of the package to this specific version, on only those Linux hosts that have the RPM architecture of the selected package.
If you select this option, Cloud Control runs the following yum command: `yum install|update epoch:packagename-ver-rel.arch`

Click **Next**.

4. On the Select Updates page, select the packages to be updated.

 **Note:**

If the **Hide obsolete updates** option was selected in the previous step, the values for **Total packages available** and **Total packages available in this view** may be different. This difference corresponds to the number of obsolete packages present in the repositories.

Click **Next**.

5. On the Select Hosts page, select the Linux hosts to be updated. You can also select a group by changing the target type to group.

By default, every discovered Linux host is displayed on this page, and can be selected. However, if you want only those hosts that have an older version of at least one of the packages (that you selected for the update operation in the previous step) to be displayed on this page, run the following

```
command: $<OMS_HOME>/bin/emctl set property -name  
'oracle.sysman.core.ospatch.filter_uptodate_hosts' -value 'true'
```

Click **Next**.

6. On the Credentials page, enter the credentials to be used for the updates.

Click **Next**.

7. On the Pre/Post script page, enter the scripts that need to be executed before/after the patching process, if any.

Click **Next**.

8. On the Schedule page, enter the details of the patching schedule that must be used.

Click **Next**.

9. On the Review page, review the update parameters.

Click **Finish**. A deployment procedure is submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

Applying Ad Hoc or Emergency Patches on Linux Hosts

To quickly apply patches on your Linux hosts in an ad hoc manner, or in case of an emergency, without using a Linux patching group, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Deployment Procedure Manager page, in the Procedure Library tab, select **Patch Linux Hosts**, then click **Launch**.
3. On the Package Repository page, in the LINUX Distribution section, select the tool that you want to use to update the RPMs.

 **Note:**

If the Linux host to be patched is running on Oracle Linux 6 (OL6) or later, then you must use the yum tool for patching. The up2date patching tool is not supported for this Linux version. If you do not use the yum tool in this scenario, the patching process fails on the *Configure Host For Patching* step with the following error: You are not selecting 'yum' as the tool to update the RPMs in this system. 'yum' is the only supported tool for updating RPMs in Oracle Linux 6 operating system

- a. If you have selected yum as the patching tool: For the tool operation mode, ensure that you select **Package update and new package installation**. Since this method of patching Linux hosts without using a Linux patching group is meant for emergencies and is not based on a compliance report, you can only use it to install new packages, and not update existing packages.
- b. In the Stage Location section, specify the location where you want the Linux patching configuration and log files to be created.
- c. In the Package Repository section, select the RPM repositories that you want to use.
- d. In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum performs a GPG signature check on the packages obtained from the specified repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.
- e. In the Advanced Options section, by default, the **Hide obsolete updates** option is selected. Selecting this option hides the obsolete packages on the Select Updates page. If you want to view these packages on the Select Updates page, ensure that you deselect this option.
- f. If you have selected yum as the patching tool, in the Advanced Options section, select one of the following patch application modes:
 - **Most suitable architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, for the suitable RPM architectures that are installed on the Linux hosts that you are patching.
If you select this option, Cloud Control runs the following yum command: `yum install|update packagename`
 - **Specific architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, on only those Linux hosts that have the RPM architecture of the selected package.
If you select this option, Cloud Control runs the following yum command: `yum install|update packagename.arch`
 - **Specific version and architecture**, if you want yum to install only the specific version of the package selected on the Select Updates page, or update the existing version of the package to this specific version, on only those Linux hosts that have the RPM architecture of the selected package.
If you select this option, Cloud Control runs the following yum command: `yum install|update epoch:packagename-ver-rel.arch`

Click **Next**.

4. On the Select Updates page, select the packages to be updated.

 **Note:**

If the **Hide obsolete updates** option was selected in the previous step, the values for **Total packages available** and **Total packages available in this view** may be different. This difference corresponds to the number of obsolete packages present in the repositories.

Click **Next**.

5. On the Select Hosts page, select the Linux hosts to be updated. You can also select a group by changing the target type to group.

By default, every discovered Linux host is displayed on this page, and can be selected. However, if you want only those hosts that have an older version of at least one of the packages (that you selected for the update operation in the previous step) to be displayed on this page, run the following command:

```
$<OMS_HOME>/bin/emctl set property -name  
'oracle.sysman.core.ospatch.filter_uptodate_hosts' -value 'true'
```

Click **Next**.

6. On the Credentials page, enter the credentials to be used for the updates.

Click **Next**.

7. On the Pre/Post script page, enter the scripts that need to be executed before/after the patching process, if any.

Click **Next**.

8. On the Schedule page, enter the details of the patching schedule that must be used.

Click **Next**.

9. On the Review page, review the update parameters.

Click **Finish**. A deployment procedure is submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

Managing Linux Configuration Files

This section describes how you can manage your Linux configuration files. It consists of the following:

- [Overview of Linux Configuration Files](#)
- [Prerequisites for Managing Configuration Files](#)
- [Creating a Linux Configuration File Channel](#)
- [Uploading Linux Configuration Files to a Particular Channel](#)
- [Importing Linux Configuration Files from One Channel to Another](#)
- [Deploying Linux Configuration Files From a Particular Channel](#)
- [Deleting a Linux Configuration File Channel](#)

Overview of Linux Configuration Files

The configuration file feature enables you to manage your Linux configuration files in an efficient and convenient manner. Using this feature (which is accessible from the Linux Patching home page), you can create a Linux configuration file channel, upload the required Linux configuration files present on your local host (or on a remote host that has a Management Agent deployed on it) to the created channel, then deploy the configuration files present in the channel to a large number of target hosts in a single operation.

This feature saves you the effort of manually copying the required Linux configuration files to each target host. For example, if a HTTP server configuration file that you want to copy to a large number of target hosts is present on your local host, you can use the Linux Patching home page to create a Linux configuration file channel, upload the HTTP server configuration file to this channel, then deploy the file from this channel to the target hosts.

Prerequisites for Managing Configuration Files

Ensure that the Software Library is already configured on the OMS.

Creating a Linux Configuration File Channel

To create a configuration file channel, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, click the **Configuration Files** tab.
3. In the Configuration Files tab, click **Create Config File Channel**.
4. On the Create Configuration File Channel page, enter a unique channel name and description for the channel, and click **OK**.

You will see a confirmation message mentioning that a new configuration file channel is created.

Uploading Linux Configuration Files to a Particular Channel

This section describes how you can upload configuration files to a particular channel. In particular, this section covers the following:

- [Prerequisites for Uploading Linux Configuration Files](#)
- [Uploading Linux Configuration Files](#)

Prerequisites for Uploading Linux Configuration Files

Before uploading configuration files to a particular channel, ensure that there exists at least one configuration file on the local host or on a remote host.

Uploading Linux Configuration Files

To upload configuration files, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.

2. On the Linux Patching page, click the **Configuration Files** tab.
3. In the Configuration Files tab, select the channel that you want to upload configuration files to, then click **Upload Configuration Files**.
4. Select an appropriate upload mode. You can either upload files from local host (where the browser is running) or from a remote host (a Management Agent should be installed on that host and the Management Agent must be communicating with the OMS).
5. In the File Upload section, enter the file name, path where the file will be deployed on the target host, and browse for the file on the upload host.
6. For uploading from remote machine, click **Upload from Agent Machine**. Click **Select Target** and select the remote machine.

Before browsing for the files on this machine, set preferred credential for this machine.
7. After selecting the files, click **OK**.

You will see a confirmation message that states that files have been uploaded.

Importing Linux Configuration Files from One Channel to Another

This section describes how you can import configuration files from one channel to another. In particular, this section covers the following:

- [Prerequisites for Importing Linux Configuration Files](#)
- [Importing Linux Configuration Files](#)

Prerequisites for Importing Linux Configuration Files

Before importing configuration files, ensure that there are at least two channels.

Importing Linux Configuration Files

To import configuration files from the source channel to the target channel, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, click the **Configuration Files** tab.
3. In the Configuration Files tab, select the source channel, and click **Import Files**.
4. Select the target channel.
5. From Source channel section, select the files and copy it to the target channel section. Click **OK**.

You will see a confirmation message stating that the selected files have been imported successfully.

Deploying Linux Configuration Files From a Particular Channel

This section describes how you can deploy configuration files from a particular channel. In particular, this section covers the following:

- [Prerequisites for Deploying Linux Configuration Files](#)
- [Deploying Linux Configuration Files](#)

Prerequisites for Deploying Linux Configuration Files

Before deploying configuration files, meet the following prerequisites:

- Ensure that the privileged patching user has write permission on the target machine location where each configuration file will be staged, and has SUDO privileges too.
- Ensure that there is at least one channel with some files uploaded.

Deploying Linux Configuration Files

To deploy configuration files from a particular channel, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, click the **Configuration Files** tab.
3. In the Configuration Files tab, select the source channel, and click **Deploy Files**.
4. In the wizard that appears, select the files you want to deploy, and click **Next**.
5. Click **Add** to select the targets where you want to deploy the files.
6. Enter the credentials for the selected targets.
7. Enter the Pre/Post scripts you want to run before or after deploying the files.
8. Review the deploy parameters and click **Finish**.

A deploy job is submitted. Follow the job's link until it completes successfully.

Deleting a Linux Configuration File Channel

This section describes how you can delete configuration file channels. In particular, this section covers the following:

- [Prerequisites for Deleting a Linux Configuration File Channel](#)
- [Deleting Linux Configuration File Channels](#)

Prerequisites for Deleting a Linux Configuration File Channel

Before deleting a configuration file channel, ensure that there is at least one configuration file.

Deleting Linux Configuration File Channels

To delete a configuration file channel, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, click the **Configuration Files** tab.
3. In the Configuration Files tab, select the channel, and click **Delete**. Click **Yes**.

You will see a configuration message stating that the channel was successfully deleted.

Additional Linux Patching Tasks You Can Perform

This section describes the additional tasks you can perform using the Linux Patching Home page:

- [Viewing Linux Patching Compliance History](#)
- [Patching Non-Compliant Linux Packages](#)
- [Rolling Back Linux Patch Update Sessions or Deinstalling Packages](#)
- [Registering a Custom Package Channel](#)
- [Cloning a Package Channel](#)
- [Copying Packages from One Channel to Another](#)
- [Adding Custom Packages to a Channel](#)
- [Deleting a Package Channel](#)

Viewing Linux Patching Compliance History

This section describes how you can view the compliance history for a selected group, for a specific time period. In particular, this section covers the following:

- [Prerequisites for Viewing Linux Patching Compliance History](#)
- [Viewing Linux Patching Compliance History](#)

Prerequisites for Viewing Linux Patching Compliance History

- Ensure that you have defined at least one Linux patching group.
- Ensure that you have *View* privileges on the Linux host comprising the patching group.

Viewing Linux Patching Compliance History

To view the compliance history of a Linux patching group, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Compliance Home page, from the Related Links section, click **Compliance History**.
3. On the Compliance History page, the Groups table lists all the accessible Linux patching groups and the number of hosts corresponding to each group.
4. If there are multiple Linux patching groups, the Compliance History page displays the historical data (for a specific time period) for the first group that is listed in that table.
5. To view the compliance history of a Linux patching group, click the View icon corresponding to that group.

 **Note:**

By default, the compliance data that is displayed is retrieved from the last seven days. To view compliance history of a longer time period, select an appropriate value from the View Data drop-down list. The page refreshes to show compliance data for the selected time period.

Patching Non-Compliant Linux Packages

This section describes how you can patch non-compliant packages from the Linux Patching home page. In particular, this section covers the following:

- [Prerequisites for Patching Non-Compliant Linux Packages](#)
- [Patching Non-Compliant Linux Packages](#)

Prerequisites for Patching Non-Compliant Linux Packages

Before patching non-compliant packages, ensure that a Linux Patching group is created and the Compliance Collection job has succeeded.

Patching Non-Compliant Linux Packages

To patch non-compliant packages, follow these steps:

1. In the Patch Linux Hosts Wizard, provide the required details in the interview screens, and click **Finish** on the Review page.
2. A deployment procedure is submitted to update the host. Check if all the steps finished successfully.

Rolling Back Linux Patch Update Sessions or Deinstalling Packages

This section describes how you can rollback a patch update session, or even uninstall the unstable version completely in case that patch version is found unsuitable for has a bug or security vulnerability. In particular, this section covers the following:

- [Prerequisites for Rolling Back Linux Patch Update Sessions or Deinstalling Packages](#)
- [Rolling Back Linux Patch Update Sessions or Deinstalling Packages](#)

 **Note:**

Rolling back upgrades is supported to a certain extent. When performing an upgrade such as from *OL 6.6* to *OL 6.7*, many RPM that are dependent on others are upgraded. When you apply RPM, this dependency can be followed. However, when rolling back patch update sessions, this dependency must be followed in reverse order. This reverse operation is not supported by yum. You can use the rollback feature to rollback a patch update session, but not to completely rollback a major upgrade.

Prerequisites for Rolling Back Linux Patch Update Sessions or Deinstalling Packages

Before rolling back patch update sessions or deinstalling packages, meet the following prerequisites:

- Ensure that a Linux Patching group is created.
- Ensure that the lower version of the packages are present in the RPM repository.

Rolling Back Linux Patch Update Sessions or Deinstalling Packages

To roll back a patch update session or uninstall packages, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, in the Compliance Report section, select a group, and click **Undo Patching**.
3. On the Undo Patching: Action page, select an appropriate option:
 - **Uninstall Packages**, deinstalls a package.
 - **Rollback Last Update Session**, reverts the effects of the previous patch update session.
4. Click **Next**.
5. Provide the required details in the wizard, and on the Review page, click **Finish**.
6. A job is submitted to rollback the updates done in the previous session.
7. Examine the job submitted to see if all the steps are successful.

Registering a Custom Package Channel

This section describes how you can register a custom channel. In particular, this section covers the following:

- [Prerequisites for Registering a Custom Package Channel](#)
- [Registering a Custom Package Channel](#)

Prerequisites for Registering a Custom Package Channel

Before registering a custom channel, meet the following prerequisites:

- Ensure that the RPM Repository is under `/var/www/html` and is accessible through HTTP protocol.
- Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.
For example: `http://h1.example.com`. If this works, then it is confirmed that Apache is installed and listening on port 80.
- Ensure that metadata files are created by running `yum-arch` and `createrepo` commands.

- Ensure that a Management Agent is installed on the RPM repository host, and ensure that Management Agent is communicating with the OMS.
- Ensure that the Enterprise Manager User logs in with Super User privileges for registering a custom channel.

Registering a Custom Package Channel

To register a custom RPM Repository, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage Repository Home page, click **Register Custom Channel**.
4. On the Register Custom Channel page, enter a unique channel name.
5. Click **Browse** and select the host where the custom RPM repository was setup.
6. Enter the path where RPM repository resides. The directory location must start with `/var/www/html/`.
7. Click **OK**.

A Package Information job is submitted. Follow the job until it completes successfully.

Cloning a Package Channel

This section describes how you can clone a channel. In particular, this section covers the following:

- [Prerequisites for Cloning a Package Channel](#)
- [Cloning a Package Channel](#)

Prerequisites for Cloning a Package Channel

Before cloning a channel, meet the following prerequisites:

1. Ensure that there is at least one channel already present.
2. Ensure that there is enough space on the target channel host.
3. Ensure that the stage location of the source host does not have a directory named `createLikeSrc`, and the *Directory* for the *Target Channel* does not exist.
4. Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.

For example: `http://h1.example.com`. If this works, then it is confirmed that Apache is installed and listening on port 80.

5. Ensure that the Enterprise Manager User logs in to the OMS with Super User privileges.

Cloning a Package Channel

To clone a channel, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the source channel you want to clone, and click **Create Like**.
4. Enter the credentials to use for the source channel. The credentials must have both read and write access.
5. Enter a unique target channel name.
6. Click **Browse** to select the target host name.
7. Enter the directory location of the target channel. This directory should be under `/var/www/html`.
8. Enter the credentials to use for the target channel. This credential should have both read and write access.
9. Click **OK**.

A Create-Like job is submitted. Follow the job until it completes successfully.

Copying Packages from One Channel to Another

This section describes how you can copy packages from one channel to another. In particular, this section covers the following:

- [Prerequisites for Copying Packages from One Channel to Another](#)
- [Copying Packages from One Channel to Another](#)

Prerequisites for Copying Packages from One Channel to Another

Before copying the packages from one channel to another, meet the following prerequisites:

1. Ensure that there are at least 2 channels.
2. Ensure that the target channel machine has adequate space.
3. Ensure that the stage location of the source host does not have a directory named `copyPkgsSrc`, and the stage location of Target Host does not have a directory named `copyPkgsDest`.
4. Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.

For example: `http://h1.example.com`. If this works, then it is confirmed that Apache is installed and listening on port 80.

5. Ensure that the Enterprise Manager User logs in to the OMS with Super User privileges.

Copying Packages from One Channel to Another

To copy the packages from one channel to another, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the source channel, and click **Copy Packages**.
4. Select the target channel.
5. From the source channel section, select and copy the packages to the target channel section.
6. Enter credentials for the source and target channels. These credentials should have read/write access to the machines.
7. Click **OK**.

A Copy Packages job is submitted. Follow the job until it completes successfully.

Adding Custom Packages to a Channel

This section describes how you can add custom packages to a channel. In particular, this section covers the following:

- [Prerequisites for Adding Custom Packages to a Channel](#)
- [Adding Custom Packages to a Channel](#)

Prerequisites for Adding Custom Packages to a Channel

Before you add custom packages to a channel, meet the following prerequisites:

1. Ensure that there is at least one channel.
2. Ensure that the stage location of the source host does not have a directory named `addPkgsSrc`, and the stage location of the destination channel does not have a directory named `addPkgsDest`.

Adding Custom Packages to a Channel

To add custom RPMs to a channel, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the channel name where you want to add the RPM, and click **Add**.
4. Select the source target name and the credentials to be used for the host. The credential you use must have write access on `emd_emstagedir` directory present on the source host.
5. On the Upload Files section, click the search icon to browse for the RPM files.
6. Select a normal host credential that has write access on the select channel.
7. Select a privileged host credential that has write access on the select channel, and has SUDO as root privilege.

8. Click **OK**.

An Add Package job is submitted. Follow the job until it completes successfully.

Deleting a Package Channel

This section describes how you can delete a channel. In particular, this section covers the following:

- [Prerequisites for Deleting a Package Channel](#)
- [Deleting a Package Channel](#)

Prerequisites for Deleting a Package Channel

Before deleting a channel, meet the following prerequisites:

1. Ensure that there is at least one channel.
2. Ensure that the Enterprise Manager User logs in to the OMS with Super User privileges.

Deleting a Package Channel

To delete a channel, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the channel name you want to delete, and click **Delete**.
4. If you want to delete the packages from the RPM Repository machine, select the check box and enter the credentials for the RPM Repository machine. Click **Yes**.
5. If you have not selected to delete the packages from RPM Repository machine, you will get a confirmation message stating *Package Channel <channel name> successfully deleted*. If you have selected the **Delete Packages** option, a job will be submitted to delete the packages from the RPM Repository machine. Follow the job until it completes successfully.

4

Monitoring and Managing Hosts

This chapter includes the following sections:

- [Overview of Host Management](#)
- [Setting Up the Environment to Monitor Hosts](#)
- [Customizing Your Host Monitoring Environment](#)
- [Monitoring Hosts](#)
- [Administering Hosts](#)

Host Consolidation Planner

Host Consolidation Planner enables you to match managed sources you want to consolidate with new or existing destinations. Host Consolidation Planner supports the following combinations:

- Consolidate source servers to generic physical machines, Oracle engineered systems (Exadata Database Machines or Exalogic Elastic Cloud systems), or Oracle Virtual Machine (VM) servers.
- Consolidate source servers to physical machines configured in the Oracle Cloud. In a consolidation of this type, an Oracle Compute Cloud configuration mimics a host, except that only memory and CPU capacity are of consequence as resources to be considered.

Host Consolidation Planner leverages data collected from managed targets by Cloud Control and factors in business and technical requirements to help you determine the optimum scenarios.

Before you begin with Host Consolidation planning, and additional Consolidation information see: Database and Host Consolidations in *Enterprise Manager Cloud Control Database Lifecycle Management Administrator's Guide*.

The following sections describe Host Consolidation Planner:

- [Overview of Host Consolidation Planner](#)
- [Creating a Host Consolidation Planner Project](#)
- [Creating a Host Consolidation Planner Scenario](#)
- [Other Host Consolidation Scenario Creation Options](#)
- [Evaluating Host Consolidation Scenarios](#)

Overview of Host Consolidation Planner

Over the years, the typical enterprise data center will grow steadily due to the addition of servers required to satisfy the needs of the business. This growth typically results in excess servers that occupy rack space, consume a lot of power for cooling, and require system maintenance such as security and patching.

Depending on the procurement cycle or the specific hardware vendor agreement in effect, enterprises may acquire different types of server hardware and operating systems, inadvertently creating a confusing array of systems that administrators need to manage, administer, patch and upgrade. This in turn increases labor and ongoing maintenance and support costs against IT budgets. Enterprises look at consolidation as a way to migrate their disparate systems onto standardized operating systems and hardware, with the ultimate goal of reducing costs.

Enterprises also are increasingly investigating virtualization technologies such as Oracle Virtual Machine by moving from physical to virtual servers. This makes it possible to use the shared hardware infrastructure while getting the benefits of isolation that virtualization provides.

The goal of consolidation is to identify such under-utilized servers and find a way to consolidate them, enabling the enterprise to free up as many servers as possible while continuing to maintain service levels. Since servers have different levels of CPU capacity, Host Consolidation Planner uses computer benchmark data to normalize the CPU usage for a given hardware in the consolidation process. Specifically, Host Consolidation Planner uses the following CPU benchmarks for different classes of hardware:

- SPECint_{base}@2006 for database hosts, application hosts, or mixed-workload hosts
- SPECjbb@2005 for middleware platforms

Host Consolidation Planner enables you to match managed target sources you want to consolidate with new or existing target destinations. Consolidate source servers to generic physical machines, Oracle engineered systems (Exadata Database Machines or Exalogic Elastic Cloud systems), or Oracle Virtual Machine (VM) servers.

You can also consolidate source servers to physical machines configured in the Oracle Cloud. In a consolidation of this type, an Oracle Cloud Compute configuration mimics a host, except that only memory and CPU capacity are of consequence as resources to be considered.

By leveraging metric and configuration data collected from managed target servers by Enterprise Manager Cloud Control, Host Consolidation Planner helps you determine the optimum consolidation scenarios that also reflect business and technical constraints in the consolidation process.

Creating a Host Consolidation Planner Project

You create a host consolidation project for each consolidation effort, then add individual consolidation scenarios within it. You can then compare consolidation scenarios to determine which consolidation strategy makes the most sense.

After the project is defined, a Cloud Control job is submitted to collect available data for the specified targets from the Management Repository. Once the job finishes, the project becomes an active project. As long as the project is in an active state, data collection will continue.

Creating a project involves the following steps:

- [Selecting the Host Consolidation Project Type](#)
- [Specifying Host Source Candidates](#)
- [Specifying Host Destination Candidates](#)

- [Setting Data Collection Parameters](#)
- [Choosing a Pre-configured Scenario](#)
- [Reviewing and Saving the Host Consolidation Project](#)

Selecting the Host Consolidation Project Type

Complete Step 1 of the host consolidation project creation process as follows:

1. From the **Enterprise** menu, select **Consolidation**, then select **Host Consolidation Planner**.
2. Click the **Create Project** button.
3. Enter the consolidation project name and, optionally, a description.
4. Select the host consolidation type:
 - From physical source servers to Oracle VM servers (P2V)
 - from physical source servers to physical servers (P2P)
5. Click **Next** to specify source candidates.

Specifying Host Source Candidates

Complete Step 2 of the host consolidation project creation process as follows:

1. Select an appropriate benchmark from the drop-down menu.
 - Specify SPECint®_base_rate2006 for database hosts, application hosts, or mixed-workload hosts
 - Specify SPECjbb®2005 for middleware platforms
2. Click **Add Source Servers** to select the source servers to be consolidated from a list of managed server candidates. Use the filtering criteria to refine your target search. Choose from the filtered results, then click **Select**.

The sources you select appear in the table. Note that you can subsequently cull the list by removing selected source servers.

3. Optionally set server I/O capacities for disk I/O request and network I/O volume capacities. Click **Specify Server I/O Capacity** to estimate these I/O capacities for all source involved in the consolidation project. Note that you can subsequently fine-tune these estimates by clicking a row and editing the values for the various capacities.
4. Click **Next** to specify destination candidates.

Specifying Host Destination Candidates

Complete Step 3 of the host consolidation project creation process as follows:

1. Optionally select one or more existing servers to consolidate the source servers to.
 - If you are consolidating from physical servers to Oracle Virtual Servers (P2V), click **Add Existing Virtual Servers as Destinations** to view a list of existing VM-based Exalogic Elastic Cloud systems and Oracle Virtual Machine destination servers to consolidate the source servers to. Use the target type filter to differentiate the two. Select the servers you want to add, then click **Select**.

The destinations you select appear in the table. Note that you can subsequently cull the list by removing selected destination servers.

- If you are consolidating from physical servers to physical servers (P2P), click **Add Existing Oracle Engineered System as Destinations** to search for the Exadata Database Machine servers or Exalogic Elastic Cloud servers to consolidate the source servers to. Use the target type filter to differentiate the two. Select the servers you want to add, then click **Select**.

The destinations you select appear in the table. Note that you can subsequently cull the list by removing selected destination servers.

 **Note:**

If you do not select existing destinations, the consolidation will be to new (phantom) destinations. For P2P projects, if you do not select existing destinations, the source servers you select are also considered destination servers. This enables you to select a subset of source servers to consolidate to.

2. Optionally edit estimated CPU capacity, I/O request, and network I/O volume by clicking in a row and changing the value. You can also edit the I/O value by clicking **Specify Server I/O Capacity** to estimate these I/O capacities for all destination servers involved in the consolidation project.
3. Click **Next** to set up data collection.

Setting Data Collection Parameters

Specify the duration over which data used to generate host consolidation scenarios will be collected for the source servers specified in the project. This data will be used to determine the resource requirements that a destination server must meet.

1. Specify the minimum number of days to collect data. The default minimum value is 21 days. To use existing historical data to run and view consolidation scenarios immediately, set the minimum number of days to 0.
2. Specify the maximum number of days to collect data. The default maximum value is 90 days.
3. Specify when to begin the data collection process.
Note that once data collection begins, you can elect to suspend and resume collecting at any time from the **Actions** menu in the Consolidation console.
4. Optionally select **Continue Data Collection Over the Maximum Days** to purge the oldest day's data when data for a new day is added.
5. Click **Next** to choose a pre-configured scenario.

Choosing a Pre-configured Scenario

When creating a host consolidation project, you can optionally choose to generate pre-configured consolidation scenarios to add to the project.

These scenarios are generated using data collected for the sources defined in the consolidation project at the time the project is created. If insufficient data is available

when the project is created, the pre-configured scenarios will be automatically generated once the minimum amount of data has been collected.

Host Consolidation Planner ships with three out-of-the-box scenarios that represent aggressive, medium, and conservative consolidation schemes.

1. Choose whether to use a pre-configured scenario by clicking the appropriate radio button. Choosing the option displays a list of the out-of-the-box consolidation scenarios available. By default, the scenarios are designated by the method used to aggregate source target resource usage:
 - **Aggressive:** Aggregate data based on *average* daily usage per hour.
This typically results in a high consolidation (source:destination) ratio, because more sources will "fit" into each destination. But because more sources are involved, the odds that one or more will not meet the resource requirements are higher.
 - **Medium:** Aggregate data based on the 80 percentile usage.
This typically results in a source:destination ratio somewhere between Aggressive and Conservative aggregations.
 - **Conservative:** Aggregate data based on *maximum* daily usage per hour.
This typically results in a lower source:destination ratio, because fewer sources will "fit" into each destination.

Usage statistics are calculated based on the following criteria:

- **Resource Requirements:** Source requirements, such as CPU, memory (GB) and disk capacity, that must be met by destinations.
 - **Applicable Dates:** The days of the week on which resource usage metrics are collected.
 - **Target Server Utilization Limit:** The maximum resource utilization (percentage) that can be used on destinations. The scenario method in effect determines utilization percentage.
2. Select the pre-configured scenario you want to use. You can select any or all scenarios.
 3. Select whether to factor new (phantom) or existing destinations in the scenario.
 4. Proceed as follows, based on project type and selections:
 - For P2V projects using new servers, specify resource requirements for CPU capacity, memory, and disk storage. Optionally, you can enter quantities of resources to hold in reserve for virtualization software. These reserves are deducted from system capacity before calculating utilization percentages.
 - For P2V projects using existing servers, there is nothing to specify. The scenario method determines resource requirements. Note that this option is available only if existing virtual servers were added as destination candidates in step three of the wizard.
 - For P2P projects using new servers and engineered systems, there is nothing to specify. The system selects an appropriate Exadata configuration based on requirements.
 - For P2P projects using new servers and generic servers, specify resource requirements for CPU capacity, memory, and disk storage. You do not, however, have the option to enter resources to hold in reserve.
 - For P2P projects using existing servers, there is nothing to specify. The scenario method determines resource requirements. Note that this option is available only if

existing virtual servers were added as destination candidates in step three of the wizard.

5. Click **Next** to review the consolidation project.

The pre-configured scenarios will be generated when the project is created using data collected for the source servers defined in the consolidation project.

You can also opt to create your own custom scenario once the consolidation project has been completed.

Reviewing and Saving the Host Consolidation Project

Review the specifics of the host consolidation project. Use the **Back** button to return to a given step to make changes. If satisfied, click **Submit**. A message confirms that the project has been created and the job has been submitted.

Once the project is created, it will show up in the Consolidation console. Consolidation scenarios can then be defined for this project.

Creating a Host Consolidation Planner Scenario

You can create custom host consolidation scenarios instead of or in addition to using the pre-configured scenarios. Multiple scenarios can be created within a project, enabling you to compare different scenarios before deciding on a solution. New consolidation scenarios are created within an existing consolidation project. You can create host consolidation scenarios for planning purposes only.

Creating a scenario involves the following steps:

- [Setting Up the Scenario](#)
- [Defining Constraints for Host Consolidation](#)
- [Planning the Destination for a Physical Server to Virtual Server Project](#)
- [Planning the Destination for a Physical Server to Physical Server Project](#)
- [Mapping Host Sources to Destinations](#)
- [Reviewing and Saving the Host Consolidation Scenario](#)
-

Setting Up the Scenario

Complete Step 1 of the host consolidation scenario creation process as follows:

1. From the **Enterprise** menu, select **Consolidation**, then select **Host Consolidation Planner**.
2. Click the consolidation project for which the scenario is intended.
3. Click the **Create Scenario** button.
4. Specify the scenario details, such as scenario name.
5. Specify the source resources to consider. Host Consolidation Planner will aggregate the specified resources to determine the total requirements.
 - Resource Type: The server requirements, such as CPU, memory (GB), and I/O capacity that must be considered.

- **Scale Factor:** Provide room for growth on the destination for each source. The resource requirement estimate uses the scale factor to pad the estimate for consolidation. So, for example, if the estimated requirement for a given source, based on usage data, is 2 GB of memory, which equates to a scale factor of 1, and you specify a scale factor of 1.1, 2.2 GB will be required to consolidate that source.
 - **Applicable Days:** The days of the week on which resource usage metrics are collected.
 - **Resource Allocation:** The method used to aggregate daily server source resource usage. Values are:
 - **Aggressive:** Aggregate data based on *average* daily usage per hour.
This typically results in a high consolidation (source:destination) ratio, because more sources will “fit” into each destination. But because more sources are involved, the odds that one or more will not meet the resource requirements are higher.
 - **Medium:** Aggregate data based on the 80 percentile usage. This typically results in a source:destination ratio somewhere between Aggressive and Conservative aggregations.
 - **Conservative:** Aggregate data based on *maximum* daily usage per hour.
This typically results in a lower source:destination ratio, because fewer sources will “fit” into each destination.
 - The date ranges should define a period of time that is typical of standard resource requirements.
6. Click **Estimated Requirements** to view the estimated total resource requirements.
For server consolidations, resource requirements are shown based on the averaged hourly requirement. The displayed requirements reflect the scale factor (if any) specified for the resource. The 24-hour requirement pattern will be used as the minimum requirements that must be met by consolidation destinations.
 7. Optionally exclude or include sources, as appropriate.
 8. Click **Next** to define consolidation constraints.

Defining Constraints for Host Consolidation

Specify business, corporate, or technical constraints that must be enforced. Constraints can guide the allocation process during automatic source-to-destination mapping. For manual mappings between sources and destinations, constraints can calculate violations.

Complete Step 2 of a host consolidation scenario as follows:

1. Select compatible server criteria.
Servers are considered compatible if they have the same specified target properties and configurations. If you are consolidating multiple source servers to a single destination server, only compatible servers can be consolidated together on the same destination server.
2. Specify mutually exclusive server criteria.
Certain types of source servers are mutually exclusive and should not be consolidated together on the same destination server due to various reasons. For example, nodes of an Oracle cluster should not be placed in the same failure group.

3. Click **Preview Effect of Constraints** to view a list of source servers that are not compatible based on the defined constraints.
4. Click **Next** to specify destination server candidates.

Planning the Destination for a Physical Server to Virtual Server Project

For a P2V project:

1. Choose destination server candidates using either of the following options:
 - Click **Use New (Phantom) Servers** if you plan to use destination servers that have yet to be provisioned or purchased, then select either of the following options:
 - **Use Oracle Engineered System** and click the search icon to select an appropriate configuration type for an Exalogic Elastic Cloud system.
 - **Use Generic Servers** and provide the estimated CPU capacity if available; otherwise, click the search icon next to the CPU capacity input field, then select a server configuration that most closely matches your needs.

Adjust the memory and storage estimates as necessary.

In a virtual environment, you can also specify a quantity of the resource to be set aside (reserved) for use by supervisory software in the database machine. This quantity is subtracted from the total capacity of the destination before consolidating source servers into the remaining resource. For example, if your estimated memory requirement is 12 GB and you specify a reserve of 2 GB, only 10 GB is available for consolidation.

Host Consolidation Planner will determine how many destination servers are required as part of the consolidation results.

- Click **Use Existing Servers** to specify a set of existing managed servers to use as destinations.

These are the servers you specified when defining the scope for the consolidation project. Host Consolidation Planner will determine the available hardware resources based on collected usage data.

By default, the consolidation process will try to use as few destination servers as possible. If you prefer, choose to balance the source load across all destinations.
- 2. Accept the defaults or edit the percentages for **Maximum Allowed Resource Utilization on Destination Servers**. Contrast these allowances, which provide headroom on destination servers, with the scale factor, which provides headroom for individual source servers.
- 3. Click **Next** to map the source servers to the destination servers.

Planning the Destination for a Physical Server to Physical Server Project

For a P2P project:

1. Choose destination server candidates using either of the following options:

- Click **Use New (Phantom) Servers** if you plan to use destination servers that have yet to be provisioned or purchased, then select one of the following options:
 - **Use Oracle Engineered System** and select either Exadata Database Machine or Exalogic Elastic Cloud. Click the search icon and select a configuration type appropriate to either choice.
 - **Use Generic Servers** and provide the estimated CPU capacity if available; otherwise, click the search icon next to the CPU capacity input field, then select a server configuration that most closely matches your needs.
 - **Use Oracle Compute Cloud** and select the cloud computing configuration, or shape, to use as the destination.
- Click **Use Existing Servers** to specify a set of existing managed servers to use as destinations for the project.

These are the servers you specified when defining the scope for the consolidation project. Host Consolidation Planner will determine the available hardware resources based on collected usage data. If you did not explicitly specify destination candidates, all source servers are potential destinations for consolidation.

By default, the consolidation process will try to use as few destination servers as possible. If you prefer, choose to balance the source load across all destinations.

2. Accept the defaults or edit the percentages for **Maximum Allowed Resource Utilization on Destination Servers**. Contrast these allowances, which provide headroom on destination servers, with the scale factor, which provides headroom for individual source servers.
3. Click **Next** to map the source servers to the destination servers.

Mapping Host Sources to Destinations

Next, map the host sources to the destinations they will be consolidated to. The objective is to fit source requirements with each destination's available resources as tightly as possible.

Oracle recommends that you allow Host Consolidation Planner to perform this mapping automatically. This will allow the tool to maximize resource utilization of destinations based on resource capabilities and the various consolidation constraints specified. If you use manual mapping, the source will be mapped to the destination even if the destination lacks sufficient capacity. In addition, manual mapping may violate previously declared constraints.

When you have chosen existing destinations, you can optionally map each source and destination manually:

1. Click a source in the list.
2. Click the flashlight icon to select the destination to map to the source. Note that you can map a single source to a destination or multiple sources to a destination, but there can be only one destination

The resulting consolidation report will show any resource and/or constraint violations due to such manual mapping.

3. Click **Next** to review the consolidation scenario.

Reviewing and Saving the Host Consolidation Scenario

Finally, review the various parameters set in the new host scenario. Use the **Back** button if you need to make changes; otherwise, proceed as follows:

- Optionally, you can save the scenario as a template, which can then be shared with other users. If you want to do this, click **Save Scenario as a Template**. In the dialog that opens, browse to a location in the local file system and save the template as an XML file.
- Click **Submit**. A message appears confirming that a job has been submitted for further analysis of the scenario. Results appear at the bottom of the Consolidation home page when done.

Other Host Consolidation Scenario Creation Options

You can create a host consolidation scenario based on an existing scenario. Select an applicable scenario under a consolidation project and then select **Create Like Scenario** from the **Actions** menu. Modify the scenario as desired, provide a meaningful name, and submit for analysis as usual.

If you saved a scenario as a template, you can subsequently import the scenario into another environment.

1. On the Host Consolidation Planner home page, select **Create Scenario from Template** from the **Actions** menu.
2. Browse to a saved template XML file in the local file system. Click **Open**.
3. Indicate the extent to which you want to replicate the saved template; that is, in terms of the resources, constraints, and destinations planning represented by the template. Click **Update** if you make any changes.
4. Click **OK** to import the saved template.

The imported scenario opens in the wizard where you can edit and save it in Host Consolidation Planner.

Evaluating Host Consolidation Scenarios

You can view details for your consolidation scenarios using the Consolidation console. After evaluating the consolidation scenario results, you can define different plans as well as rerun existing scenarios to re-evaluate them based on the previously specified conditions with the latest available data. The results of the previous analysis will be over-written. You can also create a new scenario based on an existing scenario, where you tweak certain values to customize the new scenario. This iterative process helps you obtain the optimized consolidation scenario which is generated by compromising various factors and weighing different trade-offs.

Compare the consolidation scenarios you create to determine which consolidation strategy best meets your requirements.

Your objective is to:

- Match source resource requirements with destinations best able to meet those requirements.
- Fit source requirements with each destination's available resources as tightly as possible, so you can get maximum usage of destination capacity.
- Provide room for growth on destinations by allowing for headroom as a factor of resource requirements.
- Optionally balance the source workload across all available destinations.

1. From the **Enterprise** menu, select **Consolidation**, then select **Host Consolidation Planner**.
2. First, examine the project containing the scenario you want to view.
 - The **Status** column indicates the status of data collection, based on the minimum and maximum collection days specified for the project.
 - The **General** tab summarizes the project in terms of type, collection details, number of sources, and so forth.
 - Click the **Source Candidates** tab to view various usage data collected for the sources defined in the project. Data can include utilization rates for CPU, memory, storage, and disk and network I/O, depending on the project type.
 - Click the **Source Workload** tab to view source resource usage data collected. The default display is a heat map, a grid of 24 hours by 30 days, showing the workload for a given hour as a color indicating the load relative to 100% of capacity, and a number showing the actual percentage. Alternatively, you can view the same data as a line graph. You can filter either view by source, resource type, and month.
 - Click the **Destination Candidates** tab to view a breakdown of hardware details and projected resource utilization by destination candidate, based on the sources to be consolidated.
 - Click the **Report** button above the table when the project is selected to view summarized information and more details.
3. Next, view the data for a specific scenario. The **General** tab summarizes the scenario in terms of resource type and allocation, constraints, destination types, and so forth. For a completed analysis, click any metric in the row to view details on the respective tab, as follows:
 - **Sources:** The list of sources to consolidate, including their projected CPU and memory capacities and requirements.
 - **Destinations:** The list of destinations to which the sources will be consolidated. Resource configuration and calculated utilization are shown for each destination.
For consolidations to the cloud, resources of consequence are CPU capacity and memory.
 - **Ratio:** The ratio of sources to destinations. By default, Host Consolidation Planner will try to "fit" sources into as few destinations as possible.
 - **Mapping:** The destinations to which specific sources will be mapped. The analysis includes estimated CPU and memory requirements and utilization, enhanced by suggested CPU and memory allocation figures to consider. These suggestions represent a reasonable compromise between requirements and destination server capacity.
 - **Confidence:** The percentage of the data collected for sources that meet the source usage requirements defined in the scenario. This value is aggregated for all sources defined with the project.
 - **Violations:** The number of violations of technical or business constraints defined in the scenario. This metric is applicable only if the scenario uses auto-mapping of sources to destinations.
 - **Exclusions:** The number of sources that do not have a qualified mapping to a destination. These are sources that exceed the capacity of available destinations.

A different set of constraints may result in a different optimal scenario. Modify the constraints to come up with different scenario results.

Overview of Host Management

A host is a computer where managed databases and other services reside. A host is one of many components or targets that can be monitored and managed by Oracle Enterprise Manager.

Monitoring refers to the process of gathering information and keeping track of activity, status, performance, and health of targets managed by Cloud Control on your host. A Management Agent deployed on the host in conjunction with plug-ins monitors every managed target on the host. Once hosts are discovered and promoted within Enterprise Manager, you can monitor these hosts.

Administration is the process of managing and maintaining the hosts on your system.

To view all the hosts monitored by Oracle Enterprise Manager, select **Hosts** on the **Targets** menu of the Enterprise Manager Cloud Control.

Note:

For information on discovering and promoting hosts, discovering unmanaged hosts, converting unmanaged hosts to managed hosts, and so on, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Host Statistics

The host management capabilities in Enterprise Manager provide a quick glimpse of all the hosts on your system. This includes lifecycle status and configuration changes.

Using the host UI, you can:

- Determine whether a particular host is available and whether there are incidents and problems associated with that host.
- View statistics (metrics) applicable to each host. You have over 40 metrics to choose from! Examples of metrics include CPU, memory utilization, file system, and network statistics. See the *Oracle Enterprise Manager Framework, Host, and Services Metric Reference Manual* for details about each of the host metrics.
- Add and configure individual hosts.
- Perform administrative operations on multiple hosts within the context of Enterprise Manager. This is possible by running the Host Command which enables you to type operating system commands against multiple hosts and immediately view the results.
- Analyze job activity statistics including problem job executions, suspended job executions, and running jobs.
- Analyze compliance summary information to established standards. This enables you to determine what the issues are for the host and then correct the compliance violations as soon as possible.

Diagnosing Host Problems

To diagnose a host problem, consider performing the following steps:

- Investigate the incidents and problems reported for the host.
- Determine whether the statistics reported for CPU utilization, memory utilization, file system usage, and network utilization are within acceptable levels for different periods.
- Ensure the host is compliant with the established compliance standard.
- Investigate problematic job executions and why jobs are suspended.

Viewing Targets on the Host

Enterprise Manager allows you to view summary information about the targets on the host target. You can quickly determine how the individual targets are performing by analyzing the incidents and availability information. This gives you the opportunity to make changes as needed so the targets will function at peak performance.

Examples of targets that can reside on a host are: Database Instance, Web Cache, and Oracle HTTP Server.

When working with targets on a host:

- Study incident information. The message associated with a particular incident provides a detailed description of what is wrong with the target.
- Determine whether there are any compliance violations against this target.
- If needed, remove multiple targets from the host.

This function is particularly useful when you want to eliminate, from the Management Repository and the Management Agent, those targets that no longer need to be monitored. This need can occur when a monitored target is deinstalled from the computer, or the Management Agent or host is no longer in service.

When removing multiple targets:

- Ensure that the Management Agent is up when you are removing a target. If the Management Agent is down when the target is deleted, the target will be removed from the Management Repository only and not from the Management Agent. Therefore when the Management Agent is brought back up, the target will be back again.
- Be aware that the Management Agent cannot be deleted unless it is the only target remaining on the host.

Storage Statistics and History

Tracking the storage resource allocation and usage is essential to large Information Technology departments. Unallocated and under utilized storage can be put to better use. Historical trends at a business entity level enable you to plan for future growth.

By default the storage history feature is not activated. Enabling storage history is expensive in regards to database resources. The amount of database resources used to calculate history data depends on the amount of storage data associated with the host target.

Setting Up the Environment to Monitor Hosts

Before you start monitoring and administering hosts, it is recommended that you set up credentials and install the needed software. This chapter describes:

- [Required Installations](#)

- [For Linux Hosts - Installing YAST](#)
- [Setting Up Credentials](#)
- [Setup Needed for Host Monitoring](#)
- [Target Setup Needed for Host Administration](#)



Note:

The installation of YaST is only for Oracle Linux 6.0 or below. For Oracle Linux 7.0 operating system, no other installation is needed.

Required Installations

These required installations are only applicable to hosts running Oracle Linux, Red Hat Linux, and SUSE Linux Operating Systems (x86 and x64 architectures only). To administer a host through Enterprise Manager, you need to install scripts. To determine which scripts you need to install for your host, follow these steps:

1. From the **Targets** menu, select either **All Targets** or **Hosts**.
2. Either type the name of the desired host in the Search field or scroll down to the name of the host in the Name column.
3. Click the name of the host.
4. From the **Host** menu, select **Administration**, then select **Services**.
5. Oracle Enterprise Manager now supports Oracle Linux host administration features without YAST dependency. Host agent must be running on latest Enterprise Manager Agent Bundle Patch 13.2.0.0.0 or later.



Note:

The credential user must have access to the Agent installation.

These scripts can be downloaded from <http://oss.oracle.com/projects/yast>.

For Linux Hosts - Installing YAST

YAST is an operating system setup and configuration tool that comes as a standard tool as part of SUSE Linux distribution. The Linux administration feature uses YAST to run scripts, installing YAST RPM from the following location will also install the Enterprise Manager scripts:

<http://oss.oracle.com/projects/yast>

For SUSE, you need to download the Enterprise Manager scripts and additional remote access module from the following location:

<http://oss.oracle.com/projects/yast/files/sles9>

Before you install YAST, you need to determine the following:

1. Determine the version of Linux on your machine. For example, the `uname -a` command lists the RHEL (RedHat), Oracle Linux, or SUSE versions, and if it is on a 32 or 64 bit environment.
2. Verify that you have root privileges.

To install YAST, perform the following steps:

1. Go to <http://oss.oracle.com/projects/yast>.
2. Click the **here** link. The Project Downloads: YAST page appears. Click the link that matches your Linux version.
3. Click the link associated with your machine, either *i386* for 32 bits or *x86-64* for 64 bits.
4. Click **yast_el5_x86_64.tar** to download the tar file.
5. Once the tar is downloaded, go to the directory where the tar file was downloaded.
6. Extrac the file using `tar -xvf yast_el5_x86_64.tar`
7. `cd` to the `yast_el5_x86_64` directory.
8. Type `sudo ./install.sh`
9. To verify that YAST is installed, type: `/sbin/yast2`. This should display the YAST control center. If it does not, the YAST installation has failed.
10. When you return to the Administration menu, the options should now display the available Linux administration features.

For a demonstration of how to install YAST, see the YouTube video located at: <http://www.youtube.com/watch?v=7ZiwmxZVmAw>.

Setting Up Credentials

Credentials are needed to manage target instances.

To set up various credentials, select the **Setup** menu (located at the top-right of the UI page), then select **Security**. The following options are available:

- **Named Credentials** are used for the Management Agent install. Named credentials explicitly grant you privileges on the host.
- **Preferred Credentials**
If a target has preferred credentials set, applications that log in to that target will automatically use the preferred credentials. Using preferred credentials simplifies access to managed targets.
Default credentials can be set for each target type. Default credentials are used for any targets that do not have preferred credentials explicitly set.
- **Privilege Delegation Setting** enables you to configure the Management Agent to use Sudo or PowerBroker so you can run privileged scripts.

See the online help for additional information.

Setup Needed for Host Monitoring

As you begin monitoring a host, you need to know what metrics you are allowed to monitor. You may also find that you need to set up monitoring credentials for target instances.

This section explains the required steps for these tasks.

Viewing Monitoring Configuration

The Monitoring Configuration page reports what monitoring you can do on the selected host.

To access the Monitoring Configuration page, perform the following steps:

1. From the **Targets** menu, select either **All Targets** or **Hosts**.
2. Either type the name of the desired host in the Search field or scroll down to the name of the host in the Name column.
3. Click the name of the host.
4. From the **Host** menu, select **Target Setup**, then select **Monitoring Configuration**.
5. The Monitoring Configuration page appears. Details can include, for example, Disk Activity Metrics Collection Max Rows Upload.

In addition, the Monitoring status is provided. For example, Oracle has automatically enabled monitoring for this target's availability and performance, so no further monitoring configuration is necessary.

Setting Up Monitoring Credentials

Monitoring Credentials allow you to monitor and access various target functionality. You can manage the already existing credentials for various target types using monitoring credentials.

To edit Monitoring Credentials, perform the following steps:

1. On the Enterprise Manager page, locate the **Setup** menu located at the top right of the page.
2. From the **Setup** menu, select **Security**, then select **Monitoring Credentials**.
3. On the **Monitoring Credentials** page, select **Host** and click **Manage Monitoring Credentials**.
4. On the **Host Monitoring Credentials** page, select a row and click **Set Credentials** to edit the credentials.

By default, a host has the following credential sets defined:

- Host Credentials For Real-time Configuration Change Monitoring
- Host SNMP Credentials
- Host WBEM Credentials
- Privileged Host Monitoring Credentials

You can add credential sets using the emcli `create_credential_set` verb with the `-monitoring` option.

Target Setup Needed for Host Administration

Before you start administrating the host, you need administrator access.

From the **Host** menu on the Host home page, select **Target Setup**, then select **Administrator Access**. Using this option enables you to determine target privileges for a user.

Monitoring Hosts

As a host administrator, you must have a grasp of how your host is functioning. Host monitoring can enable you to answer such questions as:

- Is host swapping occurring?
- Is the filesystem becoming full?
- Is the CPU reaching maximum capacity?
- Are the resources being used efficiently?
- What is the best way to monitor multiple hosts?
- How can I proactively schedule and purchase needed resources?

The answers to these questions are key for day to day monitoring activities and are available on the host monitoring pages explained in this chapter.

Note: This chapter explains many of the metrics available in Enterprise Manager, however it is not an exhaustive list. See the [Oracle® Enterprise Manager Framework, Host, and Services Metric Reference Manual](#) for a full description of all the host metrics available.

Overall Monitoring

When monitoring your host, the prime metrics to monitor are CPU, memory, and disk usage.

Note: To access the features explained in this section, from the **Host** menu on a host's home page, select **Monitoring**, and select the feature of interest.

CPU Details

Using the CPU statistics, you can determine whether CPU resources need to be added or redistributed. In particular, you can:

- Determine the commands that are taking the most CPU resources and perform the appropriate action on the target host to reduce contention by using an administrative tool of your choice.
- View trends in CPU Usage over various time periods including last 24 hours, last week and last month.
- Monitor all CPUs, that is, not an aggregate view but a view of all the CPUs in the system.

Note: You can use the Execute Host Command feature in Enterprise Manager to perform actions on the host.

Memory Details

Using the Memory statistics, you can determine whether memory resources need to be added or redistributed. In particular, you can determine the processes that are using the most memory resources.

Disk Details

Using the Disk statistics, you can determine whether Disk resources need to be added or whether you can distribute the load more effectively across existing resources. In particular, you can determine the disks that are over utilized or experiencing longer service times.

Correlating the disk information with the response from applications that use the underlying storage allows you to determine whether the system is properly scaled. You can then answer the questions: Should the load on the disks be redistributed? Should additional storage be added?

To redistribute the load, modify the applications that use the storage.

Program Resource Utilization

Using the Program Resource Utilization data, you can see the trends in resource usage for:

- Specific program or set of programs
- Special user or set of users
- Combination of programs and users

Log File Alerts

Enterprise Manager monitors log files and provides alerts. Once alerts are generated, you can:

- Clear open alerts selectively or clear every open alert.
- Purge open alerts selectively or purge every open alert.

Note: Clearing an alert results in the particular alert being marked as cleared but the alert is *not* deleted from the Management Repository. However, purging an alert permanently deletes the alert from the system.

Metric Collection Errors

Metric Collection Errors provide details about the errors encountered while obtaining target metrics. These details give you an idea of the metrics that may not represent the performance of the target accurately, as errors were encountered while collecting the metrics.

Storage Details

Tracking the storage resource allocation and usage is essential to large Information Technology departments. Unallocated and under utilized storage can be put to better use. Historical trends at a business entity level enable you to plan for future growth.

Storage Details are relevant to Enterprise Manager targets that are associated with one or more hosts. In particular:

- Summary attributes presented are rolled up for one or multiple associated hosts.
- A host is associated with a group either through:

- Explicit membership, or
- Implicit *hosted by* association which is inherited through a group member target

Note: The shared storage is accurately counted once when the storage is accessible from multiple systems or accessible through multiple physical paths on the same system. Globally unique identifiers have been instrumented for accurate counting of shared storage.

Refer to the online help to learn how the individual storage statistics are calculated.

Storage Utilization

Storage utilization is provided at the host level when launched in the context of a host target and associated hosts level when launched in the context of a group.

In the context of a host, the storage items are: Disks, Volumes, ASM (Automatic Storage Management), File Systems, Databases, and Vendor Distribution.

In the context of a group, the storage properties for the associated hosts are: Provisioning Summary by Host, Consumption Summary by Host, and Vendor Distribution

The graphs present historical trends over a period of time. Based on this intelligence, you can take appropriate action on the target host or group as necessary. Appropriate actions include:

- Buying and adding more storage
- Deleting underutilized application data after archiving
- Deleting unneeded application data
- Altering the storage deployment configuration for optimal use

Note: The storage information shown in a group is the aggregate of the individual host information of the associated hosts in the group.

Overall Utilization

Overall Utilization represents summary attributes (unallocated, overhead, used, and free) that provide a system level view of storage resource utilization. The overall statistics enable you to determine:

- How much storage is unallocated?
- How much space is still free among deployed applications?

Provisioning Summary

Provisioning Summary represents allocation related summary attributes (allocated, unallocated, and overhead) for File Systems (Writeable NFS part), ASM, Volumes, and Disks for the associated hosts.

Note that Writeable NFS is shown in Provisioning Summary to account for the storage attached to the host over NFS. These layers are managed by IT administrators who are responsible for provisioning space to applications.

Allocation related attributes do not change frequently, change typically results from an administrative action taken by an IT administrator. See Provisioning Summary section in the About Storage Computation Formula help topic for details on how this information is calculated.

The bar chart summarizes the allocated, unallocated, and overhead for all entities present in Disk, Volume, Oracle ASM, and Writeable Network File Systems (NFS) portion of File System layer for the host or associated hosts of the group.

If a specific layer is not deployed, the corresponding bar is omitted from the chart. The bar chart answers the following questions.

- How much space is available for allocation from the entities present in the given layer?
- How much space was allocated from the entities present in the given layer?
- What is the overhead of deployed Volume Management software?
- What is the overhead of deployed Oracle ASM software?

Note: When launched in the context of a group, rollup information shown in the charts excludes NFS mounts that are based on Local File Systems present in the associated hosts.

Consumption Summary

Consumption Summary provides usage related summary attributes (used and free) for Databases, File Systems (Local File Systems and Writeable NFS parts).

Usage related attribute values tend to change more frequently relative to allocation related attributes. See Consumption Summary section in the About Storage Computation Formula help topic for details on how this information is calculated.

The bar chart shows used and free space summary information for all Databases, all Local File Systems, and all Writeable Network File Systems (NFS) in the host or the associated hosts of the group.

Note: When launched in the context of a group, rollup information shown in the charts excludes NFS mounts that are based on Local File Systems present in the associated hosts.

ASM

Oracle Automatic Storage Management (ASM) is a simple storage management solution that obviates the need for using volumes layer technologies for Oracle databases.

Databases

Databases refer to Oracle databases (including Real Application Cluster (RAC) databases) on top of which other applications may be running. Databases can consume space from disks, volumes, file systems, and Oracle Automatic Storage Management (ASM) layers.

Disks

Disks statistics provide the allocated and unallocated storage for all the disks and disk partitions on a host. All disks are listed including virtual disks from external storage systems such as EMC Storage Array.

Note: Overhead information for virtual disks is not instrumented nor presented.

For a disk to be deployed for usage, the disk must first be formatted. After formatting, the disk can be configured (using vendor-specific configuration utilities) to have one or more partitions.

A disk or disk partition can be associated (using vendor-specific configuration utilities) with exactly one entity from one of the upper layers (volumes, Oracle ASM, databases, and file systems) on the host. When an association exists for a disk or disk partition to an upper layer entity, it is reported as allocated space in Enterprise Manager.

File Systems

File Systems Layer contains directories (also known as folders) and files that are accessed, managed, and updated through the use of databases, middle tier applications, and end-user tools. They can be broadly categorized into local file systems that are disk based and remote file systems like NFS. In Enterprise Manager, summary attributes are provided for local file systems and the Writeable NFS part of File Systems layer.

Local File Systems

Local File Systems are based on disk storage visible to the host. Various operating systems support different types of local file systems. The following table provides examples:

Local File System	Operating System
lofs	Solaris (Monitored only if NMUPM_SUPPORT_LOFS property is set to 1 for the target instance.)
nfs	Solaris, Linux
tmpfs	Solaris
ufs	Solaris, Linux, AIX, HP
vxfs	Solaris, Linux, AIX, HP
zfs	Solaris, Linux, AIX
ext2	Linux, AIX
ext3	Linux, AIX

NFS

Network File Systems (NFS) are accessible over the network from the host. A remote server (NFS Server) performs the I/O to the actual disks. There are appliances that provide dedicated NFS Server functionality, such as Network Appliance Filer. There are also host systems, for example, Solaris and Linux, that can act as both NFS Server and Client.

Writeable NFS refers to the NFS mounted on a host with *write* privilege.

Suggestions for Monitoring NFS Mounts

The following are suggestions on monitoring NFS mounts.

- Monitor the remote host if NFS exports are coming from another host supported by Enterprise Manager. The Filesystems metric will monitor the local file systems on the remote host.
- Monitor the Netapp Filer if NFS exports are coming from a remote Netapp Filer. Volumes and Qtree metrics will monitor the exports from the remote Netapp Filer.

- Use the 'File and Directory Monitoring' metric if any of the previous choices do not meet the need. Set the threshold against the 'File or Directory Size' metric to monitor specific remote mounts.

Volumes

Various software packages are available in the industry that are either generically known as Volume Manager technology or Software*RAID (Redundant Arrays of Independent Disks) technology. These technologies are deployed to improve the RAS (Reliability, Availability, and Scalability) characteristics of the underlying storage. For example, Veritas Volume Manager is a popular product used across multiple operating systems. Such technologies are referred to as Volumes in Enterprise Manager.

The Volumes option displays the allocated and unallocated storage space for all the entities present in the Volumes layer, including relevant attributes for the underlying Volumes layer technology.

Types of Entities

The Volumes layer can have entities of various types present internally. Entity type shown in Enterprise Manager is based on the terminology as defined by the deployed Volumes layer technology. For example, a Veritas volume manager defines and supports the following entity types: Volume, Plex, Sub Disk, VM Disk, VM Spare Disk, and Diskgroup. Refer to the vendor documentation for more details about the Volumes technology deployed on your system.

Top-Level Entities

Top-level Volumes layer entities provide storage space to the upper layers for usage. If a top-level entity does not have an association to an entity from an upper layer, the top-level entity is unallocated and it is available for further allocation related activity.

For each vendor technology, entities of specific types from their layer can be associated with entities from the upper layers. File Systems, Databases, and ASM are examples of upper layers. For example, entities of type 'Volume' in Veritas Volume Manager are such entities. These entities are referred to as top-level Volumes layer entities in this documentation.

Bottom-Level Entities

For each vendor technology, entities of specific types from their layer can be associated with entities from the disk layer. For example, VM Disk and VM Spare Disk entities in Veritas Volume Manager are such entities. These entities are considered to be bottom-level Volumes layer entities in this documentation.

Bottom-level Volumes layer entities consume storage space from the disk layer and provide storage space to the rest of the entities in the Volumes layer. Bottom-level entities of 'reserve' or 'spare' type are always allocated and no space is available from them for allocation purposes. Note that spare entities are utilized by the Volumes technology for handling disk failures and they are not allocated to other entities present in the Volumes layer by way of administrator operations.

Non-spare bottom-level entities can have an association to an intermediate or top-level entity configured using respective vendor administration utilities. If no association exists for a non-spare bottom-level entity, then it is unallocated. If one or more associations exist for the non-spare bottom-level entity, then the space consumed through the existing associations is allocated. It is possible that some space could be left in the bottom-level entity even if it has some associations defined for it.

Storage space in non-spare bottom-level entities not associated with intermediate or top-level entities is available for allocation and it is accounted as unallocated space in the bottom-level entity.

Intermediate Entities

Non top-level and bottom-level entities are considered to be intermediate level entities of the Volumes layer. For example, Volume (layered-volume case), Plex and Sub Disk entities in Veritas Volume Manager are such entities.

If an intermediate entity has association to another intermediate or top-level entity, the storage space consumed through the association is allocated. Space present in the intermediate entity that is not consumed through an association is unallocated.

The following vendor products are instrumented:

Platform	Product
Solaris	Solaris Volume Manager
Linux	mdadm, raidtool, Suse LVM

Vendor Distribution

The Vendor Distribution statistic reflects the host-visible storage for associated hosts, that is:

```
Sum of the size of all disks
+ Sum of the size of all Writeable NFS mounts
```

Storage History

Enterprise Manager provides historical trends for its storage statistics. Historical trends can be viewed over last month, last three months, or last year. Using this historical trend, you can predict how much storage your organization may need in the future.

In the case of a group, history is not enabled by default. The user interface allows you to enable or disable the history for each group. Computation of history for a group consumes resources in the Enterprise Manager Repository database. It is not anticipated that a given deployment would find it useful to have the history for all instances of groups, so the control is given to you to choose for which groups it is worth keeping the history.

Storage Layers

The stack of storage management technologies is deployed on a host. Deployed technology at any layer can provide storage resources to any layer above it and consume the storage resources from any layer below it.

The ultimate consumer of the storage is application level software such as an Oracle database or the end users. In Enterprise Manager, Volumes refers to Volume Management and Software*RAID (Redundant Arrays of Independent Disks) technologies offered by various vendors.

In Enterprise Manager, the following storage layers and their associations have been modeled.

Storage Layer	Can Provide Storage To:
Disks	Volumes, File Systems, Database, ASM
Volumes	File Systems, Database, ASM
ASM	Database
File Systems	Database

Storage Refresh

Storage Refresh is performed in the context of two types of targets: host target and group target.

Storage Refresh in Context of Host Target

Storage Refresh functionality, in the context of a host target, allows you to refresh the storage data in your Enterprise Manager repository by:

- Forcing Enterprise Manager to perform a real-time collection of all storage attributes from the host, and
- Uploading the storage attributes into the Enterprise Manager repository

Once the refresh operation is complete, the Storage UI pages display the latest information about the host.

Storage Refresh in Context of Group Target

Storage Refresh functionality, in the context of a group target, allows you to refresh the storage data in your Enterprise Manager repository by:

- Forcing Enterprise Manager to do a real-time collection of all storage attributes from all the member hosts of the group, and
- Uploading the storage attributes into the Enterprise Manager repository

Since this refresh could take some time, depending on the number of hosts involved, the functionality is provided as an Enterprise Manager job submission.

Once the refresh job is complete, the Storage UI pages display the latest information about the group.

Ksplice for Oracle Linux

Ksplice updates the Linux operating system (OS) kernel and key user space libraries, while the OS is running, without a reboot or any interruption. For more information, see [Oracle Ksplice](#). All the Oracle Linux Hosts where Enterprise Manager Agent is installed and Ksplice software is configured can be monitored and managed through the Oracle Linux Home Ksplice region. For more information, see [Oracle Linux Ksplice User's Guide](#)

For more detailed reporting and use of these metrics on Oracle Linux Systems see, [Oracle Linux Home](#)

Customizing Your Host Monitoring Environment

To facilitate your use of host monitoring, Enterprise Manager enables you to customize your host pages and environment. This chapter describes:

- [Customizing the Host Home Page](#)
- [Using Groups](#)

Customizing the Host Home Page

By default, the Host home page displays the following regions:

- Summary
- Configuration
- Job Activity
- CPU and Memory
- FileSystem and Network
- Incidents and Problems
- Compliance Standard Summary

These regions are displayed in the two column format with the left column being narrower than the right column.

For many customers, these regions and column formats meet their needs. However, additional or different regions may be needed.

To customize the Host Home page, click the Personalize Page icon located next to the Page Refreshed text at the top right of the page.

You can also add or remove regions. Regions you can add include:

- Compliance Summary
- Configuration Details
- CPU and Memory Performance charts for host targets
- File system and Network Performance charts for host targets
- Incident List
- Job Activity Region
- Job Summary Region
- Performance Metric Chart

In addition, you can change the layout of the page to one column, two equally-sized columns, two columns above a wide area, and so on.

Using Groups

Groups are an efficient way to logically organize, manage, and monitor the targets in your global environments. Each group has its own group home page. The group home page shows the most important information for the group and enables you to drill down for more information. The home page shows the overall status of the group and other information such as current availability, incidents, and patch recommendations for members of the group.

The Managing Groups chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide* explains the different types of groups, as well as how to manage, edit, and view groups.

To create a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**.
2. Click **Create**. Select **Group**, **Dynamic Group**, or **Administration Group**.

The Enterprise Manager Console displays a set of Create Group pages that function similarly to a wizard. See the online help for a description of the fields on each page.
3. On the General tab of the Create Group page or the Create Dynamic Group page, enter the **Name** of the Group you want to create. If you want to make this a privilege propagating group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target privileges granted on the group to an administrator (or a role) are propagated to the member targets.

Note: The Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner *must* be able to view any target that can potentially match the membership criteria.
4. Configure each page, then click **OK**. You should configure all the pages before clicking **OK**.

After you create the group, you always have immediate access to it from the Groups page.

You can edit a group to change the targets that comprise the group, or change the metrics that you want to use to summarize a given target type. To edit a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**.
2. Click the group **Name** for the group you want to edit.
3. From the **Group** menu, click **Target Setup**, then choose **Edit Group**.
4. Change the configuration for a page or pages, then click **OK**.

Alternatively, you can select the group you want to edit from the list of groups on the Groups page and click **Edit** from the top of the groups table.

Administering Hosts

As you monitor your host, you will find that the host needs to be fine-tuned to perform at optimum levels. This chapter explains how to administer your host to reap the best performance. In particular, this chapter explains:

- [Configuration Operations on Hosts](#)
- [Administration Tasks](#)
- [Using Tools and Commands](#)
- [Adding Host Targets](#)
- [Running Host Command](#)
- [Miscellaneous Tasks](#)

Configuration Operations on Hosts

There are a number of configuration operations you can perform on hosts to enhance their effectiveness. These operations include:

- [Configuring File and Directory Monitoring Criteria](#)
- [Configuring Generic Log File Monitor Criteria](#)
- [Configuring Program Resource Utilization Monitoring Criteria](#)

To access the configuration operations explained in this section, perform the following steps:

1. From the **Targets** menu, select either **All Targets** or **Hosts**.
2. Either type the name of the desired host in the Search field or scroll down to the name of the host in the Name column.
3. Click the name of the host.
4. From the **Host** menu, select **Monitoring**, then **Metric and Collection Settings**.

Follow the instructions for each configuration explanation.

Configuring File and Directory Monitoring Criteria

Enterprise Manager monitors the files and directories for the operator-specified criteria on hosts running various flavors of the UNIX operating system. The operator should configure the criteria for monitoring the desired files and directories.

Operator should specify the criteria for file and directory monitoring using the Monitoring Settings Page.

To configure the file and directory monitoring criteria, do the following:

1. On the Metric and Collection Settings page, select **All metrics** in the View menu. Locate the File and Directory Monitoring metrics. The metrics are:
 - File or Directory Attribute Not Found
 - File or Directory Permissions
 - File or Directory Size (MB)
 - File or Directory Size Change Rate (KB/minute)
2. After reviewing each metric, decide which metrics need to change. Click the pencil icon to navigate to the corresponding Edit Advanced Settings page.
3. You can specify new criteria for monitoring by clicking **Add** on this page. Refer to Notes about Specifying Monitored Objects for details on configuring the criteria.
4. You can edit or remove existing criteria by selecting the row from the Monitored Objects table and clicking **Edit** or **Remove**.

Notes about Specifying Monitored Object

File or Directory Name specifies the criteria to be monitored. Each row in the Monitored Object table specifies a unique criteria to be monitored.

File or Directory Name is the name of the file or directory being monitored from the host operating system. Specified value should correspond to the absolute path for the desired file or directory.

Either exact name or name with SQL wild cards (%) and _) can be specified for File or Directory Name. SQL wild card matches 0 or more characters. SQL wild card _ matches exactly one character.

Configuring Generic Log File Monitor Criteria

Enterprise Manager monitors the log files for the occurrence of operator-specified patterns that the owner of the Management Agent software is able to read. You can use this facility for monitoring abnormal conditions recorded in the log files present on the host.

Log files are periodically scanned for the occurrence of desired patterns and an alert is raised when the pattern occurs during a given scan. During a scan, new content created since the last scan is searched for the occurrence of the desired patterns.

The operator should specify the criteria for log file monitoring using the Metric and Collection Settings Page. To configure the log file monitoring criteria, first identify the monitoring criteria using the form <log file name, match pattern in perl, ignore pattern in perl>.

Perform the following steps using the Enterprise Manager console:

1. Search for Log File Pattern Matched Line Count in the table displayed for Metrics with Thresholds filter. Click the pencil icon in this row to navigate to the Edit Advanced Settings: Log File Pattern Matched Line Count page.
2. You can edit or remove existing criteria by selecting the row from the Monitored Objects table and clicking **Edit** or **Remove**. Refer to Notes about Specifying Monitored Objects for details on configuring the criteria.

Optionally, perform the following steps directly in the ORACLE_HOME directory of the Management Agent present on the managed host.

1. By default, matching number of lines is reported through log file monitoring. To enable upload of matching content for a specific file, add the absolute path for the file to the \$ORACLE_HOME/sysman/admin/lfm_ifiles file. The \$ORACLE_HOME/sysman/admin/lfm_ifiles.template file is a template needed for creating the \$ORACLE_HOME/sysman/admin/lfm_ifiles file.
2. For security purposes, you may want to disable monitoring of sensitive files by Enterprise Manager permanently by adding the names of the sensitive files to the \$ORACLE_HOME/sysman/admin/lfm_efiles file. The \$ORACLE_HOME/sysman/admin/lfm_efiles.template file is a template needed for creating the \$ORACLE_HOME/sysman/admin/lfm_efiles file.

Notes about Specifying Monitored Object

The set of columns (Log File Name, Match Pattern In Perl, Ignore Pattern In Perl) uniquely specifies the criteria to be monitored. Each row in the Monitored Object table specifies a unique criteria to be monitored. Multiple criteria can exist against the same log file.

Column	Description
Log File Name	In this column, specify the absolute path for the log file to be monitored. SQL wild characters can be used for specifying multiple file names. Examples: (a) /orahome/log/f1.log This value would monitor single log file.(b) /orahome/log/%.log This value would monitor all files with suffix .log in /orahome/log directory.
Match Pattern in Perl	In this column, specify the pattern to be matched for. Perl expressions are supported. This column specifies the pattern that should be monitored in the log file. During each scan, the file is scanned for occurrence of the specified match pattern [with case ignored]. Example: (a) Pattern Value = ERROR This pattern will be true for any line containing error (b) Pattern Value = .*fan.*error.* This pattern will be true for lines containing fan and error
Ignore Pattern in Perl	This column specifies the ignore pattern. In the given Log file, line containing the match pattern will be ignored if the ignore pattern is contained in that line. In this column, specify any pattern that should be ignored. Perl expressions are supported. If nothing needs to be ignored, specify %
Time Stamp	If this column is present, always specify it to be %.

Configuring Program Resource Utilization Monitoring Criteria

Enterprise Manager monitors the CPU resources consumed by the combination of <program name, owner> on hosts running various flavors of UNIX operating systems. The operator should configure the criteria for monitoring the resources consumed. This facility can be used for usage tracking of CPU resources.

The operator should specify the criteria for program resource utilization monitoring by using the Monitoring Settings Page.

To configure the program resource utilization criteria, do the following:

1. On the Metric and Collection Settings page, select **All metrics** in the View menu. Locate the Program Resource Utilization metrics. The metrics are:
 - Program's Max CPU Time Accumulated (Minutes)
 - Program's Max CPU Utilization (%)
 - Program's Max Process Count
 - Program's Max Resident Memory (MB)
 - Program's Min Process Count
 - Program's Total CPU Time Accumulated (Minutes)
 - Program's Total CPU Utilization (%)
2. After reviewing each metric, decide which metrics need to change. Click the pencil icon to navigate to the corresponding Edit Advanced Settings page.

3. You can edit or remove existing criteria by selecting the row from the Monitored Objects table and clicking **Edit** or **Remove**. Refer to Notes about Specifying Monitored Objects for details on configuring the criteria.

Notes about Specifying Monitored Object

Set of <Program Name, Owner> specifies the criteria to be monitored. Each row in the Monitored Object table specifies an unique criteria to be monitored.

Column	Description
Program Name	Program name is the name of the command being executed on the host operating system. On UNIX systems, ps command displays the name for each process being executed. Either exact name or name with SQL wild cards (% and _) can be specified for program name. SQL wild card matches 0 or more characters. SQL wild card _ matches exactly one character.
Owner	Owner is the name of the user running the given process on the host operating system. On UNIX systems, ps command displays the name for each process being executed. Either exact name or name with SQL wild cards (% and _) can be specified for owner. SQL wild card matches 0 or more characters. SQL wild card _ matches exactly one character.

Administration Tasks

The Administration tab gives you access to all the administrative tasks you can perform on this host. With the categories listed, you can easily access the appropriate pages for system services, network connections, and user and group settings. The tasks include starting services, setting users, and configuring network cards.

Using the Administration tab, you can manage:

- Services
View statistics of individual service and edit their services.
Note: This feature is only available on hosts running Oracle Linux, Red Hat Linux and SUSE Linux Operating Systems (x86 and x64 architectures only).
- Default System Run Level
- Network Cards
Manage routing configuration, view configuration statistics, and view network file system clients.
- Host Lookup Table
- NFS Client
- User and Group Administration
Manage user and group settings.

Note: For Oracle Linux 7 systems, to perform administration tasks on these items, your host must have patched with latest 13.2.0.0.0 bundle release. For Oracle Linux 5 and Oracle Linux 6, in case host can not be patched then you should have YaST and EM Wrapper scripts installed. See the [Required Installations](#) section of the [Setting Up the Environment to Monitor Hosts](#) chapter for information.

To access the administration tasks explained in this section, perform the following steps:

1. From the **Targets** menu, select either **All Targets** or **Hosts**.
2. Either type the name of the desired host in the Search field or scroll down to the name of the host in the Name column.
3. Click the name of the host.
4. From the **Host** menu, select **Administration**, then the entity on which you want to make changes.

Note: To see a video showing the navigation in the Administration menu option, see http://www.youtube.com/watch?v=ROfqR2GhQ_E.

Services

The Services page provides a list of all the services and their statistics for this host. This page enables you to:

- Start, stop, and restart services.
- Access the page that allows you to edit the properties of individual services.
- View the current system run level. When no run level is defined for the service, the service uses the current *system* run level.
- Refresh updates the services administration page with latest list of services available on the host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.
- Determine whether the service is enabled and view the service run levels:

Run Level	Description
0	System halt
1	Single user mode Only one user can be logged on at any point in time. Additional users will not be allowed to log on until the user using the system logs off.
2	Basic multiuser mode without network
3	Full multiuser mode with network
4	This run level is for future Oracle use.
5	Full multiuser mode with network and X display manager
6	System reboot

Note: Be aware that you must restart the system for the run level to take effect.

Default System Run Level

The Default System Run Level page allows you to change the run level that the system uses when it reboots. The change to the system run level takes effect after the system reboots.

Run Level	Description
0	System halted
1	Single user mode
2	Basic multiuser mode without network
3	Full multiuser mode with network

Run Level	Description
4	Unused. This run level is for future Oracle use.
5	Full multiuser mode with network and X display manager
6	Reboot the system.

Note the following:

- Click **Change** to change the host credentials. You must have SUDOER credentials to complete the Default System Run Level operation. If you do not have SUDOER credentials, this button provides the opportunity to change credentials.
- Click **Cancel** to abandon the changes and return to the Host Administration page.
- Click **Save** to keep the changes made to the default system run level and return to the Host Administration page.
- You need to install the YAST toolkit to use the Default System Run Level feature. See Required Installations.

The run levels are not immediately affected and hence the default run level and current run level may be different if the system has not been rebooted.

 **Note:**

The default run level is a powerful tool. You should only change the default system run level if you have sufficient knowledge and experience. Changing the default system run level inappropriately could result in improper system functionality after rebooting.

Network Card

The Network Card page provides detailed information on the network cards in your enterprise. With this information, you can decide whether edits need to be made to global domain name system (DNS) settings and routing table configuration.

Using the Network Card page, you can:

- Configure, enable, and disable network cards.
- View the device name and IP address of the network card used by Enterprise Manager.
- View the DNS settings and click **Edit** to change the global DNS settings for the listed domains.
- Edit a default gateway if it is available, or click **Add** to define a routing configuration.

Notes

Note the following:

- Click **Done** to exit the page without making any changes.
- Click **Change** to edit the credentials used for this page.

Configuring the Network Card

Use the Configure Network Card page to change the specifications of the network card. Using this page, you can:

- Opt to use either of the following setup methods: Static Address Setup or Automatic Address Setup using DHCP.
- Add the IP Address
- Add Subnet Mask information
- Maximum Transfer Unit (Bytes)

Adding Routing Configuration

Use the Add Routing Configuration page to add either a gateway or network device for routing requests. Using this page, you can:

- Specify the **Gateway**
- Specify the network **Device**
- When configuring routing to a network, specify the **Netmask** and **Destination**

Host Lookup Table

The Host Lookup Table page displays the mapping of IP address to a host name or its aliases. Using this page, you can:

- Edit hostname and aliases.
- Delete a lookup table entry for a host.
- Add a lookup table entry for a host by accessing the Add Host Configuration page.
- Refresh updates the **Host Lookup Table** page with latest host Look up data from host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.

Note the following:

- A host can have one or more aliases.
- Click **Done** to exit the page without making any changes.
- Click **Change** to edit the credentials used for this page. You do not need to reboot for changes to take effect.
- Each alias should be comma-separated.

NFS Client

The NFS Client page provides a list of all the network file system (NFS) clients mounted on the current host. Using this page, you can:

- Mount, unmount, and delete clients.
- Access the page that allows you to edit the properties of individual clients.
- Access the page that allows you to add NFS clients to the host.
- Refresh updates the network file system (NFS) client page with latest list of network file clients mounted on current host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.

- View the statistics of the various clients.

Client Statistic	Description
Server	Hostname of the remote NFS server
Remote File System	Location of the remote file system
Mount Point	Local mount point
Mounted	Indicates whether the remote file system is mounted
Persist Over Reboot	Retains mount points between reboots
Options	Displays mount options

Note the following:

- Click **Done** to exit the page without making any changes and returning to the previous page.
- Click **Change** to edit the host credentials used for this page.

Adding and Editing an NFS Client

The Add and Edit NFS Client pages provide the ability to mount a file system on a remote NFS server to a location on a local host. Using these pages, you can create and edit an NFS mount by providing:

- The local mount point
- The name of the NFS Server host name
- The location of the remote file system
- Mount options

Options

```
ro
rsize=32768
wsize=32768
acregmin=1200
acregmax=1200
acdirmin=1200
acdirmax=1200
hard
intr
tcp
lock
rw
nosuid
nodev
```

Note the following:

- Click **Cancel** to ignore all changes and return to the NFS Client page.
- Click **OK** to accept all changes made. All changes are implemented immediately.
- Check **Persist Over Reboot** to ensure mounts are available between reboots.

User and Group Administration (Users)

The User and Group Administration (Users) page provides a list of all the user accounts on this host along with their statistics. On this page you can:

- Access the page to edit user account statistics
- Access the page to add a user account to the host
- Delete a user account from the host
- Refresh updates user administration page with latest list of user accounts available on the host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.
- View the statistics of specific user accounts

User Statistic	Description
Login	User account name that allows you to access the software
Name	Full name of user account Many logins can have the same name. For example, logins aim1, aim2, and aim3 can all have the same name - AIM Manager.
UID	User account identifier This identifier is unique to the login.
Groups	Categories to which the user account belongs User account inherits the permissions given to the group.

Adding or Editing a Local User

This option enables you to add and edit a user account to this host. The following table describes the fields.

Group Information	Description
User's Full Name	Full name of user account
Username	Name used as login
Password	In conjunction with the username, a set of characters that allows access to this host The password must be no shorter than 5 characters and no longer than 72 characters. If you have changed the user's password, ensure you inform the user of this change.
Confirm Password	The password typed in this field must be exactly as the password typed in the Password field If the confirm password does not match the password typed in the Password field, either retype the password or define a new password and confirm it.
User ID (UID)	User identifier This identifier is unique to the user account. The ID must be a whole number greater than 499.
Home Directory	Ensure the home directory begins with a slash (/)
Additional User Information	Enter any additional user information
Login Shell	Select the Login Shell from the list of available shells from the drop-down list

Group Information	Description
Default Group	Select the default group from the drop-down list of available groups
Group Memberships	Groups to which the user account belongs Group names are separated by a comma. Do not include any spaces. An example: <i>adm,daemon,root</i>

When editing a local user, you can:

- Change the password
- Change the profile information, for example, the default group

User and Group Administration (Groups)

The User and Group Administration (Groups) page provides a list of all the groups on the host and their statistics. Using this page, you can:

- Access the page to edit group statistics.
- Access the page to add a group to the host.
- Delete a group from the host.
- Refresh updates group administration page with latest list of groups available on the host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.
- View the statistics of particular groups.

Group Statistic	Description
Group Name	Name of the group.
Group ID	Group identifier. This identifier is unique to the group.
Group Members	Groups that belong to the group. Group shares the permissions given to the subordinate groups.

Note: This feature is only available on Linux.

Adding or Editing a Local Group

The Add New Local Group page provides you the opportunity to add a group to this host. On the Add New Local Group page, you can add information for the fields listed in the following table:

Group Information	Description
Group Name	Name of the group.
Group ID	Group identifier. This identifier is unique to the group.
Group Members	Groups that belong to this group. Group shares the permissions given to the subordinate groups. Group names are separated by a comma. Do not include any spaces: for example, <i>adm,daemon,root</i> .

When editing a local group, you can:

- Change the group ID

- Add, delete, or change group members

Using Tools and Commands

There are a number of tools and commands available to you to facilitate your administration of hosts. This section introduces you to:

- Sudo and Power Broker
- Host Command
- Remote File Editor

Enabling Sudo and Power Broker

The sudo command is a program for UNIX-like operating systems that allows users to run programs with the security privileges of another user (normally the root user). It also provides auditing capabilities.

PowerBroker is a tool used for restricting the type of commands that can be run by users and maintains an audit trail of what users have done or have tried to do. PowerBroker allows for policy-defined authorization controls which allow administrators to define where and when their end-users can access other accounts they are authorized to use, up to and including root.

To enable Sudo or PowerBroker, perform these steps:

1. From the **Setup** menu located at the top-right of the page, select **Security** then select **Named Credentials**.
2. On the Named Credentials page, click **Create**.
3. On the Create Credential page, provide host credentials with root privileges.
4. In the Credential Properties section, select **Sudo** or **PowerBroker** from **Run Privilege**.
5. Provide the details for Sudo or PowerBroker and the system performs the administrative task.

Using Sudo or Power Broker

To use Sudo or Power Broker as Linux administrator, perform the following steps:

- Navigate to the Host target page.
- View list of administration activities on Host target.
- Select an administration task.
- Provide host credentials with root privileges. Provide information for Sudo (runas) support or Power Broker (profile) support.
- Provide the details and the system performs the administrative task.

Executing the Host Command Using Sudo or PowerBroker

To execute Host command using Sudo or PowerBroker, perform these steps:

1. Navigate to the Host target page.

From the **Targets** menu, select **Hosts**. On the Hosts page, highlight the row containing the name of the host in which you are interested.

2. Click the **Run Host Command** button.
3. Provide host target credentials. Provide information for Sudo (runas) support or PowerBroker (profile) support.
Note: On the target host, the `/etc/sudoers` file needs to be present with the target user information inserted.
4. Type the specific command to be run on the system and view the command output.

Using Remote File Editor

The Remote File Editor enables you to view and edit text files on the remote host. For example, using this utility, you can update the contents of configuration files on the remote host. In addition, you can:

- With the appropriate privileges, view and edit any text file present on the remote host.
- Save a file that has been edited on the remote host by clicking **Save**.
- Save the contents to a different file on the remote host by clicking **Save a Copy**.
- Change to another user account or use another set of Host Preferred Credentials by clicking **Change** next to User.
- After you have opened a file for editing, select a new file for editing by clicking **Change** next to File Name.
- Revert to text at the time of the last successful save operation by clicking **Revert**.

Accessing Remote File Editor

To navigate to the Remote File Editor, perform the following steps:

1. From the **Targets** menu, select **Hosts**. On the Hosts page, click the name of the host in which you are interested.
2. On the resulting Host home page, select **Remote File Editor** from the Host menu (located at the top-left of the page).
3. If the preferred credentials are not set for the host target, the Host Credentials page appears. Three options are available: Preferred, Named, and New.
You must have permissions to perform operations on a target host.
4. Once credentials are set, you can perform the following on the Remote File Editor page:
 - Perform operations on files, for example, listing of files in a directory, opening a file for reading, editing, and saving.
 - Provide host target credentials.
 - Provide details of the file and type of operation to be done.

Note the following:

- The file must be an ASCII text file and cannot be larger than 100 KB.
- In the case where the credential check is successful, the file exists and you have read privilege on the file, the file content is loaded for editing.
- If you do not have write privilege, you will not be able to save the file. Click **Save a Copy** and save the file to a directory on which you have write privilege.

- In the case the file does not exist but you have write privilege on the directory, a new empty file is opened for text input.

Adding Host Targets

To add a host target, install the Oracle Management Agent on the host computer you want to manage.

Detailed information is available in:

- Installing Oracle Management Agents chapter of the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*. The Add Host Targets wizard is described in detail.
- Discovering, Promoting, and Adding Targets chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Running Host Command

The Host Command enables you to interactively perform administrative operations on a single host, multiple hosts, or group composed of multiple hosts. For example, using this command the DBA can list the contents of some common directory of a set of hosts.

Accessing Host Command

To access the Execute Host Command page to perform administrative operations on hosts, do the following:

1. From the **Targets** menu, select either **All Targets** or **Hosts**.
2. Either type the name of the desired host in the Search field or scroll down to the name of the host in the Name column.
3. Click the name of the host. Enterprise Manager displays the Home page for the host.
4. From the **Host** menu, select **Execute Host Command**.
5. On the Host Credentials page, type the user name and password for the host.
6. The Execute Host Command page appears.

Executing Host Command Using Sudo or Power Broker

To execute the Host Command using Sudo or Power Broker, perform the following steps:

1. Navigate to the Host target page.
2. Run the command on the host target.
3. Provide host target credentials. Provide information for Sudo (runas) support or Power Broker (profile) support.
4. Type a specific command to be run on the system and view the command output.

Note: If the credentials (with runas Sudo/PowerBroker) are not set, then you will be prompted to create credentials. To create credentials, select **Security** from the **Setup** menu, then select either **Named Credentials** or **Preferred Credentials**.

Execute Host Command - Multiple Hosts

The Execute Host Command page enables you to type operating system commands against multiple hosts and immediately view the results. This gives you the opportunity to perform administrative operations on multiple hosts within the context of Enterprise Manager.

On this page, you can:

- Refine the command, reexecute the command, and view the execution results, all on the same page.
- Either type operating system commands, load the commands from a script on the browser machine or on the host, or load host commands from a job defined in the job library.
- Select hosts individually or through the use of a group. You can also switch to Single Target Mode where only one host target is acted upon.
- Interactively view command execution results or hide the results to be viewed at a later time.
- Use preferred credentials or override preferred credentials.
- Add targets and modify the targets list.

Note the following:

- Re-execute the command by clicking **Execute**.
- Cancel execution of the command is possible when the Processing: Executing Host Command page appears.
- Execution history reflects the host commands that have been executed in the current Execute Host Command session, as well as any host commands executed in previous sessions that were executed with the 'Keep the execution history for this command' option chosen.
- Clicking **Add** launches the Target Selection page with the target type list limited to host targets and any groups that contain host targets.
- When saving the OS script or execution results, the saved file is located on the browser machine.
- No more than 500 lines appear in the history list.
- At most, 10 rows of command execution results data will be displayed in the targets table. If more data is returned, click the **Execution Status** icon in the table or click **Complete Execution Results**.
- When switching from multiple to single target mode, the first host in the targets table will be used.

Target Properties

The host command is executed using the Enterprise Manager job system. The job system allows you to specify system variables called *target properties*. The supported target properties are listed in the following table. Note that the available properties change according to the type of target the job is run against.

Name	Description	Target Type
%emd_root%	Location of Management Agent	Host, Database Instance
%perlbin%	Location of Perl binary used by Management Agent	Host, Database Instance
%TargetName%	Target Name	Host, Database Instance
%TargetType%	Target Type	Host, Database Instance
%orcl_gtp_comment%	Comment	Host, Database Instance
%orcl_gtp_contact%	Contact	Host, Database Instance
%orcl_gtp_deployment_type%	Deployment Type	Host, Database Instance
%orcl_gtp_line_of_business%	Line of Business	Host, Database Instance
%orcl_gtp_location%	Location	Host, Database Instance
%OracleHome%	Oracle home path	Database Instance
%Port%	Port	Database Instance
%SID%	Database SID	Database Instance
%Role%	Database Role	Database Instance
%MachineName%	Listener Machine Name	Database Instance

Note the following:

- Property names are case-sensitive.
- Properties can be used in both the Command and OS Script fields.
- To use the % character without a target property, escape it with a second %.
- To use the Database Instance target type, launch Execute Host Command from a group containing one or more host targets and switch the Target Type.

Examples:

To execute a Perl script, passing in the target name as an argument, enter the following in the Command field: %perlbin%/perl myPerlScript %TargetName%

To execute a program in the directory identified by the TEMP environment variable on a Windows host: %%TEMP%%/myProgram

Execute Host Command - Group

The Execute Host Command page enables you, in the context of a group, to type operating system commands against multiple hosts and immediately view the results. This gives you the opportunity to perform administrative operations on multiple hosts within the context of Enterprise Manager.

On this page, you can:

- Choose the target type. You can choose hosts directly, or choose hosts by way of database instance targets.
- Refine the command, re-execute the command, and view the execution results, all without leaving the page.

- Either type operating system commands, load the commands from a script on the browser machine or on the host, or load host commands from a job defined in the job library.
- Select hosts individually or through the use of a group. You can also switch to Single Target Mode where only one host target is acted upon.
- Interactively view command execution results or hide the results to be viewed at a later time.
- Use preferred credentials or override preferred credentials.
- Add targets and modify the targets list.

Note the following:

- Re-execute the command by clicking **Execute**.
- Cancel execution of the command by clicking **Cancel** on the Processing: Executing Host Command page.
- If the current target type is Host, clicking **Add** launches the Target Selection page with the target type list limited to host targets and any groups that contain host targets.
- If the current target type is Database Instance, clicking **Add** launches the Target Selection page with the target type list limited to database targets and any groups that contain database targets.
- Execution history reflects the host commands that have been executed in the current Execute Host Command session, as well as any host commands executed in previous sessions that were executed with the 'Keep the execution history for this command' option chosen.
- Changing the target type re-initializes the host credentials, command, OS script, and targets table.
- When saving the OS script or execution results, the saved file is located on the browser machine.
- At most, 10 rows of command execution results data will be displayed in the targets table. If more data is returned, click the **Execution Status** icon in the table or click **Complete Execution Results**.
- When switching from multiple to single target mode, the first host in the targets table will be used.

Execute Host Command - Single Host

The Execute Host Command page enables you to type operating system commands against one host and immediately view the results. This gives you the opportunity to perform administrative operations on the host within the context of Enterprise Manager.

On this page, you can:

- Refine the command, reexecute the command, and view the execution results, all without leaving the page
- Switch to Multiple Target Mode where multiple host targets are acted upon
- Change credentials

Note the following:

- Reexecute the command by clicking **Execute**.
- Cancel the execution of the command by clicking **Abort**.
- Context will be preserved if you switch to Multiple Target Mode, including the host command, host, and credentials.

Load OS Script

The Load OS Script page is used to load commands from a script into the **OS Script** field on the Execute Host Command page.

On this page, you can:

- Click **Browse** to launch the browser's file selector window to locate and choose a script file.
- Click the Host field's search icon to choose which host to search, then click the Host File field's search icon to locate and choose a script file on that host.

Load From Job Library

The Load From Job Library page provides the mechanism by which to search the Job Library directly for an existing job. This encourages the reuse of existing jobs.

On this page, you can click the icon in the 'Load' column of any row to return to the Execute Host Command page loading the complete context of the library job in that row. The complete context includes the host command, OS script, targets, and credentials.

Note: Jobs displayed in the table on this page will be host command jobs from the Job Library that are owned by the current Enterprise Manager user.

Execution History

The Execution History page lists the host commands executed during the current Enterprise Manager session, as well as any host commands executed in previous sessions that were executed with the 'Keep the execution history for this command' option chosen.

On this page, you can:

- Click the icon in the 'Load' column of any row to return to the Execute Host Command page loading the complete execution context of the host command in that row. The complete execution context includes the host command, OS script, targets, credentials, and results.
- Click the icon in the 'Load Command And OS Script Only' column of any row to return to the Execute Host Command page loading only the host command and the OS script in that row. Any targets, credentials, and most recent results will remain.
- Click the icon in the 'Remove' column of any row to remove the host command in that row, along with all its execution context, from the Execution History. This will delete the job that was used to execute the host command.

Execution Results

The Execution Results page provides the full listing of the results of the executed host command, for a specific host. This listing chronicles all the information from the run.

On this page, you can:

- Cut the text from the listing and paste it into another script.
- Study the results uninterrupted and separate from all the other executions.

Note: The extent of the editing features is dependent upon the browser displaying the results.

Miscellaneous Tasks

This section explains the following:

- [Enabling Collection of WBEM Fetchlet Based Metrics](#)
- [Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts](#)
- [Adding and Editing Host Configuration](#)

Enabling Collection of WBEM Fetchlet Based Metrics

To enable the Web-Based Enterprise Management (WBEM) Fetchlet based collections for the host target, configure the WBEM Host Username and WBEM Host Password properties.

Note: Host targets running with Linux and Windows operating systems do not, by default, have WBEM Fetchlet based metric collections. This is due to the fact that these operating systems, by default, do not run a DMTF (Distributed Management Task Force) WBEM-compliant Common Information Model (CIM) Object Manager. If the systems have been configured to run a WBEM-compliant CIM Object Manager, then WBEM Fetchlet based metric collections will be possible.

To configure the collections on systems with WBEM compliant CIM Object Managers, use the following steps:

1. Navigate to the home page of the specific host.
2. From the **Host** menu, select **Target Setup**, then **Monitoring Configuration**.
3. Set the Username and Password values.

Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts

Hardware-specific monitoring is available for Dell PowerEdge Linux hosts with Enterprise Manager. To enable the hardware monitoring of your Dell PowerEdge Linux hosts, perform the following steps:

1. Download the latest Dell OpenManage Server Administrator (OMSA) software certified for your Linux OS by accessing the Dell FTP site at <http://downloads.dell.com>.

To identify the latest OMSA software, search for the latest `om??_lnx_managed_system*.gz` file on the website:<http://downloads.dell.com>

Note:

At the time of writing this is the current Dell site, contact your Dell support contact if you cannot find the latest Dell OMSA software.

2. Install the software using the instructions provided by Dell.

3. Verify that the installation was successful by performing the following steps:
 - Verify that snmp daemon is up and running.

```
% ps -ef | grep snmpd
```
 - Verify that the following commands execute without errors:

```
% /usr/bin/omreport about
```

```
% /usr/bin/omreport system version
```
4. Verify the Dell OMSA software is functioning correctly by verifying that the Dell OpenManage Server Administrator website is up and running.
 - Using your web browser, access the URL - `https://target_hostname:1311`.
 - Log in using an operating system account. Check that you are able to successfully log in and navigate in the website.
5. If the SNMP Community String for the SNMP daemon running on the Linux host is not *public*, set the SNMP Community String property in Enterprise Manager using the following steps:
 - a. Log in into Enterprise Manager Cloud Control.
 - b. Navigate to the home page of the specific host.
 - c. From the **Host** menu, select **Target Setup**, then **Monitoring Configuration**.
 - d. Set the SNMP Community String property to the correct value on this page.
6. Restart the Management Agent on the host.

```
% login into host as the owner account of emagent software
```

```
% emctl stop agent
```

```
% emctl start agent
```

7.  **Note:**

The following step is not required if the Dell OMSA software was functional prior to the previous startup of the Management Agent.

Verify that hardware monitoring is working correctly using the following steps:

- a. Log in to Enterprise Manager Cloud Control.
- b. Navigate to the home page of the specific host.
- c. From the **Host** menu, select **Monitoring**, then **All Metrics**.
- d. Verify that the following metrics are present on this page: Fans, Memory Devices, PCI Devices, Power Supplies, Processors, Remote Access Card, System BIOS, and Temperature.
- e. You can navigate to the metric data page by clicking on one of the metrics listed in the previous step and view the data that Enterprise Manager is able to fetch.

Adding and Editing Host Configuration

The Add Host Configuration page provides you the opportunity to add a host to the `/etc/hosts` file. When you add a host to the `/etc/hosts` file, Enterprise Manager can then translate host names into IP addresses. In most cases, host names are much easier to remember than IP addresses.

On the Add Host Configuration page, you can:

- Associate a host's IP address with a host name.
- Add, delete, or change the aliases associated with a host

Note: Changes take effect after you click **OK**.

When editing a host configuration, you can:

- Change the name of the host
- Add, delete, or change the aliases associated with the host

Note: Changes are in effect immediately after you click **OK**.

Managing Oracle Linux Homes

Oracle Linux Home target enables the user to access and perform complete management and monitoring of Oracle Linux hosts.

Oracle Linux Home includes:

- Oracle Linux host administration and management
- Bare Metal Provisioning (BMP)
- Linux Patching
- Oracle Ksplice patching to update Oracle Linux operating system kernel and key user space libraries, while the OS is running, without a reboot or any interruption

This chapter covers the following:


- [Oracle Linux Home Target](#)
- [Oracle Linux Home page](#)
- [Ksplice for Oracle Linux](#)

Oracle Linux Home

The Oracle Linux Main menu contains the following options:

- Add Linux Hosts —
The Add Linux Hosts Page redirects to the add targets page.
- Linux Patching — Linux Patching contains sub-items such as, Setup and Patching.
 - Setup redirects to existing setup linux patching page.
 - Patching redirects to the actual linux patching page.
- Bare Metal Provisioning — Bare Metal Provisioning redirects to the existing BMP page.


Figure 4-1 Oracle Linux Home



Target Navigation Tree (TNT) of "Oracle Linux Home"

The TNT contains "Oracle Linux Home" as the main item. It contains different groups/folders based on Oracle Linux Release Version, that is, Oracle Linux 6.7, Oracle Linux 6.9, and so on. The folders contain list of hosts with same Oracle Linux Release version. When you click any particular host name, it redirects to the existing home page of the host to perform monitoring and administration work.

Figure 4-2 Target Navigation Tree of Oracle Linux Home



Oracle Linux Home Target

You can open the Oracle Linux Home page from Enterprise Manager by performing the following steps:

1. From Enterprise Manager Cloud Control home page, click **Enterprise** select **Cloud**, and then click **Oracle Linux Home**.
The **Oracle Linux Home** page appears.
2. A new system repository target "Oracle Linux Home" is created which is associated with all the Oracle Linux Hosts (Physical or VM's). Clicking on that will redirect to the "Oracle Linux Home" page.

Oracle Linux Home page

Oracle Linux Home Page has the following sections:

- [General](#)
- [Overview of Incidents and Problems](#)
- [Host Flux](#)
- [CPU](#)
- [Memory](#)
- [Linux Patching Compliance/Summary](#)
- [Ksplice Patching Compliance/Summary](#)

General Region

The General Region shows the summary of added Oracle Linux Hosts like Total Hosts, Total Up Hosts, Total Down Hosts, and so on, for each Oracle Linux Released Version.

Figure 4-3 General Region

General

Total Linux Hosts 7 (↑ 6 🚩 1)

OS Version	Total	↑	🚩
Oracle Linux 6.7	1	-	1
Oracle Linux 6.9	1	1	-
Oracle Linux 7.3	4	4	-
Oracle Linux 7.4	1	1	-

Overview of Incidents and Problems

Overview of Incidents and Problems displays any incidents and problems occurred to any Oracle Linux Hosts.

Figure 4-4 Overview of Incidents and Problems

Overview of Incidents and Problems

Incidents

Updated in the last 24 hours 1

Updated in last 7 days 1

Breakdown of incidents updated in the last 7 days

Category	🔴	🚫	⚠️	🚩
Availability	-	-	1	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	-	-


Problems

Total Open 0

Host Flux

This region displays the number of hosts that have been newly added and retired as on a particular date for a time period of last 30 days.


Figure 4-5 Host Flux



CPU


This region displays the range of CPU utilization by the Linux host machines.

Figure 4-6 CPU Graph



For example, in [Figure 45-4](#) six hosts have a CPU utilization of 0-25%. To view the list of hosts along with the percentage utilization of CPU by each of them, click the bar graph. The CPU Utilization Table displays the details. See [Figure 45-5](#).


Figure 4-7 CPU Utilization Table



Memory

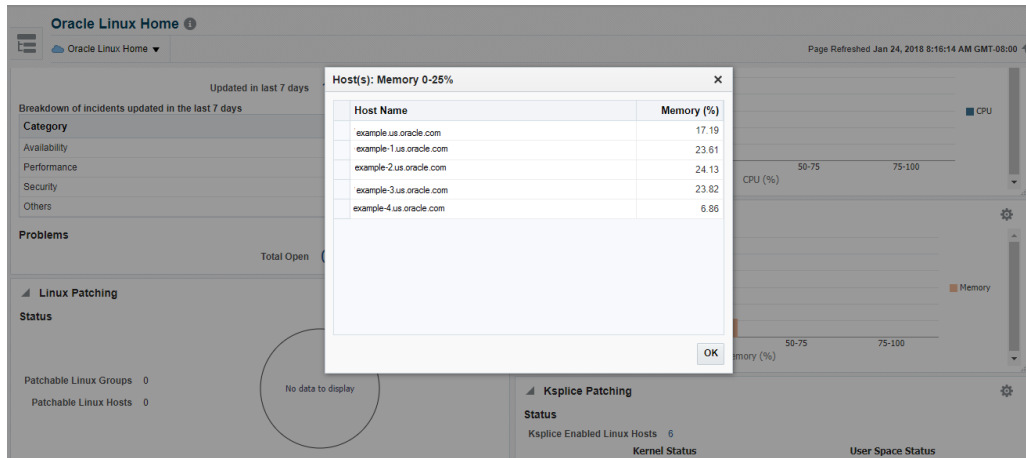
This region displays the range of memory utilization by the Linux host machines.

Figure 4-8 Memory Graph



For example, in [Figure 45-6](#) five hosts have a memory utilization of 0-25% and one host is utilizing in the range of 25-50%. To view the list of hosts along with the percentage of memory utilization by each of them, click the bar graph. The Memory Utilization Table displays the details. See [Figure 45-7](#)

Figure 4-9 Memory Utilization Table



Oracle Linux Patching

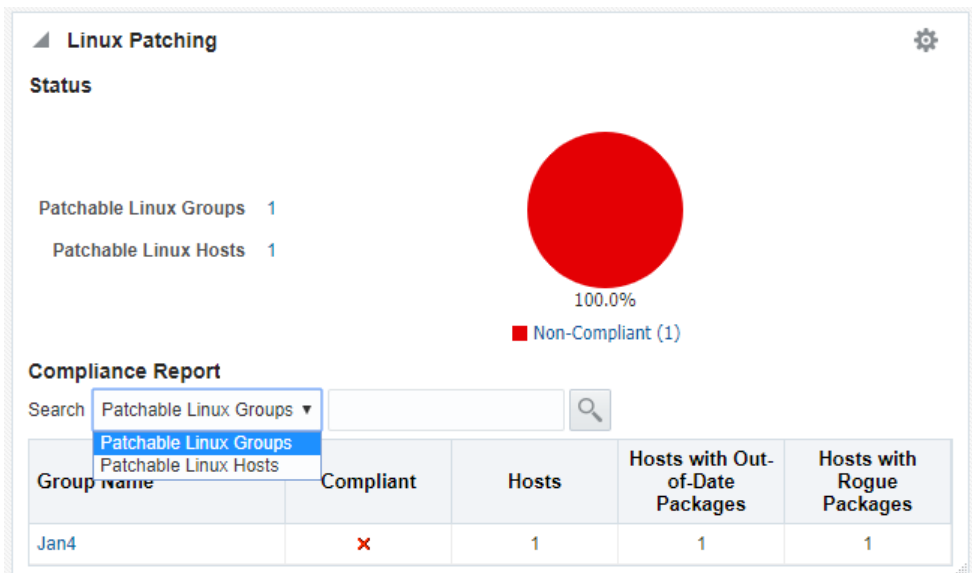
The Oracle Linux Patching screen contains:

- Linux Patching Status Region — This region displays the total Patchable Linux Groups, total Patchable Linux hosts, and the compliance pie chart.
- Linux Patching Compliance Region — This region displays a table that lists the names of the Patchable Linux Groups and Linux hosts.

Patchable Linux Groups

With Patchable Linux Groups, the following functions can be performed:

- If you choose groups then it will display group name, group compliance status, number of hosts present in the group, number of hosts with Out-of-Date Packages, and number of hosts with Rogue Packages.



- If you click **Patchable Linux Groups**, in the **Search** box, you will see the list of groups with details such as group name, group compliance status, number of hosts present in that group, number of compliant hosts in that group, number of non-compliant hosts in that group, number of compliance unknown hosts in that group, number of hosts with Out-of-Date Packages, and number of hosts with Rogue Packages.

ORACLE Enterprise Manager Cloud Control 13c

Oracle Linux Home

Oracle Linux Home > Patchable Linux Groups

Patchable Linux Groups

Group Name: Advanced Search

An out-of-date package is a host package for which an updated version is available. A "rogue" package is a host package that is not supposed to be installed on the host. A host without a single out-of-date or rogue package is a compliant host. A patching group comprising only compliant hosts is a compliant group.

View

Group Name	Compliant	Hosts	Compliant Hosts	Non-Compliant Hosts	Compliance Unknown Hosts	Hosts with Out-of-Date Packages	Hosts with Rogue Packages
Jan4	✘	1	0	1	0	1	1

- If you click the **Group** you will see the details of the selected group.

ORACLE Enterprise Manager Cloud Control 13c

Oracle Linux Home

Oracle Linux Home > Patchable Linux Groups > Patchable Group Details: Jan4

Patchable Group Details: Jan4

Host Name: Advanced Search

An out-of-date package is a host package for which an updated version is available. A "rogue" package is a host package that is not supposed to be installed on the host. A host without a single out-of-date or rogue package is a compliant host. A host needs reboot if an updated version of a package is available that can only be applied when the host reboots.

Host Name	Compliant	Needs Reboot	Out-Of-Date Packages	Rogue Packages
example.us.oracle.com	✘	-	6	1325

Total Hosts : 1

100.0%

■ Non-Compliant (1)

Patchable Linux Hosts

With Patchable Linux Hosts, the following functions can be performed:

- If you choose **Hosts** then it will display host name, group name, host is compliance status, host needs reboot status, number of Out-of-Date Packages, and number of hosts with Rogue Packages.

Linux Patching

Status

Patchable Linux Groups 1
Patchable Linux Hosts 1

100.0%
■ Non-Compliant (1)

Compliance Report

Search Patchable Linux Hosts

Host Name	Group	Compliant	Needs Reboot	Out-Of-Date Packages	Rogue Packages
example.us.oracle.com	Jan4	✘	-	6	1325

- If you click **Patchable Linux Hosts**, in the **Search** box you will see the list of hosts along with details such as the host name, group name, host is compliance status, host needs reboot status, number of Out-of-Date Packages, and number of hosts with Rogue Packages.

ORACLE Enterprise Manager Cloud Control 13c

Oracle Linux Home

Oracle Linux Home > Patchable Linux Hosts

Patchable Linux Hosts

Host Name example.us.oracle.com

An out-of-date package is a host package for which an updated version is available. A "rogue" package is a host package that is not supposed to be installed on the host. A host without a single out-of-date or rogue package is a compliant host. A host needs reboot if an updated version of a package is available when the host reboots.

Host Name	Group	Compliant	Needs Reboot	Out-Of-Date Packages	Rogue Packages
example.us.oracle.com	Jan4	✘	-	6	1325

Total Hosts : 1

100.0%
■ Non-Compliant (1)

- If you click **Host** name, you will see the details of the selected host.

ORACLE Enterprise Manager Cloud Control 13c

Oracle Linux Home

Oracle Linux Home > Patchable Linux Hosts > Host Non-Compliance: example-1.us.oracle.com

Host Non-Compliance: example-1.us.oracle.com

Package Name

An out-of-date package is a host package for which an updated version is available. A "rogue" package is a host package that is not supposed to be installed on the host. Some out-of-date packages can only be updated when the host reboots.

Package Name	Version	Status	Needs Reboot
abrt	2.0.8-40.0.1.el6	⊙	✘
abrt-addon-ccpp	2.0.8-40.0.1.el6	⊙	✘
abrt-addon-kerneloops	2.0.8-40.0.1.el6	⊙	✘
abrt-addon-python	2.0.8-40.0.1.el6	⊙	✘
abrt-cil	2.0.8-40.0.1.el6	⊙	✘
abrt-desktop	2.0.8-40.0.1.el6	⊙	✘
abrt-gui	2.0.8-40.0.1.el6	⊙	✘
abrt-libs	2.0.8-40.0.1.el6	⊙	✘
abrt-python	2.0.8-40.0.1.el6	⊙	✘
abrt-tui	2.0.8-40.0.1.el6	⊙	✘
abyssinica-fonts	1.0-5.1.el6	⊙	✘

Total Packages : 1344

Needs Reboot : No

98.59%
0.9673%
0.4449%

■ Up-to-Date (13) ■ Out-of-Date (6) ■ Rogue (1325)

Ksplice for Oracle Linux

Ksplice updates the Linux operating system (OS) kernel and key user space libraries, while the OS is running, without a reboot or any interruption. For more information, see [Oracle Ksplice](#). All the Oracle Linux Hosts where Enterprise Manager Agent is installed and Ksplice software is configured can be monitored and managed through the Oracle Linux Home Ksplice region. For more information, see [Oracle Ksplice User's Guide](#).



Note:

The minimum required ksplice version for Enterprise Manager is 1.0.29.

This section covers the following topics:

- [Ksplice Metrics](#)
- [Ksplice Patching](#)
- [Ksplice Linux Hosts Page](#)
- [Oracle Linux Home](#)
- [Target Navigation Tree \(TNT\) of "Oracle Linux Home"](#)

Ksplice Metrics



Note:

To access the Ksplice Metrics, from the **Host** menu on a host's home page, select **Configuration**, and then click **Latest**. The Oracle Linux Host needs the latest Ksplice software to be installed and configured with Oracle Premium Support before these metrics can report data. For more information, see [Oracle Ksplice User's Guide](#).

The following attributes are covered under Ksplice:

1. Ksplice
 - Ksplice Version
 - This attribute reports the version of the Ksplice software installed on the target host.
 - Ksplice Status
 - This reports if the host is configured to get updates from the Ksplice Server or if Ksplice offline.
 - Base Kernel Version
 - This attribute queries the stock Kernel running in the system.
 - This version does not represent the patched version, it shows only the one that booted the system.
 - Effective Kernel Version

- This attribute reports the Effective Kernel which means the Kernel version after the live patching including security fixes and others.
 - It also reports the date when the last applied patch was built.
 - Kernel Status
 - This attribute reports if the kernel of the host is up-to-date or out-of-date. A system is updated if it has all the available Ksplice patches installed.
 - In case Ksplice is offline, it is based on the `uptrack-updates-`uname -r`` package installed on the system.
 - Kernel Patches Installed
 - This reports the count of Ksplice packages installed on the system.
 - User Space Status
 - This attribute reports if the host's User Space Ksplice aware packages are up to date or if it is out of date.
 - In case Ksplice is offline, it is based on the local repository configured on the system.
 - User Space Packages Installed
 - This attribute reports the count of Ksplice User Space packages installed on the system.
2. Kernel Installed Patches
 - This attribute reports the installed Ksplice patches in the system.
 3. Kernel Available Patches
 - This attribute lists the available Ksplice patches for the system, in short it list the patches that have not yet been installed.
 - This information is derived based on the Ksplice configuration.
 - In case a Ksplice host is configured with a Ksplice server, it gets that information from the server.
 - In case the Ksplice is offline, it reflects the data based on the `uptrack-updates-`uname -r`` package installed on the system.
 4. User Space Installed Packages
 - This attribute reports the Ksplice User Space packages installed on the system.

Ksplice Patching

Ksplice Patching region on the Oracle Linux Home Page uses the metrics collected in Ksplice metrics to collate the Ksplice status over all the monitored Oracle Linux Hosts. For information on metrics, see [Ksplice Metrics](#). It contains the following two sub-regions:

- Ksplice Status Region

This region displays the total number of

- Ksplice Patching Configured Linux Hosts: If you click this link, it opens a pop-up window that lists the Oracle Linux Hosts which are configured with the Ksplice software.

- Ksplice Patching Non-Configured Linux Hosts: If you click this link, it opens a pop-up window that lists the Oracle Linux Hosts which are not configured with Ksplice software.

The Ksplice Status region also contains the following two pie charts:

- Kernel Status
- User Space Status

Each pie chart shows the status of all the hosts that is, how many hosts are compliant, non compliant, or non-compliant. If you click on a particular host, it opens another page with appropriate hosts list.

- Ksplice Summary Region

This region displays the table of hosts that lists the following:

- Ksplice Client (Online/Offline)
- Kernel Status (Compliant/Non-Compliant/Compliance unknown)
- User Space Status (Compliant/Non-Compliant/Compliance unknown)
- Effective Kernel Version

Figure 4-10 Ksplice Patching Status



Ksplice Linux Hosts Page

Ksplice Linux Hosts page contains the following tables:

- Ksplice Enabled Hosts with Ksplice Software
- Ksplice software Version
- Ksplice Client (Online/Offline)
- Kernel Status (Compliant/Non-Compliant/Compliance unknown-in case of unconfigured/offline systems)
- Number of Kernel Installed Patches
- User Space Status (Compliant/Non-Compliant/Compliance unknown-in case of unconfigured/offline systems)
- Number of User Space Installed Patches
- Base Kernel Version
- Effective Kernel Version

Host Ksplice Patches/Packages details

If you click the host name in the Ksplice detail table, a new page opens. This page lists the installed Ksplice patches on that host. If this host is a Ksplice Online host, it also lists the updates that are available on that host. These updates can be added or removed from this page.

If the host is a Ksplice Offline page, this page shows all the Ksplice Kernel patches available in the local repository or user space installed packages if there is any available in the local offline repository.

If Ksplice Enhanced Client Software is installed on the host, then it displays the list of installed/available user space patches. Otherwise, it shows the message “Install/Upgrade/Configure Ksplice Enhanced Client Software”.

The **Refresh** button in the screen helps in pulling the latest data to the dashboard. As this dashboard works in asynchronous mode, so in case there are some database changes or patches updated, you should use this **Refresh** button to sync the system. When you click this button, a dialog box appears asking for a confirmation.

For Install Updates, and Remove Updates you have to select/enter host root/sudo privileged credentials.

Note:

The Kernel Space Compliance status is based on in-memory kernel patches. It shows “Complaint”, if all the patches are installed, otherwise it shows “Non-Complaint”. When the patches are removed from the Ksplice Online System, it renders the system non-compliant. User Space Compliance Status is based on on-disk user space packages. When user installs the user space packages, in-memory process is updated on the on-disk packages as well and hence, the status shows “Compliant”. When user removes user space packages, it only removes the in-memory updates. When user space updates are on disk but not in memory, it continues to show “Compliant”. Though it does not remove the on-disk packages, the table still shows the user space updates to be “Compliant”.

Additional Setup for Real-time Monitoring

This section describes Oracle Enterprise Manager Cloud Control's (Cloud Control) Compliance features include the ability to monitor certain elements of your targets in real time to watch for configuration changes or actions that may result in configuration changes. These features include Operating System level file change monitoring, process starts and stops, Operating System user logins and logouts, Oracle database changes and more.

The real-time monitoring for these features takes place from the Cloud Control agent. Some of these monitoring capabilities require specific setup steps depending on the type of monitoring you will do and what Operating System is being monitored.

This chapter outlines the specific requirements and pre-requisites that exist to use the Compliance Real-time Monitoring features. For details on how to use Real-time monitoring from Cloud Control, see the chapter Compliance Management in this document. This chapter covers the following topics:

- [Overview of Real-Time Monitoring](#)
- [Overview of Resource Consumption Considerations](#)
- [Configuring Monitoring Credentials](#)
- [Preparing To Monitor Linux Hosts](#)
- [Preparing To Monitor Windows Hosts](#)
- [Preparing To Monitor Solaris Hosts](#)
- [Preparing to Monitor AIX Hosts](#)
- [Preparing To Monitor the Oracle Database](#)
- [Setting Up Change Request Management Integration](#)
- [Overview of the Repository Views Related to Real-time Monitoring Features](#)
- [Modifying Data Retention Periods](#)
- [Real-time Monitoring Supported Platforms](#)

Overview of Real-Time Monitoring

Real-time monitoring is configured through the Cloud Control Server. Users with the EM_COMPLIANCE_DESIGNER role create Compliance Standard Rules that are of type "Real-time Monitoring Rule." These rules are then associated with Compliance Standards and these standards are subsequently associated with one or more targets.

After the Compliance Standard to target association is complete, the set of monitoring rules are sent to the agent to enable real-time monitoring. All monitoring for Real-time monitoring occurs on the agents and all observed action data is sent from the agent to the Cloud Control server for reporting and data management.

Overview of Resource Consumption Considerations

The Real-time monitoring features are built into the Cloud Control agent. There are some specific resource considerations if you use the Real-time monitoring features. The following sections describe issues you should consider when using Real-time monitoring features.

OS File Monitoring Archiving

An optional setting when monitoring for file changes in real time is to make an archive copy of the file on the agent. When monitoring first begins, a copy of the file at that time is made and stored into a private directory in the ORACLE_HOME directory of the agent. Then, any subsequent changes to that file will result in additional copies of the file being archived in that same directory. This feature allows you to later perform a file diff from the user interface or to issue a job to roll back a file to a previous version.

This feature however will use disk space to make copies of the file. Care should be taken to ensure that this feature is only enabled for files that are critical. During rule creation, the user can specify how many copies of a file to save. The default is five historic versions. This can also be adjusted to tune potential resource consumption.

Select the checkbox in the Ignore Events Prior to Rule field to ignore all previous Oracle database change events when the Oracle database monitoring module runs the first time.

OS File Read Monitoring

The Operating System File level monitoring can monitor many types of changes to files, but can also monitor reads to files. If you have a Rule to monitor a facet that has file patterns that are read frequently, this may result in a very large number of observations. You can reduce the number of observations by ensuring that your Rule includes a filter on either time or a user that you want to ensure does not read the file.

For instance, monitoring the */etc/passwd* file for reads for All users will result in many observations being created. However, if you only monitor the */etc/passwd* file by a specific user, you can create a user filter for this specific user during rule creation. You will then only receive an observation when that specific user attempts to read the file.

Creating Facets That Have Very Broad Coverage

It is important to remember that facets are created to specify files that are very important to monitor for security/compliance purposes. For instance, monitoring all modifies to a log file that change every few seconds will result in reporting many file changes making it harder for you to identify the critical file changes you care about. Instead, in this case, it may be appropriate to create a rule to monitor the log file for all changes, but filter only when the log change is made by a non-application user. This would only capture the log file change if a regular user attempted to change or tamper with the log rather than when the log is simply being updated by an application.

Cloud Control Repository Sizing

Database sizing considerations for Real-time monitoring depend on several factors. The most important factor is the number of observations expected in a month. The second factor is the number of months data will be retained in the repository. Repository retention rates are explained in the Enterprise Manager Administration Guide.

In general, each observation consumes roughly 1.5KB of space in the database. This is a guideline and this number can vary depending on many factors for each installation.

For example, if a customer expected a total of 10 million Real-time observations per month across all targets and wanted to retain the data for 12 months, then the database size required for this would be roughly 180GB.

10,000,000 Observations x 12 Months x 1500 Bytes = 180,000,000,000 Bytes

This size represents Real-time monitoring data only and does not include database storage needs for other areas of Cloud Control.

The number of observations to expect per month can vary from environment to environment and can also depend on what types of monitoring are configured. You may be required to tune the expected size over time after Rules and Facets have been enabled for some time and configured to fit the organizations requirements. You can easily find your observations usage over a month by selecting **Compliance** from the **Enterprise** menu, then choosing **Browse By Systems UI Report** from the **Real-time Observations** page to select your systems and see the related counts of observations for each system over a period of time.

Configuring Monitoring Credentials

Many of the real-time monitoring capabilities require monitoring credentials that maintain the ability to launch monitoring programs with root privileges. These processes that Real-time monitoring uses begin with the prefix *nmxc*. Low-level monitoring uses operating system APIs that are not available to regular users.

Before starting to use the Real-time monitoring features on a target host for the first time, the following settings must be configured from the Enterprise Manager Console.

1. Ensure that the agent's *root.sh* script is run after agent installation.

After installing the agent, the *root.sh* script must be run as the root user. This script must be run before configuring the rest of these credential steps.

2. Configure Privilege Delegation.

Privilege Delegation settings are found from the **Setup** menu by choosing **Security**, then **Privilege Delegation**. On this page you can either set privilege delegation for each host manually or you can create a Privilege Delegation Setting Template.

Privilege delegation for each host that will have real-time monitoring must have SUDO setting enabled with the appropriate SUDO command filled in (for example, */usr/local/bin/sudo*).

3. Configure Monitoring Credentials.

Monitoring Credential settings are found from the Setup menu. Choose **Security** then **Monitoring Credentials**. From this page, select the Host target type and click **Manage Monitoring Credentials**.

For each entry with the credential "Host Credentials For Real-time Configuration Change Monitoring", select the entry and click **Set Credentials**. You will be asked for a credential set to use. Ensure you also add "root" to the Run As entry. If "Run As" is not visible, then the privilege delegation was not set properly in the previous step.

To set monitoring credentials in bulk on multiple hosts at once, you can use EMCLI. For more information on using EMCLI to set monitoring credentials, see the section, *Managing Credentials Using EMCLI* in the Security chapter of *Oracle Enterprise Manager Administration*. Likewise, for more information about configuring monitoring credentials in Cloud Control, the Security chapter of *Oracle Enterprise Manager Administration*.

Preparing To Monitor Linux Hosts

The following sections describe how to prepare Linux hosts for monitoring.

OS File Monitoring

Before using Real-time file monitoring for Linux, a loadable kernel module must be installed on the host. This loadable kernel module provides you with the most efficient way of monitoring the host. This loadable kernel module is referred to as the File Audit Module, or Audit Module for short.

Acquiring the Kernel Module

The kernel audit module is available from <http://oss.oracle.com/projects/fileauditmodule>. There are two ways to get the file audit kernel module:

1. **Prebuilt .ko files** for which Oracle has already prebuilt, you can use this in your environment. You can look for the Prebuilt kernel modules under the **Downloads** link. To find the matching prebuilt version, run the `uname -r` command on the host being monitored and compare that version to the version of the prebuilt modules. The complete version string must match perfectly. For 32-bit machines, the post-fix of the .ko file name will be .ko. For 64-bit machines, the post-fix of the .ko file name will be .k64.ko.
2. **Build your own kernel module.** To build your own kernel module, you can download the following RPM from the **Downloads** link:

Fileauditmodule-emversion-revision-noarch.rpm

You should always retrieve the latest revision available at the time you are installing this module. The emversion field must match the version of Cloud Control agent and server you are using.

Install this RPM on the host you want to monitor as root. The installation of this RPM depends on the kernel-devel package matching your running kernel also existing on the host. This kernel-devel package comes with the same media as the Linux installers.

In addition to installing this package, you must ensure that the version of gcc available on your host matches the version with which the kernel was built. To do this, view the `/proc/version` file to see what gcc version the kernel was built with and then run the command `gcc -v` to see what version of gcc is being used. These two versions should match.

Also check that the file `/boot/System.map-{version}` exists where {version} must match the kernel version you see when you run the `uname -r` command. This file contains system symbols that are required to decode the kernel symbols we are monitoring for real-time changes. Without this file, real-time file monitoring will not function. This file is standard on all default Linux installations.

After installing this package and checking prerequisites successfully, go to the directory where the package contents were installed (defaults to `/opt/fileauditmodule`) and run the following script:

```
compmod.sh
```

This will build the kernel module file (.ko, .k64, or .o extension depending on the OS version) and place it in the `/opt/fileauditmodule` directory.

If the audit module file is not created, check the *make.log* and *build.log* files for any errors in building the module.

If all of your hosts have the exact same kernel version as shown using the command `uname -r`, then you only need to compile the module on one machine. You can then copy the `.ko`, `.k64`, or `.o` file to the other servers without having to build on that specific host.

Deploying the Kernel Module

Once you have either the prebuilt `.ko` file or a `.ko` file that exists from building it from the source RPM, the `.ko` file must be located in the proper directory. The default location for this file is in the `bin` folder under the agent home directory. You can also place the file in any location on the host and change the *nmxc.properties* file under the `AGENT_INST/sysman/config` directory of the agent home. The property *nmxcf.kernel_module_dir* specifies the absolute path to the `.ko` directory.

Install Kernel Module Job

In addition to manually placing the `.KO` file on the agent, there is a Cloud Control job named *Real-time Monitoring Kernel Module Installation*. This job is configured with a list of Linux hosts on which you can install the kernel module. It will search in a directory locally on the Cloud Control server disk for prebuilt `.ko` files or the source RPM file. If it finds a matching prebuilt `.ko` file, it will send this to the matching agents; otherwise it will send the RPM to the agent and install and compile it resulting in a new `.KO` file.

Prior to using this job, files from OSS.ORACLE.COM must be manually retrieved by the user and placed into the `%ORACLE_HOME%/gcompliance/fileauditmodule/resources/linux` directory. This directory already exists on the server with a `README` file indicating this is the location to place these files. The files that must be placed here are either prebuilt `.KO` files or the source RPM file. If you have built your own `.KO` files in your environment, you can also place those `.KO` files into this directory on the server and deploy it to other hosts in your environment.

Special Considerations for Enterprise Linux 5 and Greater

For Enterprise Linux 5 and greater, the kernel audit module is not required. The monitoring will use the built-in audit subsystem if a kernel module is not detected at startup time. However, the functionality of the audit subsystem is not as robust as the capability that the kernel audit module can provide.

You will lose the functionality that provides the granularity of what type of change there has been to a file, whether it was a create action or a modify action. Without the kernel module, all changes to a file will appear as a modify action. Additionally, monitoring a directory that does not exist yet or a directory that may exist now and gets removed later may be disrupted since the underlying Linux audit subsystem does not handle these cases.

It is recommended that you use the kernel audit module even with the newer versions of Linux, if possible.

Debugging Kernel Module Or Other File Monitoring Issues

You may detect a problem with the kernel module in a few different ways:

1. You may have noticed that you do not receive real-time file changes on the Enterprise Manager console for file changes that you know should occur.
2. In the Compliance Standard Target Associations or Real-time Observations page on the user interface, you may see an agent warning indicating a kernel module problem.

3. When examining the *nmxcf.log* file under *AGENT_INST/sysman/logs*, you may see errors indicating that the kernel module could not be loaded or used for various reasons.

If you encounter any of these issues, most likely there was a problem with compiling or inserting the Linux kernel module at run time.

You can confirm whether the auditmodule was loaded properly by running the following command.

```
grep -i auditmodule /proc/modules
```

If you do not receive any output, then the auditmodule is not loaded and the agent will not perform real time file monitoring.

If the audit module file was generated properly and it does not show up in the module list above, you can attempt to manually load the module to see if there are any errors. Use the following command where you replace {audit module file name} with the entire name of the .ko file that was created from *compmo.d.sh*:

```
insmod {audit module file name}
```

If you experience no errors during this command, you can check the module list again by using the *grep* command above. If the audit module now appears, then the file monitoring capability should work. An agent restart is necessary; however there still may be a problem with the file monitoring process finding the .ko file which you will experience again next time your host is rebooted.

One additional step to debug any issues with the file monitoring process is to try to run it manually. To do this, follow these steps:

1. Get the process ID of the agent TMMain process:

```
ps -eaf |grep TMMain
```

2. Execute the *nmxcf* process using the following command replacing the values in {} with the proper path elements or the process ID from the previous command:

```
sudo {agent_home}/core/{agent_version}/bin/nmxcf -e {agent_home}/  
agent_inst/ -m {agent_home}/agent_inst/sysman/emd/state/  
fetchlet_state/CCCDDataFetchlet -w {process id of TMMAIN}
```

Running the *nmxcf* process this way will not work in the long term since it will not start up again when the agent is restarted, but this can help in trying to debug any issues as to why the process cannot start.

If the module still is not able to load and if you need to contact Oracle support about the issue, please be sure to include the following information with your support ticket:

- Output of the command: `uname -a`
- Output of the command: `grep -i auditmodule /proc/modules`
- Output of the command: `rpm -q -a |grep -i kernel-devel`
- The *make.log* and *build.log* files from the */opt/fileauditmodule* directory where you ran *compmo.d.sh* if you built your own .ko file
- The files *AGENT_INST/sysman/log/nmxc*.log*
- Any warnings or errors you received when trying to start *nmxcf* manually.

This information will help Oracle Support to determine if the real time file monitoring audit module of the agent can be built on your environment.

 **Note:**

Be careful when using the Linux OS command `rmmod` which is used to unload a kernel module. If the `nmxcf` binary is running and you use `rmmod`, there is a chance that a kernel panic can arise by trying to unload a kernel module in use. The use of `rmmod` in Linux should be done carefully no matter which module you are unloading.

Preparing To Monitor Windows Hosts

The Real-time monitoring features support Windows 2003 and 2008 Server along with Windows XP. The Real-time monitoring modules for Windows rely on various capabilities of the operating system to collect all of the information on actions. One part of this is to capture the user that made changes from the Windows Event Log. If you do not configure Windows to capture users that make changes, the agent will not capture this information. However it will still capture that a change occurred and when it occurred.

To configure the event log to work with real time monitoring, perform the following steps:

1. From Windows Explorer, select the directory that is being monitored by a Real-time Monitoring Rule, right-click and select **Properties**.
2. Go to the Security tab.
3. Click **Advanced**.
4. Select the Auditing tab.
5. Click **Add**. (In Microsoft XP, double-click the **Auditing Entries** window).
6. Select the Name **Everyone**, then click **OK**. You can also choose specific users if you are only monitoring for changes by specific users in Configuration Change Console rules. The rules filter the results by user as well, so even if you enable audit for everyone, only users that you want to monitor changes of in your rules will be captured.
7. Select the following options (Successful and/or Failed) from the Access window. For Windows XP and Windows 2003:
 - Create Files/Write Data
 - Create Folders/Append Data
 - Delete Files Subfolders and Files
 - Delete

For Windows 2008 and Windows 7:

- Create Files/Write Data
- Create Folders/Append Data
- Write Attributes
- Write Extended Attributes
- Delete Files Subfolders and Files
- Delete
- Change Permissions

- Take Ownership
8. Click **OK** to exit.
 9. Repeat steps 1 through 7 for all other monitored directories and/or files.
 10. From the **Start** menu, select **Settings**, then **Control Panel**, then **Administrative Tools**, then **Local Security Policy**, then **Local Policies**, then **Audit Policy**. Double-click and turn on the following policies (Success and/or Failure):
 - Audit account logon events
 - Audit logon events
 - Audit object access
 11. Close the Local Security Settings screen.
 12. From the **Start** menu, select **Settings**, then **Control Panel**, then **Administrative Tools**, and finally **Event Viewer**.
 13. Select **System Log**, then click **Action** from the menu bar and select **Properties**.
 14. From the System Log Properties panel, on the General tab, set the Maximum log size to at least 5120 KB (5 megabytes) and select **Overwrite Events as Needed**. Note that the log size depends on the number of events generated in the system during a two-minute reporting interval. The log size must be large enough to accommodate those events. If you extend the monitoring time for file events because you expect the change rate to be lower, you need to ensure that the audit log in Windows is large enough to capture the events.
 15. Click **Apply** then **OK** to exit.

If Windows auditing is not configured properly, you will see warnings on the Compliance Standard Target Association page on the Cloud Control user interface. This is the same page where you associated your Real-time Monitoring compliance Standards to your targets.

Verifying Auditing Is Configured Properly

To verify that the host records login and logout events, follow these steps:

1. Log out of the host and then log back into the host.
2. From **Start**, select **Settings**, then **Control Panel**, then **Administrative Tools**, and finally **Event Viewer**.
3. Select **Security Log** and choose **Filter** from the **View** menu. Select **Security for the Event Source** and **Logon/Logoff** for the Category fields.
4. Click **OK**.

The Event Viewer should have the activity recorded as Event 528.

Subinacl External Requirements

As mentioned earlier, the agent will send warnings to the server when audit settings are not set properly. It, however, can only do this if the windows feature SUBINACL is installed. If this feature is not installed on the host, a warning will be sent to the server saying that the agent cannot detect whether audit settings are correct. This warning will be visible from the Compliance Standard Associate Targets page.

You can specify the absolute path to the directory that contain subinacl by setting the following property in the *AGENT_INST/sysman/config/nmxc.properties* file:

```
nmxcf.subinacl_dir=
```

SubInACL is available for download from Microsoft's Web site.

Preparing To Monitor Solaris Hosts

Real-time monitoring on Solaris systems utilizes the Solaris audit system which is part of the Solaris Basic Security Model (BSM). BSM auditing allows system administrators to monitor events and to detect user account logins and logouts as well as file changes.

Verify that BSM auditing is enabled by running the following command with root privilege:

```
/usr/sbin/auditconfig -getcond
```

You should see the following output:

```
audit condition = auditing
```

If the output is different from the above, it means the BSM auditing needs to be enabled through different methods in different Solaris releases.

Enabling BSM Auditing

You can enable BSM auditing using the steps below for each of the following environments.

Enabling BSM Auditing Using Solaris Versions 9 and 10

To enable BSM auditing, you can use the following command with root privilege:

```
/etc/security/bsmconv
```

See the Solaris BSM Auditing manuals for additional details on setting up BSM auditing.

If auditing is already enabled on the server, simply verify that the audit system configuration matches the configurations detailed below.

The audit file can be configured to include specific events. The */etc/security/audit_control* file controls which events will be included in the audit file. This section summarizes the configuration; for further details, refer to the Sun Product Online Documentation site.

For monitoring entity types OS FILE (file changes) and OS USER (user logins/logouts), the flags line in the file */etc/security/audit_control* should be set as follows:

```
flags: +fw,+fc,+fd,+fm,+fr,+lo
```

This configuration enables success/fail auditing for file writes (fw), file creates (fc), file deletes (fd), file attribute modifies (fm), file reads (fr) and login/logout events (lo); where '+' means to only log successful events.

If you are interested in logging the failed events as well, remove the "+" sign before each event in the flag.

 **Note:**

Installing BSM on an existing host has the requirement that the host is rebooted.

Auditing Users: The `audit_user` file controls which users are being audited. The settings in this file are for specific users and override the settings in the `audit_control` file, which applies to all users.

Audit Logs and Disk Space: The `audit_control` file uses entries to control where the audit logs are stored and the maximum amount of disk space used by the audit system. The minimum requirement for file monitoring is approximately 10 minutes worth of data stored on the hard drive or the configured reporting interval time.

Enabling BSM Auditing Using Solaris 11

Auditing is enabled by default on Solaris 11, but only user login/logout events are monitored by default. For monitoring both the OS File change events and OS USER logins/logout events, you can execute the following command with root privilege:

```
/usr/sbin/auditconfig -setflags fw,fd,fc,fm,fr,lo
```

The configuration flags have the same meaning as defined in the last section.

 **Note:**

This configuration will not affect the existing sessions in which users already log into the host, so you must terminate all the existing sessions and then re-login or simply reboot the machine to ensure this change takes effect.

As the `bsmconv` command has been removed on Solaris 11, you can use the following command to enable the auditing feature, if needed:

```
audit -s
```

Managing Audit Log Files

Cloud Control Real-time Monitoring only reads the audit logs; it does not delete the logs. This might flood the system with log files and prevent it from logging additional events. To manage and delete old audit events while maintaining minimum monitoring requirements, follow these steps:

1. The auditing policy can be set to automatically drop new events (keeping only a count of the dropped events) rather than suspending all processes by running the following command:

```
# auditconfig -setpolicy cnt
```

2. Run the following command to force the audit daemon to close the current audit log file and use a new log file:

```
/usr/sbin/audit -n
```


3. Run the following command to merge all existing closed auditing log files into a single file with an extension of `.trash` and then delete the files:

```
/usr/sbin/auditreduce -D trash
```
4. Create a cron job to periodically run the commands in Step 2 and 3 above. The frequency at which these two commands are run can be adjusted based on the anticipated event volume and the amount of disk space allocated to auditing. The only requirement is that the time between the `audit -s` command and the `auditreduce -D trash` command is at least 15 minutes or twice the reporting interval if that is changed.

Preparing to Monitor AIX Hosts

Real-time monitoring on AIX systems utilizes the underlying AIX audit subsystem provided by the OS. IBM AIX 5.3 and 6.1 are the only currently supported versions.

Installation Prerequisite for AIX 5.3

Before using Real-time monitoring on AIX 5.3 hosts, ensure that you are using AIX 5.3 5300-08 service pack or higher. This maintenance package is available from IBM.

Administering AIX Auditing

The AIX auditing subsystem allows an administrator to record security-relevant information, such as User Logins, Logouts, and file changes, for analysis against existing security policies and detection of security violations.

Setting up auditing involves modification of the existing auditing configuration files. To set up auditing, follow these steps:

1. Log into the AIX machine as the root user.
2. Open a terminal window and change directory to `/etc/security/audit`
3. Open the config file in `vi`.
4. Locate the following sections, and update or add the listed values:

```
start:
    binmode = off
    streammode = on

stream:
    cmds = /etc/security/audit/ccstream

classes:
...
    filewatch = PROC_Create,PROC_Delete,FILE_Open,FILE_Write,FILE_Close,FILE_
Link,FILE_Unlink,FILE_Rename,FILE_Owner,FILE_Mode,FILE_Fchmod,FILE_Fchown,FS_Chdir,
FS_Fchdir,FS_Chroot,FS_Mkdir,FS_Rmdir,FILE_Symlink,FILE_Dupfd,FILE_Mknod,FILE_Utime
s

users:
    root = filewatch
    default = filewatch
```

 **Note:**

- In this case, `default` refers to all users that are not root. Further note that the last line of the config file should be a blank line.
- Each parameter (`binmode`, `streammode`, `filewatch`, `root`, and `default`) must have a tab in front of them. You can verify that the audit system has used all variables properly by using the `audit query` command. Make sure the `filewatch` property appears in the output.

5. Save your modifications and exit vi.
6. In the same directory (`/etc/security/audit/`) open the file `streamcmds` in vi.
7. Clear all text from the file. The default configuration for this file is not necessary, as the File Monitoring agent module (`nmxcf` process) will operate as a direct audit reader. Clearing the file helps to reduce CPU usage and improve overall auditing performance.
8. Save the file and exit vi.
9. At the terminal prompt, enter the following command to initialize Auditing at system startup:

```
mkkitab "audit:2:once:/usr/sbin/audit start"
```
10. At the prompt, restart audit using the command `/usr/sbin/audit shutdown` and `/usr/sbin/audit start` or directly reboot the host to make the auditing effective.
11. At the prompt, use the command `audit query` to view the configuration the audit system is using. Ensure that the properties are set correctly and that the required settings for `filewatch` are set.

Verifying AIX System Log Files for the OS User Monitoring Module

The OS User monitoring module relies on the following system log files:

- `/etc/security/failedlogin`
- `/var/adm/wtmp`
- `/var/adm/sulog`

Be sure the log files exist before running the OS User monitoring module on an AIX host. If any of the log files is missing, refer to the AIX System documentation for more information about how to generate it.

Preparing To Monitor the Oracle Database

This section describes the steps involved in setting up auditing within an Oracle database. If you are going to monitor an Oracle database with any of the Oracle entity types, you will need to perform these steps before events will be captured.

Before configuring auditing it is suggested you review the Auditing Database Use section of the *Oracle Database Administrator's Guide*. This document provides an overview of Oracle's auditing functionality, as well as basic concepts and guidelines for auditing configurations. Note that this document does not cover all details of

configuring and fine tuning the Oracle audit system. Instead, this document serves as an example of the basic steps involved to configure the Oracle audit system to enable Real-time monitoring through Real-time monitoring rules.

Setting Auditing User Privileges

When you create a Real-time Monitoring Compliance Standard Rule to monitor an Oracle instance target, the agent will read the audit trail to perform its monitoring.

Real-time monitoring for Oracle entity types requires the audit trail to be stored in the database as opposed to a file. To verify if a setting is correct, follow these steps:

1. In Cloud Control, go to the target home page for the Oracle Database target for which you want to enable Real-time Monitoring.
2. From the **Administration** menu, select **Initialization Parameters**.
3. Log in to the database as a sys user, connecting as SYSDBA.
4. Find the parameter *audit_trail* and ensure it is set to DB. If not, this parameter needs to be changed in the Oracle Database.
5. This change will require a restart of the database.

Specifying Audit Options

Through SQL plus, an Oracle DBA can use audit and noaudit statements to configure audit options for the database. The audit statement allows you to set audit options at three levels:

Table 4-1 Audit Options Table

Level	Effect
Statement	Audits specific SQL statements or groups of statements that affect a particular type of database object. For example, AUDIT TABLE audits the CREATE TABLE, TRUNCATE TABLE, COMMENT ON TABLE, and DELETE [FROM] TABLE statements.
Privilege	Audits SQL statements that are executed under the umbrella of a specified system privilege. For Example, AUDIT CREATE ANY TRIGGER audits statements issued using the CREATE ANY TRIGGER system privilege.
Object	Audits specific statements on specific objects, such as ALTER TABLE on the employee table

To use the audit statement to set statement and privilege auditing options a DBA must be assigned AUDIT SYSTEM privileges. To use the audit statement to set object audit options, the DBA must own the object to be audited or be assigned the AUDIT ANY privilege within Oracle. Privilege assignments are covered in the following section.

Audit statements that set statement and privilege audit options can also include a BY clause to supply a list of specific users or application proxies to audit, and thus limit the scope of the statement and privilege audit options.

Some examples of audit statements can be seen below. Feel free to use these as a basis for the audit settings you specify within your database. Once all audit settings are in place you can create application policies, using the Oracle (SQL Trace) agent module with which to monitor the Oracle database instance.

Statement Audit Options (User sessions)

The following statement audits user sessions of users Bill and Lori.

```
AUDIT SESSION BY scott, lori;
```

Privilege Audit Options

The following statement audits all successful and unsuccessful uses of the DELETE ANY TABLE system privilege:

```
AUDIT DELETE ANY TABLE BY ACCESS WHENEVER NOT SUCCESSFUL;
```

Object Audit Options

The following statement audits all successful SELECT, INSERT, and DELETE statements on the dept table owned by user jward:

```
AUDIT SELECT, INSERT, DELETE ON jward.dept BY ACCESS WHENEVER SUCCESSFUL;
```

Example Oracle Audit Monitor Configurations

The following command audits all basic statements. Extra statements are not audited.

```
Audit all by access;
```

The following statement audits all extra statements:

```
audit ALTER SEQUENCE, ALTER TABLE, DELETE TABLE, EXECUTE PROCEDURE, GRANT
DIRECTORY, GRANT PROCEDURE, GRANT SEQUENCE, GRANT TABLE, GRANT TYPE,
INSERT TABLE, LOCK TABLE, UPDATE TABLE by access;
```

The following command displays audit settings for statements:

```
SELECT * FROM DBA_STMT_AUDIT_OPTS;
```

Once you have specified your audit configuration you can then create real-time monitoring rules from the Cloud Control Server that uses the Oracle Database entity types.

Oracle Database Table Monitoring

The following table displays the platforms that support Oracle Database Table Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-2 Oracle Database Table Monitoring

Actions to Monitor	Oracle Database		
	10g	11g	12g
Insert (successful)	X	X	X
Select (successful)	X	X	X
Update (successful)	X	X	X
Delete (successful)	X	X	X
Create (successful)	X	X	X
Drop (successful)	X	X	X

Table 4-2 (Cont.) Oracle Database Table Monitoring

Actions to Monitor	Oracle Database		
	10g	11g	12g
Truncate (successful)	X	X	X
Alter (successful)	X	X	X
Comment (successful)	X	X	X
Rename (successful)	X	X	X
Lock (successful)	X	X	X
Grant (successful)	X	X	X
Revoke (successful)	X	X	X
Audit (successful)	X	X	X
NOAUDIT usage	X	X	X
Flashback (successful)	X	X	X

Setting Up Change Request Management Integration

This section explains how to install and configure integration with a Change Management Server and to be able to determine whether changes that occur are authorized automatically.

- [BMC Remedy Action Request System 7.1 Integration](#)

BMC Remedy Action Request System 7.1 Integration

Remedy ARS 7.1 is a supported Change Management system for automatic reconciliation of observations. The following steps outline how to setup Remedy and also configure the integration with Cloud Control.

- [Installing and Customizing Remedy ARS](#)

Installing and Customizing Remedy ARS

Follow these steps to install and customize Remedy ARS 7.1.

1. Install Remedy ARS 7.1. Ensure the following components are all installed and properly licensed:
 - ARS 7.1.00 Patch 011
 - Midtier 7.1.00 Patch 011
 - Flashboard Server 7.0.03
 - Assignment Engine 7.1
 - Asset Management 7.0.03*
 - CMDB 2.1.00 Patch 4
 - CMDB Extension Loader
 - Approval Server 7.1
 - Change Management Server 7.0.03 Patch 008*

Problem Management Server 7.0.03*

Incident Management Server 7.0.0*3

User Client

Administrator Client

These packages all come with the IT Service Management Pack. Oracle provides example customizations for the Remedy under ITSM 7.0.03 Patch 008 environment. For different versions, the customizations may need to be adjusted to account for changes in the version of Remedy.

2. Install the Cloud Control EMCLI_Client on the same host on which Remedy is installed. This will need to be able to communicate to your Cloud Control Server.
 - a. Log in to the Enterprise Manager console.
 - b. Choose **Setup**, then select **Command Line Interface** from the **My Preferences** menu.
 - c. Click **Download the EM CLI kit to your workstation** and download the jar to your Remedy server.
 - d. Follow the steps given on the page to install the EMCLI client on the Remedy server.
3. Get the latest version of the Change Request Management connector self-update package. Also acquire the latest version of the example Remedy ARS customizations for Cloud Control version 12c.

These definition files provide a guideline of customizations that must be made in your environment for the integration. These customization files assume a fresh install of Remedy ARS. When integrating with a production instance of Remedy, care should be taken to make sure these customizations are compatible with any previous customizations that have been made to the Remedy instance.

- ActiveLinks_Customization.def
- Forms_Customization.def
- Menus_Customization.def
- Webservices_Customization.def

To get these definition files, in the Enterprise Manager Self Update user interface, export the connector. The definition files are inside this connector package.

4. Install the four definition files (.DEF) files in the running Remedy environment by completing these steps:
 - a. Log into the Remedy Administrator tool.
 - b. Select the **Remedy** instance from the hierarchy on the left.
 - c. From the **Tools** menu, select **Import Definitions**, then select **From Definition File...**
 - d. Select the definition file to import from the list above.
 - e. Check the box labeled **Replace Objects on the Destination Server**.
 - f. Choose the drop down option **Replace With New Type**.
 - g. Click **Import**.
 - h. You should not encounter any errors during this process. At the end of import there should be an Import Complete message.

- i. When done, repeat for the rest of the customization files.
5. Customize Web Services.
 - a. Log into Remedy Administrator tool.
 - b. Select **Webservices**, then select the webservice **EMCCC_GetCR**. Right click, then select **Open**.
 - c. Select the **WSDL** tab.
 - d. In the input on top, modify the `midtier_server` and `servername` values in the **WSDL Handler URL**.
 - e. If `midtier` is on localhost, you can enter localhost right after `http://`.
 - f. If the `midtier` uses port 80, you can omit the port, otherwise include the port after the server name.
 - g. For the `servername` after "`public/`", enter the name of the Remedy server.
 - h. Click **View**.
 - i. You should see an XML representing the webservice WSDL file for the webservice.
 - j. If you see an error, check the `midtier_server` name, port, or `servername`. Also, you can try adding/removing the domain part of the `servername`. Another possible issue occurs when the `midtier` password set in Remedy's System > General > Serverinfo > Connection Settings may not be set correctly. Be sure to check this also if the WSDL XML is not returned.
 - k. If you see the XML content after clicking View, then close this window and save the changes.
 - l. Repeat all above steps with the webservices `EMCCC_PublishCSDData` and `EMCCC_UpdateChangeRequest`.
6. Customize Active Links.
 - a. Log in to Remedy Administrator tool.
 - b. Select active links and then select the active link **EMCCCC_ApprovedCR**. Right click, then select **Open**.
 - c. Click the **If Action** tab.
 - d. Click the Current Action **Run Process** at the end of the list of actions.
 - e. In the Command Line field, change the path to `emcli.bat` to match that of where you installed the `emcli` on the local host.
7. Create a user in Remedy that will be used for creating requests that will be used for automatic observation reconciliation:
 - a. Log in to BMC Remedy User Client as an administrative user.
 - b. Click **Application Administration Console** on the User Client Home Page.
 - c. Click **Create** for each step 1 through 4 in this wizard.
 - d. When adding the person, add the support group under the Support Groups tab.
 - e. Under the Support Groups Tab, select sub tab **Support Group Functional Roles**.
 - f. Add Support groups with functional role of *Infrastructure Change Management*. Without this, you will not be able to create change requests as the Infrastructure Change Manager fields support group will not have values.
 - g. Go to AR System Administrator Console.

- h. On the left side bar, select **Application**, then **Users/Groups/Roles**, then **Select Users**.
- i. This will load the user search page. Click **Search** at the top right.
- j. Double-click the newly created user above to bring up the user form.
- k. Click the down arrow next to "Group List" field and select **Infrastructure Change Master**.
- l. Repeat the previous step and add the following Groups to this user as well.
 - Infrastructure Change Submit
 - Infrastructure Change User
 - Infrastructure Change Viewer
- m. Save the changes to this user by clicking the **Save** button in the upper right hand corner of the window.

Adding the Connector to Cloud Control

Follow these steps to add the connector to Cloud Control.

1. Add the Change Management Connector to Cloud Control.
 - a. Log into Enterprise Manager as an Administrative user that has privileges to create a connector.
 - b. From the **Setup** menu, select **Provisioning and Patching**, then choose **Software Library**.
 - c. Click **Actions**, then select **Administration**.
 - d. Click **Add**.
 - e. Provide a name, such as "self update swlib".
 - f. Provide a location where the swlib files will be located on the Cloud Control server. This can be anywhere, but must be a path that the Cloud Control user can access. You must put the full absolute path in this input.
 - g. This process will take several minutes to complete.
 - h. Locate the connector self-update package file.

The connectors jar can be downloaded from the Cloud Control store to EM@Customer using the Self Update console, and can be exported to any local directory using the export functionality of Self Update.
 - i. Run: `emcli import_update -file=<full path>/connector.zip -omslocal` (where *connector.zip* is an example name of the self update package)
 - j. If you have errors with the previous step, make sure the user you run emcli as has permissions to access this directory and file. Also, be sure you are using absolute path for the *-file* switch.
 - k. When successful, you will receive the following message:

Operation completed successfully. Update has been uploaded to Cloud Control. Please use the Self Update Home to manage this update.
 - l. Log into the Enterprise Manager console.
 - m. From the **Setup** menu, select **Extensibility**, then select **Self Update**.

- n. Find the type "Management Connector" and click the link "1" under "Downloaded Updates" for this entry.
- o. Select the Connector from the table and click **Apply**.
2. Create a Change Management Connector instance.
 - a. Log in into Enterprise Manager console.
 - b. From the **Setup** menu, select **Extensibility**, then select **Management Connectors**.
 - c. Select "Remedy Change Management Connector" from the drop-down after "Create Connector", then click **Go**.
 - d. Provide a name and description for the connector. This name is used to choose the connector when creating a Real-time Monitoring Compliance Standard Rule.
 - e. After returning to the management connector listing page, select the newly added row, then click **Configure**.
 - f. Under the Web Service End Points label, change the [servername] and [port] to match that of your Remedy instance Web Services. The values you put here will be similar to what you configured in the Web Services step earlier in these instructions.
 - g. Enter the Remedy username and password you are using for the connector integration.
 - h. Enter the locale ('en', for example).
 - i. Enter the time zone offset of the remedy server from UTC, ('-08:00', for example).
 - j. Enter the Change ID to use as a test. This should be a valid Change Request ID currently existing in Remedy that is used to test the connectivity between Cloud Control and Remedy.

Using Automatic Reconciliation Rules

Once Remedy is customized and the Cloud Control connector is configured, to utilize the automatic reconciliation features you need to create Real-time Monitoring Rules that are configured to use automatic reconciliation. Use the following steps:

1. Create a Real-time monitoring Rule:
 - a. Follow the normal steps to create a Real-time monitoring Rule.
 - b. On the Settings page, choose **Authorized Observations Automatically** using Change Request Management System. This configures Cloud Control to use this change request from Remedy for reconciliation of Real-time Observations that are detected.
 - c. Select the connector from the drop-down.
 - d. Click to annotate change requests with authorized observations check box.
 - e. Continue to save the rule after this. The Real-time Monitoring Rule can be used like any other Real-time Monitoring rule. The integration with a new Change Management server will not begin until at least one Real-time Monitoring Standard with a rule using Automatic Reconciliation is associated to a target. Create a Compliance Standard, add this rule to the Compliance Standard, and associate this compliance standard to one or more targets.

For more information on the configuration of rules, see Managing Compliance in *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

Creating Change Requests for Upcoming Changes

Now that integration is set up and Real-time monitoring rules have been created, Change Requests can be created by Remedy users in the Remedy interface. These Change Requests will be compared to observations that occur to automatically determine if these observations are from actions that were authorized by change requests or not.

To make this correlation, some new fields that have been added to the Change Request form must be filled out by the change request filer. Not all fields are required; correlation only occurs on the fields that are present in the Change Request.

For instance, the following fields have been added to the Change Request form under the Oracle Enterprise Manager Integration tab:

- **Connector:** Choose the Cloud Control connector this Change Request will use to integrate with Cloud Control.
- **Hostname:** the hostname(s) this change request is for. These are the hosts that this change request is specifying someone needs to make changes to. An empty value in this field indicates that all hosts will be correlated to this change request.
- **Target User List:** the user name(s) this change request is for based on target users. These are the target users you expect to log in to the target to make a change. An empty value in this field means that all users on the target will be correlated to this change request.
- **Target Type:** the target type this change request is against. An empty value in this field means that any target type will be correlated to this change request.
- **Target:** The target this change request is specifically for. An empty value in this field indicates that any target will be correlated to this change request.
- **Facet:** The facet this change request is specifically for. An empty value in this field indicates that all facets on the above target type and target will be correlated for this change request.

When creating a change request that you want to use to authorize changes detected by Real-time monitoring rules, follow these steps in addition to whatever requirements your organization implements for creation of Change Requests:

1. Under the Dates tab of the Change Request form, fill out the Scheduled Start date and Scheduled End Date. These are the date ranges the request is valid for reconciliation. If an action occurs outside this time, it is marked as unauthorized by the Real-time Monitoring feature.
2. Select the Oracle Enterprise Manager **Integration** tab.
3. Select the Cloud Control connector from the drop-down list.
4. Optionally select values for the five reconciliation criteria as described above: Hostname, Target User List, Target type, Target and Facet. The last three -- Target Type, Target, and Facet -- will be Choice lists based on content in Real-time Monitoring Rules that have been created in Cloud Control that belong to Compliance Standards which are associated to targets. You can add multiple values separated by commas.

 **Note:**

This form can be customized in Remedy to look differently. The example form elements from the customizations loaded earlier are only examples.

5. Change the auditable status to True. This configures Remedy to allow Cloud Control to use this change request for reconciliation of Real-time Observations that are detected.
6. Save the change request.
7. A popup displays, notifying you that active links will send the content to Cloud Control. You will see a DOS command window open and then close.

Overview of Reconciliation Functionality

After creating a change request that references a target and/or facet that is being monitored by Real-time Monitoring rules, any observations that happen against that rule will be correlated to all open and matching change requests.

When the observation arrives at the Cloud Control server, all open change requests that were active (based on Scheduled Start/Stop time) and have matching correlation criteria from the Cloud Control Integration tab will be evaluated. If any change request exists that matches the criteria of the observation, this observation will be marked with an "authorized" audit status. If the annotation check box was checked in the Rule configuration, details of these authorized observations will be put into a table in the Enterprise Manager Integration tab of the Remedy Change Request.

If no open change requests can be correlated to the observation and the rule was configured to use automatic reconciliation, then this observation is set to an Unauthorized audit status. The Observation bundle to which this observation belonged will be in violation and results in a Cloud Control event being created. This event can further be used through creation of a Cloud Control Event Rule.

An observations audit status can be seen whenever looking at observation details either by selecting Compliance, then Real-time Observations, then Observation Search, or either of the Browse By screens. A user with the proper role can also override the audit status for individual observations from these pages.

Any bundles that are in violation because they contain unauthorized observations will be reflected as violations in the Compliance Results page. These violations cause the compliance score skew lower. If these violations are cleared, the score becomes higher; however, the history of these audit status changes will be retained for the given observation.

Overview of the Repository Views Related to Real-time Monitoring Features

The following views exist to allow access to Real-time Monitoring data.

View: mgmt\$ccc_all_observations

Description: This view returns all observations that have occurred. Any query against this view should ensure that filtering is done on appropriate fields with *action_time* being the first to take advantage of partitions.

Fields:

Field	Description
OBSERVATION_ID	Unique ID given to the observation when detected by the agent
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings

Field	Description
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
ENTITY_TYPE	Entity type of the entity that had an action against it
ACTION	Action that was observed
ACTION_TIME	Time the action occurred
USER_TYPE	Type of user that performed the action (for example, OS user versus DB user)
USER_PERFORMING_ACTION	Name of the user that performed the action
ORIGINAL_USER_NAME	Previous user name in the case of a SU/SUDO action (only applicable to some entity types)
AFFECTED_ENTITY_NAME	Name of the entity that was affected by this action (file name, and so on)
AFFECTED_ENTITY_PREVIOUS_NAME	Name of the entity prior to the action. For instance for file rename actions, this would be the old file name.
SOURCE_HOST_IP	Source IP of a connection when an action comes from another host (only applicable to some entity types)
ACTION_PROCESS_ID	PID of the process that performed the action (only applicable to some entity types)
ACTION_PROCESS_NAME	Name of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_ID	PID of the parent process of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_NAME	Name of the parent process of the process that performed the action (only applicable to some entity types)
ENTITY_PREVIOUS_VALUE	Previous value of the entity (only applicable to some entity types)
ENTITY_NEW_VALUE	New value of the entity (only applicable to some entity types)
FILE_ENTITY_PREVIOUS_MD5_HASH	Previous MD5 hash value of the entity (only applicable to some entity types)
FILE_ENTITY_NEW_MD5_HASH	New MD5 hash value of the entity (only applicable to some entity types)
AUDIT_STATUS	Current audit status of the observation (unaudited, authorized, unauthorized, and so on)
AUDIT_STATUS_SET_DATE	Date the most recent audit status was set
AUDIT_STATUS_SET_BY_USER	User who set the most recent audit status

View: mgmt\$ccc_all_obs_bundles

Description: This view returns a summary of all observations bundles. Any query against this view should ensure that filtering is done on appropriate fields with *bundle_start_time* being the first to take advantage of partitions.

Fields:

Field	Description
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
RULE_NAME	Name of the Real-time Monitoring Compliance Standard Rule
ENTITY_TYPE	Entity type of the entity that had an action against it
USER_PERFORMING_ACTION	Name of the user that performed the action
BUNDLE_IN_VIOLATION	Boolean value if the bundle currently is in violation. This means at least one observation in the bundle is unauthorized. True indicates the bundle is in violation.
BUNDLE_START_TIME	Date of the first observation in this bundle
BUNDLE_CLOSE_TIME	Date when this bundle was closed
BUNDLE_CLOSE_REASON	Explanation of why this bundle was closed
DISTINCT_OBS_COUNT	Total number of observations in this bundle
AUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently authorized
UNAUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently unauthorized
UNAUTH_CLEARED_OBS_COUNT	Number of observations in this bundle that are currently cleared (that were at one point unauthorized)
UNAUDITED_OBS_COUNT	Number of observations in this bundle that are currently unaudited. They have not been evaluated manually or with Change Management integration to determine audit status.

View: mgmt\$ccc_all_violations

Description: This view returns all real-time monitoring violations caused by an observation bundle having at least one unauthorized observation in it.

Fields:

Field	Description
ROOT_CS_ID	Root Compliance Standard GUID. This is used for internal representation of the violation context.
RQS_ID	Runtime compliance standard GUID. This is used for internal representation of the violation context.
RULE_ID	Rule GUID. Internal ID of the rule having a violation.
TARGET_ID	Target GUID. Internal ID of the target having a violation.
ROOT_TARGET_ID	Root Target GUID. Internal ID of target hierarchy.
RULE_TYPE	Type of rule (Repository, Weblogic Server Signature, Real-time Monitoring)
SEVERITY	Severity Level of the rule (Info, Warning, Critical)

Field	Description
BUNDLE_ID	Internal ID of the Observation Bundle that is in violation. This observation bundle has one or more unauthorized observations in it
BUNDLE_START_TIME	Time the Observation Bundle started
BUNDLE_CLOSE_TIME	Time the Observation Bundle closed
TARGET_TYPE	Target Type of the Observation Bundle and all observations inside that bundle.
ENTITY_TYPE	Entity Type of the Observation Bundle and all observations inside that bundle.
USER_NAME	User name that performed the actions in this bundle
AUTHORIZED_OBS_COUNT	Number of Authorized observations in the observation bundle involved in this violation.
UNAUTHORIZED_OBS_COUNT	Number of Unauthorized observations in the observation bundle involved in this violation.
UNAUDITED_OBS_COUNT	Number of unaudited observations in the observation bundle involved in this violation.
RULE_NAME	Rule Name this violation is against.
COMPLIANCE_STANDARD_NAME	Compliance Standard Name this violation is against.
TARGET	Target Name this violation is against.

View: mgmt\$compliant_targets

Description: This view returns all evaluation and violation details for all targets. This is the same data that is shown in the Compliance Summary dashboard regions for targets.

Fields:

Field	Description
TARGET_ID	Internal representation of the Target
TARGET_NAME	Name of the Target
TARGET_TYPE	Target Type of the Target
TARGET_TYPE_INAME	Internal representation of the Target Type
CRIT_EVALS	Number of Critical-level Evaluations
WARN_EVALS	Number of Warning-level Evaluations
COMPLIANT_EVALS	Number of Compliant Evaluations
CRIT_VIOLATIONS	Number of Critical-level Violations
WARN_VIOLATIONS	Number of Warning-level Violations
MWARN_VIOLATIONS	Number of Minor Warning-level Violations
COMPLIANCE_SCORE	Current Compliance Score for the target

View: mgmt\$compliance_summary

Description: This view returns all evaluation and violation details for Compliance Standards and Frameworks. This is the same data that is shown in the Compliance Summary dashboard regions for Standards and Frameworks.

Fields:

Field	Description
ELEMENT_NAME	Display name of the Compliance Standard or Compliance Framework
ELEMENT_ID	Internal ID of the compliance standard or compliance framework
FRAMEWORK_ID	Internal ID of the Compliance Framework
CRIT_EVALS	Number of Critical-level Evaluations
WARN_EVALS	Number of Warning-level Evaluations
COMPLIANT_EVALS	Number of Compliant Evaluations
CRIT_VIOLATIONS	Number of Critical-level Violations
WARN_VIOLATIONS	Number of Warning-level Violations
MWARN_VIOLATIONS	Number of Minor Warning-level Violations
COMPLIANCE_SCORE	Current compliance score for the standard or framework
NON_COMPLIANT_SCORE	Current non-compliant score for the standard or framework
ELEMENT_TYPE	Type of element (1=Compliance Standard, 4=Compliance Framework)
AUTHOR	Author of the standard or framework
VERSION	Version of the standard or framework
ELEMENT_INAME	Internal representation of the standard or framework

View: mgmt\$compliance_trend

Description: This view returns the last 31 days compliance trend information for compliance frameworks and standards. This is the same data that is shown in the Compliance Summary dashboard trend regions for Standards and Frameworks.

Fields:

Field	Description
ELEMENT_ID	Internal ID representation of the standard or framework
FRAMEWORK_ID	Internal ID representation of the compliance framework
ELEMENT_NAME	Display name of the Compliance Standard or Compliance Framework
ELEMENT_INAME	Internal representation of the standard or framework
AVG_COMPLIANCE_SCORE	Average compliance score over last 31 days
DAILY_AVG_VIOLATIONS	Average number of violations per day over last 31 days
SNAPSHOT_TS	The snapshot timestamp
TOTAL_EVALS	Total evaluations over last 31 days
ELEMENT_TYPE	Type of element (1=Compliance Standard, 4=Compliance Framework)

Modifying Data Retention Periods

Real-time Monitoring features use partitioning and data retention configuration.

The following are the tables along with their default retention periods. When changing any retention periods, all tables related to Real-time monitoring must be changed to the same value to ensure that data is consistent across various features.



Note:

For more information about modifying data retention values, see the chapter "Maintaining and Troubleshooting the Management Repository" in the book *Oracle Enterprise Manager Administration*.

Table Name	Default Retention Period	Description
EM_CCC_WATCHDOG_ALERTS	366 Days	This table stores warnings from the agents when we detect that monitoring was not active.
EM_CCC_HISTORY_JOBEXEC	366 Days	This table stores history of all Enterprise Manager Jobs that are run as part of the Real-time Monitoring functionality.
EM_CCC_OBSERVATION	366 Days	This table stores each individual observation of a user action (for example, each file change, login/logout, process start/stop, each database object change, and so on).
EM_CCC_OBSGROUP	366 Days	This table stores information about how a single observation is related to a bundle based on the bundle settings set in the Real-time Monitoring Rule's user interface.
EM_CCC_OBS_GROUP_MAP	366 Days	This table stores the relationship between each single observation bundle and the target, rule, and standard that was monitoring for that observed action.
EM_CCC_HISTORY_OBS_STATUS	366 Days	This table stores the state change history for audit status (unaudited, unauthorized, authorized) for each observation.
EM_CCC_HA_OBS	366 Days	This table stores analytic summaries of counts of observations by hour and other attributes for reporting.
EM_CCC_HA_OBSGROUP	366 Days	This table stores analytic summaries of counts of observations bundles by hour and other attributes for reporting.
EM_CCC_FILEOBS_DIFF	366 Days	This table stores past file comparison for OS File based observations.
EM_CCC_AUTHOBS_CR_MAP	366 Days	This table stores the mapping between Change Management Request System change requests that were used to authorize an observation.

Table Name	Default Retention Period	Description
EM_CCC_CMPUBACTION	366	This table stores requests to publish data from EM server to an integrated Change Management Server using the connector.
EM_CCC_CMPUBACTION_DETAIL	366	This table stores additional details for requests to publish data from EM server to an integrated Change Management Server using the connector.

Real-time Monitoring Supported Platforms

The following tables display the various platforms that support Real-time monitoring. For all tables, an X indicates support for the listed action and NS indicates "Not Supported".

The following Operating System platform combinations are not supported at this time:

- Microsoft Windows -- IA64
- Any Linux -- IA64, PA-RISC, POWER

OS User Monitoring

The following table displays the platforms that support OS User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-3 OS User Monitoring

Actions to Monitor	Oracle/Redhat Linux					Windows					
	V4		V5		V6	XP		2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit
Telnet Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Telnet Logout (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Telnet Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Logout (Successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Console Login (successful)	X	X	X	X	X	X	X	X	X	X	X
Console Logout (successful)	X	X	X	X	X	X	X	X	X	X	X
Console Login (failed)	X	X	X	X	X	X	X	X	X	X	X

Table 4-3 (Cont.) OS User Monitoring

Actions to Monitor	Oracle/Redhat Linux					Windows					
	V4		V5		V6	XP		2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit
FTP Login (successful)	NS	NS	NS	X	X	NS	NS	NS	NS	NS	NS
FTP Logout (successful)	NS	NS	NS	X	X	NS	NS	NS	NS	NS	NS
FTP Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Logout (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SUDO (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SUDO (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
RDP Login (Successful)	NS	NS	NS	NS	NS	X	X	X	X	X	X
RDP Logout (Successful)	NS	NS	NS	NS	NS	X	X	X	X	X	X
RDP Login (failed)	NS	NS	NS	NS	NS	X	X	X	X	X	X

Table 4-4 OS User Monitoring

Actions to Monitor	SUSE Linux			Solaris				AIX			
	V10		V11	V9		V10		V11		V 5.3	V 6.1
	X86 32 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sparc	X86 64 bit	Sparc	X86 64	Sparc	POWER	POWER
Telnet Login (successful)	X	X	X	X	X	X	X	X	X	X	X
Telnet Logout (successful)	X	X	X	X	X	X	X	X	X	X	X
Telnet Login (failed)	X	X	X	X	X	X	X	X	X	X	X
SSH Login (successful)	X	X	X	X	X	X	X	X	X	X	X
SSH Logout (Successful)	X	X	X	X	X	X	X	X	X	X	X
SSH Login (failed)	X	X	X	X	X	X	X	X	X	X	X
Console Login (successful)	NS	X	X	X	X	X	X	X	X	NS	NS

Table 4-4 (Cont.) OS User Monitoring

Actions to Monitor	SUSE Linux			Solaris					AIX			
	V10		V11	V9		V10		V11		V 5.3		V 6.1
	X86 32 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sparc	X86 64 bit	Sparc	X86 64 bit	Sparc	POWER	POWER	
Console Logout (successful)	NS	X	X	X	X	X	X	X	X	NS	NS	
Console Login (failed)	NS	X	X	X	X	X	X	X	X	NS	NS	
FTP Login (successful)	X	NS	NS	X	X	X	X	X	X	X	X	
FTP Logout (successful)	NS	NS	NS	X	X	X	X	X	X	X	X	
FTP Login (failed)	X	X	X	X	X	X	X	X	X	X	X	
SU Login (successful)	X	X	X	X	X	X	X	X	X	X	X	
SU Logout (successful)	NS	X	X	NS	NS	NS	NS	X	X	NS	NS	
SU Login (failed)	X	X	X	X	X	X	X	X	X	X	X	
SUDO (successful)	X	X	X	NS	NS	NS	NS	NS	NS	NS	NS	
SUDO (failed)	X	X	X	NS	NS	NS	NS	NS	NS	NS	NS	
RDP Login (Successful)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	
RDP Logout (Successful)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	
RDP Login (failed)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	

OS Process Monitoring

The following table displays the platforms that support OS User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-5 OS Process Monitoring

Actions to Monitor	Oracle/Redhat Linux					Windows					Solaris							
	V4		V5		V6	XP		2003 Server			2008 Server (R1 and R2)		V9		V10		V11	
	X86 32 bit	X86 32 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	X86 64 bit	Sparc	X86 64 bit	Sparc	X86 64 bit	Sparc	X86 64 bit	Sparc
Process Start (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Table 4-5 (Cont.) OS Process Monitoring

Actions to Monitor	Oracle/Redhat Linux			Windows					Solaris							
	V4	V5	V6	XP	2003 Server		2008 Server (R1 and R2)		V9	V10	V11					
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	Spa rc	X86 64 bit	Sparc	X86 64 bit	Sp arc	
Process Stop (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Table 4-6 OS Process Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
Process Start (successful)	X	X	X	X	X
Process Stop (successful)	X	X	X	X	X

OS File Monitoring

For Linux v5, there are two possible ways monitoring can occur. Some actions to monitor below will work only on one or the other method. The two methods are to use the Loadable Kernel Module. Actions that are detectable ONLY with this method are annotated with "(KO)". The other option is to not use the loadable kernel module, which will result in using the Linux built-in audited method. The actions that can only be monitored using this method are annotated with "(non-KO)". The actions that have no annotation other than the check mark can be monitored using either approach.



Note:

Monitoring remote file systems on Unix-based platforms is not supported. Likewise, monitoring remote file systems on Windows platforms is also not supported.

When restoring a file from the Recycle Bin on the Microsoft Windows operating system, capturing the user that made the change is not available since that feature is not available from the Operating System.

When using the audited monitoring method on Linux operating systems, not the Oracle kernel audit module method, directory creations are reported as file creation. Additionally, file create activity will be reported as a file modification instead of create. These are limitations of using the audited method of monitoring. If you use the Oracle kernel audit module approach for OS file monitoring on Linux, these limitations will not exist.

An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-7 OS File Monitoring

Actions to Monitor	Linux			Windows						Solaris							
	V4	V5	V6	XP	2003 Server			2008 Server (R1 and R2)			V9	V10		V11			
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sp arc	X86 64 bit	Spa rc	X86 64 bit	Spa rc
File Read (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Delete (Successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Rename (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Create (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Content Modified (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Modified without content change	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Modified (failed)	NS	X (No n-KO)	NS (No n-KO)	X	X	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
File Permission Change (successful)	NS	X (non-KO)	X (no n-KO)	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File Ownership Change (successful)	NS	X (non-KO)	X (no n-KO)	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File content modified (successful) Archive File	NS	X (non-KO)	X (no n-KO)	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Read (failed)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File Delete (failed)	NS	X (No n-KO)	X (No n-KO)	NS	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X

Table 4-7 (Cont.) OS File Monitoring

Actions to Monitor	Linux			Windows						Solaris							
	V4	V5	V6	XP	2003 Server			2008 Server (R1 and R2)			V9	V10		V11			
	X86 32 bit	X86 32 bit	X86 64 bit	X8 32 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sp arc	X8 64 bit	Spa rc	X8 64 bit	Spa rc	
File Rename (failed)	NS	X (No n-KO)	X (No n-KO)	X (no n-KO)	X (non -KO)	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File Create (failed)	NS	X (non -KO)	X (no n-KO)	X (no n-KO)	X (non -KO)	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File Permission Change (Failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X
File Ownership Change (failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X

Table 4-8 OS File Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
File Read (successful)	X	X (KO)	X (KO)	X	X
File Delete (Successful)	X	X (KO)	X (KO)	X	X
File Rename (successful)	X	X	X	X	X
File Create (successful)	X	X	X	X	X
File Content Modified (successful)	X	X	X	X	X
File Modified without content change (successful)	X	X	X	X	X
File Modified (failed)	NS	NS	NS	X	X
File Permission Change (successful)	X	X (KO)	X	X	X
File Ownership Change (successful)	X	X (KO)	X	X	X
File content modified (successful) Archive File	X	X	X	X	X
File Read (failed)	NS	NS	NS	X	X
File Delete (failed)	NS	NS	NS	X	X

Table 4-8 (Cont.) OS File Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
File Rename (failed)	NS	X (Non-KO)	X (Non-KO)	X	X
File Create (failed)	NS	NS	X (Non-KO)	X	X
File Permission Change (Failed)	NS	X (Non-KO)	X (Non-KO)	X	X
File Ownership Change (failed)	NS	X (Non-KO)	X (Non-KO)	X	X

OS Windows Registry Monitoring

The following table displays the platforms that support OS Windows Registry Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-9 OS Windows Registry Monitoring

Actions to Monitor	Windows					
	XP		2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit
Create Key (successful)	X	NS	X	X	X	X
Delete Key (successful)	X	NS	X	X	X	X
Create Value (successful)	X	NS	X	X	X	X
Modify Value (successful)	X	NS	X	X	X	X
Delete Value (successful)	X	NS	X	X	X	X
Create Key (failed)	X	NS	X	NS	NS	NS
Create Value (failed)	X	NS	X	NS	NS	NS
Modify Value (failed)	X	NS	X	NS	NS	NS
Delete value (failed)	X	NS	X	X	X	X

OS Windows Active Directory User Monitoring

The following table displays the platforms that support OS Windows Active Directory User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-10 OS Windows Active Directory User Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
User Create (successful)	X	X	X	X
User Delete (successful)	X	X	X	X

Table 4-10 (Cont.) OS Windows Active Directory User Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
User Attribute Modify (successful)	X	X	X	X

OS Windows Active Directory Computer Monitoring

The following table displays the platforms that support OS Windows Active Directory Computer Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-11 OS Windows Active Directory Computer Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
Computer Create (successful)	X	X	X	X
Computer Delete (successful)	X	X	X	X
Computer Attribute Modify (successful)	X	X	X	X

OS Windows Active Directory Group Monitoring

The following table displays the platforms that support OS Windows Active Directory Group Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 4-12 OS Windows Active Directory Group Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
Group Create (successful)	X	X	X	X
Group Delete (successful)	X	X	X	X
Group Attribute Modify (successful)	X	X	X	X
Group Member Add (successful)	X	X	X	X
Group Member Delete (successful)	X	X	X	X

A

Understanding PXE Booting and Kickstart Technology

This appendix explains PXE booting and kickstart technology in the following section:

- [About PXE Booting and Kickstart Technology](#)
- [Subnet Provisioning Usecases](#)

About PXE Booting and Kickstart Technology

One of the key requirements of provisioning is the hardware server's ability to boot over the network instead of a diskette or CD-ROM. There are several ways computers can boot over a network, and Preboot Execution Environment (PXE) is one of them. PXE is an open industry standard supported by a number of hardware and software vendors. PXE is part of the "Wired for Management" (WfM) specification, which is part of a bigger PC98 specification defined by Intel and Microsoft in 1998. A detailed document on PXE specification can be found at <http://www.pix.net/software/pxeboot/archive/pxespec.pdf>.

PXE works with Network Interface Card (NIC) of the system by making it function like a boot device. The PXE-enabled NIC of the client sends out a broadcast request to DHCP server, which returns with the IP address of the client along with the address of the TFTP server, and the location of boot files on the TFTP server. The following steps describe how it works:

1. Target Machine (either bare metal or with boot sector removed) is booted.
2. The Network Interface Card (NIC) of the machine triggers a DHCP request.
3. DHCP server intercepts the request and responds with standard information (IP, subnet mask, gateway, DNS etc.). In addition, it provides information about the location of a TFTP server and boot image (pxelinux.0).
4. When the client receives this information, it contacts the TFTP server for obtaining the boot image.
5. TFTP server sends the boot image (pxelinux.0), and the client executes it.
6. By default, the boot image searches the pxelinux.cfg directory on TFTP server for boot configuration files on the TFTP server using the following approach:

First, it searches for the boot configuration file that is named according to the MAC address represented in lower case hexadecimal digits with dash separators. For example, for the MAC Address "88:99:AA:BB:CC:DD", it searches for the file 01-88-99-aa-bb-cc-dd.

Then, it searches for the configuration file using the IP address (of the machine that is being booted) in upper case hexadecimal digits. For example, for the IP Address "192.0.2.91", it searches for the file "C000025B".

If that file is not found, it removes one hexadecimal digit from the end and tries again. However, if the search is still not successful, it finally looks for a file named "default" (in lower case).

For example, if the boot file name is /tftpboot/pxelinux.0, the Ethernet MAC address is 88:99:AA:BB:CC:DD, and the IP address 192.0.2.91, the boot image looks for file names in the following order:

```
/tftpboot/pxelinux.cfg/01-88-99-aa-bb-cc-dd
/tftpboot/pxelinux.cfg/C000025B
/tftpboot/pxelinux.cfg/C000025
/tftpboot/pxelinux.cfg/C00002
/tftpboot/pxelinux.cfg/C0000
/tftpboot/pxelinux.cfg/C000
/tftpboot/pxelinux.cfg/C00
/tftpboot/pxelinux.cfg/C0
/tftpboot/pxelinux.cfg/C
```

7. The client downloads all the files it needs (kernel and root file system), and then loads them.
8. Target Machine reboots.

The Provisioning application uses Redhat's Kickstart method to automate the installation of Redhat Linux on target machines. Using kickstart, the system administrator can create a single file containing answers to all the questions that will usually be asked during a typical Red Hat Linux installation.

The host specific boot configuration file contains the location of the kickstart file. This kickstart file would have been created earlier by the stage directive of the OS image based on the input from user.

Subnet Provisioning Usecases

Following are examples of subnet provisioning usecases:

Subnet of size 256

IP Prefix: 192.168.1.0

Subnet Mask: 255.255.255.0

Covers IPs from 192.168.1.0 - 192.168.1.255

Subnet of size 16

IP Prefix: 192.168.1.0

Subnet Mask: 255.255.255.240

B

Troubleshooting Issues

This appendix provides solutions to common issues you might encounter when using provisioning and patching Deployment Procedures. In particular, this appendix covers the following:

- [Troubleshooting Linux Provisioning Issues](#)
- [Troubleshooting Linux Patching Issues](#)
- [Frequently Asked Questions on Linux Provisioning](#)

Troubleshooting Linux Provisioning Issues

I cannot see my stage, boot server in the UI to configure them with the provisioning application?

Either Management Agents have not been installed on the Stage or Boot Server machine, or it is not uploading data to the OMS. For more information see *Troubleshooting in Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide* for troubleshooting information and known issues.

Bare metal machine is not coming up since it cannot locate the Boot file.

Verify the DHCP settings `/etc/dhcpd.conf` and TFTP settings for the target machine. Check whether the services (DHCPD, Xinetd, Portmap) are running. Make the necessary setting changes if required or start the required services if they are down.

Bare Metal box is not getting booted over network, or DHCP server does not get a DHCPDISCOVER message for the MAC address of the bare metal machine.

Edit the DHCP configuration to include the IP address of the subnet where the bare metal machine is being booted up.

Agent Installation fails after operating system has been provisioned on the bare metal box, or No host name is assigned to the bare metal box after provisioning the operating system

This might happen if the `get-lease-hostnames` entry in the `dhcpd.conf` file is set to true. Edit the `dhcpd.conf` file to set `get-lease-hostnames` entry to false. Also, ensure that length of the host name is compatible with length of the operating system host name.

Bare metal machine hangs after initial boot up (TFTP error/kernel error).

This may happen if the TFTP service is not running. Enable the tftp service. Go to the `/etc/xinetd.d/tftp` file and change the `disable` flag to: `disable=no`. Also verify the DHCP settings.

Kernel panic occurs when the Bare Metal machine boots up.

Verify the DHCP settings and TFTP settings for the target machine and make the necessary changes as required. In a rare case, the **intird** and **vmlinuz** copied may be corrupted. Copying them from RPM repository again would fix the problem.

Bare metal machine hangs after loading the initial kernel image.

This may happen if the network is half duplex. Half duplex network is not directly supported, however a modification in the kickstart file will resolve this issue. Modify `ethtool -s eth0 duplex half` to `ethtool -s eth0 duplex full` in the kickstart file.

Bare Metal machine cannot locate the kickstart file (Redhat screen appears for manually entering the values such as 'language', 'keyboard' etc).

This happens if `STAGE_TOP_LEVEL_DIRECTORY` is not mountable or not accessible. Make sure the stage top level is network accessible to the target machine. Though very rare but this might also happen because of any problem in resolving the stage server **hostname**. Enter the IP address of the stage or the NAS server instead of **hostname** on which they are located, and try the provisioning operation again.

Bare Metal machine does not go ahead with the silent installation (Redhat screen appears for manually entering the network details).

Verify that DNS is configured for the stage server host name, and that DHCP is configured to deliver correct DNS information to the target machine. If the problem persists, specify the IP address of the stage or NAS server instead of **hostname**, and try the provisioning operation again.

After provisioning, the machine is not registered in Enterprise Manager.

This happens if Enterprise Manager Agent is not placed in the `STAGE_TOP_LEVEL_DIRECTORY` before provisioning operation. Place the Enterprise Manager agent in this directory, and try the operation again. It might also happen if the OMS registration password provided for securing the agents is incorrect. Go to the agent oracle home on the target machine, and run the `emctl secure agent` command supplying the correct OMS registration password. Check the time zone of the OMS and the provisioned operating system. Modify the time zone of the provisioned operating system to match with the OMS time zone.

With 64-bit OS provisioning, agent is not installed.

During OS provisioning, specify the full path of the agent RPM in the Advanced Operating System Properties page.

Provisioning operations cannot be initiated since either one or all of Stage Server, Boot Server, and RPM Repository have not been configured in the Infrastructure page.

Set up at least one stage server, boot server, and RPM repository to proceed with Linux Provisioning.

Submitting the deployment operations shows error: "An unexpected error has occurred. Please check the log files for details." Logs have the corresponding message: "ComponentType with internal name BMPTType not found"

Set up Software Library from the Software Library console.

The deployment procedure fails with directory permission error.

This error occurs because of insufficient user privileges on the stage server machine. `STAGE_TOP_LEVEL_DIRECTORY` should have write permission for the stage server user. In case of NAS, the NAS directory should be mounted on the staging server. If the error appears while writing to the boot directory, then the boot server user must have the write permission.

Bare metal box fails to boot with "reverse name lookup failed" error.

Verify that the DNS has the entry for the IP address and the host name.

Fetching properties from reference machine throws the error: " Credentials specified does not have root access"

Verify if the credentials specified for the reference machine has `sudo` access.

Following Package/Package Group are not available in the RPM Repository. Either update the Package List or select the correct RPM Repository in the Deployment page.

Verify that the RPM packages mentioned in the error message are present in the repository, and that they are spelled correctly. If not, either copy the packages to the repository or do not install them.

Troubleshooting Linux Patching Issues

My Staging Server Setup DP fails at "Channels Information Collection" step with the error message "Could not fetch the subscribed channels properly". How do I fix this?

This error is seen if there is any network communication error between `up2date` and ULN. Check if `up2date` is configured with correct proxy setting by following https://linux.oracle.com/uln_faq.html - 9. You can verify if the issue is resolved or not by using the command, `up2date -nox -show-channels`. If the command lists all the subscribed channels, the issue is resolved.

My "up2date -nox -show-channels" command does not list the subscribed channels properly. How do I fix this?

Go to `/etc/sysconfig/rhn/sources` files, uncomment `up2date default` and comment out all the local RPM Repositories configured.

How can I register to channels of other architectures and releases?

Refer to https://linux.oracle.com/uln_faq.html for this and more such related FAQs.

After visiting some other page, I come back to "Setup Groups" page; I do not see the links to the jobs submitted. How can I get it back?

Click **Show** in the details column.

Package Information Job fails with "ERROR: No Package repository was found" or "Unknown Host" error. How do I fix it?

Package Repository you have selected is not good. Check if metadata files are created by running `yum-arch` and `createrepo` commands. The connectivity of the RPM Repository from OMS might be a cause as well.

Even after the deployment procedure finished its execution successfully, the Compliance report still shows my Group as non-compliant, why?

Compliance Collection is a job that runs once in every 24 hour. You should wait for the next cycle of the job for the Compliance report to update itself. Alternately, you can go to the **Jobs** tab and edit the job to change its schedule.

Package Information Job fails with "ERROR: No Package repository was found" or "Unknown Host" error. How do I fix it?

The package repository you have selected is not good. Check if the metadata files are created by running `yum-arch` and `createrepo` commands. The connectivity of the RPM Repository from OMS might be a cause as well.

I see a UI error message saying "Package list is too long". How do I fix it?

Deselect some of the selected packages. The UI error message tells you from which package to unselect.

Frequently Asked Questions on Linux Provisioning

What is PXE (Pre-boot Execution Environment)?

The Pre-boot Execution Environment (PXE, aka Pre-Execution Environment) is an environment to bootstrap computers using a network interface card independently of available data storage devices (like hard disks) or installed operating systems. See: [Understanding PXE Booting and Kickstart Technology](#) for more information.

Can my boot server reside on a subnet other than the one on which the bare metal boxes will be added?

Yes. But it is a recommended best practice to have boot server in the same subnet on which the bare metal boxes will be added. If the network is subdivided into multiple virtual networks, and there is a separate DHCP/PXE boot server in each network, the Assignment must specify the boot server on the same network as the designated hardware server.

If one wants to use a boot server in a remote subnet then one of the following should be done:

-- Router should be configured to forward DHCP traffic to a DHCP server on a remote subnet. This traffic is broadcast traffic and routers do not normally forward broadcast traffic unless configured to do so. A network router can be a hardware-based router, such as those manufactured by the Cisco Corporation or software-based such as Microsoft's Routing and Remote Access Services (RRAS). In either case, you need to configure the router to relay DHCP traffic to designated DHCP servers.

-- If routers cannot be used for DHCP/BOOTP relay, set up a DHCP/BOOTP relay agent on one machine in each subnet. The DHCP/BOOTP relay agent relays DHCP and BOOTP message traffic between the DHCP-enabled clients on the local network and a remote DHCP server located on another physical network by using the IP address of the remote DHCP server.

Why is Agent rpm staged on the Stage server?

Agent rpm is used for installing the agent on the target machine after booting over the network using PXE. With operating system provisioning, agent bits are also pushed on the machine from the staging location specified in the Advanced Properties.

Can I use the Agent rpm for installing Agent on Stage and Boot Server?

This is true only if the operating system of the Stage or Boot Server machine is RedHat Linux 4.0, 3.1 or 3.0 or Oracle Linux 4.0 or later.

Can the yum repository be accessed by any protocol other than HTTP?

Though the rpm repository can be exposed via file:// or ftp:// as well, the recommended method is to expose it via http://. The latter is faster and more secure.

What is the significance of the Status of a directive? How can one change it?

Look at the following table to know the possible Status values and what they signify.

Table B-1 Status Values

Status	Description
Incomplete	This Status signifies that some step was not completed during the directive creation, for example uploading the actual script for the directive, or a user saved the directive while creating it and still some steps need to be performed to make complete the directive creation.
Ready	This signifies that the directive creation was successful and the directive is now ready to be used along with any component/image.
Active	A user can manually change the status of a Ready directive to Active to signify that it is ready for provisioning. Clicking Activate changes the Status to Active.

What is a Maturity Level of a directive? How can one change it?

See [Table B-2](#) to know the possible Status values and what they signify:

Table B-2 Maturity Levels

Maturity Level	Maturity Level Description
Untested	This signifies that the directive has not been tested and is the default maturity level that is assigned to the directive when it is created.
Beta	A directive can be manually promoted to Beta using the Promote button after testing the directive.
Production	A directive can be manually promoted to Production using the Promote button after a user is satisfied that the directive can be used for actual provisioning on production systems.

Can a same component be used in multiple deployments?

Yes. Components are reusable and a given component can be a part of multiple deployments at the same time.

Do I need to edit scheduled deployments associated with a component, if the component is edited?

Yes.

For creating the Linux OS component does the Reference Machine need to have a management agent running on it?

Yes. Reference Machine has to be one of the **managed targets** of the Enterprise Manager.

What is the significance of the Status of a component? How can one change it?

Status of a component is similar to that of a directive. Refer to [What is the significance of the Status of a directive? How can one change it?](#).

What is a Maturity Level of a component? How can one change it?

Maturity Level of a component is similar to that of a directive. Refer to [What is a Maturity Level of a directive? How can one change it?](#).

Index

A

adding

- host configuration, [4-45](#)
 - hosts targets, [4-39](#)
 - local host group, [4-36](#)
 - local user, [4-35](#)
- administering, hosts, [4-26](#)
administration tasks, [4-30](#)
ASM, storage, [4-20](#)

B

bare metal provisioning

- concepts
 - boot server, [2-4](#)
 - reference host, [2-4](#)
 - RPM repository, [2-5](#)
 - stage server, [2-4](#)
 - overview, [2-2](#)
 - setting up, [2-5](#)
 - supported releases, [2-5](#)
- boot server, [2-4](#)
overview, [2-4](#)
setting up, [2-9](#)

C

commands

- hosts, [4-37](#)
- configuring hosts, [4-27](#)
consumption summary, storage, [4-20](#)
CPU statistics
 - hosts monitoring, [4-17](#)creating Disk Layout component, [2-15](#)
credentials
 - host, [4-15](#)
 - setting up monitoring, [4-16](#)customizing
 - hosts environment, [4-24](#)

D

databases

- storage, [4-20](#)

default system run level

- in hosts, [4-31](#)

Dell PowerEdge Linux hosts monitoring, [4-44](#)

dhcp server

- setting up, [2-9](#)

discovering hosts, [1-1](#)

- automatically, [1-1](#)
- manually, [1-1](#)

disks

- statistics, hosts monitoring, [4-18](#)
- storage, [4-20](#)

E

editing

- host configuration, [4-45](#)
 - local host group, [4-36](#)
 - local user, [4-35](#)
- EM_LINUX_PATCHING_ADMIN role, [3-6](#)
Execute Host Command
 - group, [4-41](#)
 - multiple hosts, [4-40](#)
 - single host, [4-42](#)execution history
 - host commands, [4-43](#)execution results
 - host command, [4-43](#)

F

file systems

- storage, [4-21](#)

G

- GPG keys, [3-4](#), [3-6](#)
- GPG signatures, [3-6](#)
- group administration, [4-35](#), [4-36](#)
- groups
 - in hosts, [4-25](#)

H

- history
 - statistics, [4-13](#)
 - storage, [4-23](#)
- Host command
 - executing using sudo or PowerBroker, [4-37](#)
 - running, [4-39](#)
- hosts
 - adding host targets, [4-39](#)
 - administering, [4-26](#)
 - administration, target setup, [4-16](#)
 - commands, [4-37](#)
 - execution history, [4-43](#)
 - execution results, [4-43](#)
 - configuration
 - adding and editing, [4-45](#)
 - configuring, [4-27](#)
 - customizing environment, [4-24](#)
 - default system run level, [4-31](#)
 - Dell PowerEdge Linux monitoring, [4-44](#)
 - diagnosing problems on, [4-12](#)
 - group administration, [4-35](#)
 - groups, [4-25](#)
 - installing YAST on, [4-14](#)
 - log file alerts, [4-18](#)
 - lookup table
 - host administration, [4-33](#)
 - metric collection errors, [4-18](#)
 - monitoring, [4-17](#)
 - monitoring setting up, [4-15](#)
 - NFS clients, [4-33](#)
 - overview, [4-12](#)
 - preferred credentials, [4-15](#)
 - running Host command, [4-39](#)
 - services, [4-31](#)
 - setting up credentials, [4-15](#)
 - statistics, [4-12](#)
 - CPU, [4-17](#)
 - disk, [4-18](#)
 - memory, [4-17](#)
 - program resource utilization, [4-18](#)
 - storage, [4-18](#)
 - tools, [4-37](#)
 - PowerBroker, [4-37](#)
 - Remote File Editor, [4-38](#)
 - sudo command, [4-37](#)
 - user administration, [4-35](#)
 - viewing targets on, [4-13](#)

K

- Ksplice for Oracle Linux, [4-24](#), [4-55](#)
- Ksplice Metrics, [4-55](#)

L

- layers, storage, [4-23](#)
- Linux hosts, installing YAST, [4-14](#)
- linux patching
 - package compliance, [3-7](#)
 - prerequisites, [3-2](#)
 - registering with ULN, [3-5](#)
 - setting up group, [3-6](#)
 - setting up infrastructure, [3-2](#)
 - setting up linux patching groups for compliance reporting, [3-6](#)
 - setting up RPM repository, [3-3](#), [3-4](#)
- linux patching groups, [3-1](#)
 - jobs, [3-6](#)
- local file systems, storage, [4-21](#)
- log file alerts, hosts monitoring, [4-18](#)

M

- memory statistics
 - hosts monitoring, [4-17](#)
- metric collection errors
 - hosts monitoring, [4-18](#)
- monitoring
 - credentials, setting up, [4-16](#)
 - hosts, [4-17](#)
 - NFS mounts, [4-21](#)

N

- named credentials, host, [4-15](#)
- network cards
 - configuring, [4-32](#)
 - in hosts, [4-32](#)
- network file systems See NFS, [4-21](#)
- NFS (network file systems)
 - clients
 - adding and editing, [4-33](#)
 - host administration, [4-33](#)
 - monitoring mounts, [4-21](#)
 - storage, [4-21](#)

O

- Oracle Cloud Compute configuration, [4-2](#)
- OS script, load, [4-43](#)

P

- patching
 - linux patching, [3-1](#)
 - patching linux hosts, [3-1](#)

patching linux hosts
 concepts, [3-1](#)
 deployment procedures, [3-2](#)
 meeting prerequisites, [3-2](#)
 overview, [3-1](#)
 package compliance, [3-7](#)
 setting up infrastructure, [3-2](#)
 setting up linux patching groups for
 compliance reporting, [3-6](#)
 setting up RPM repository, [3-3](#), [3-4](#)
 supported linux releases, [3-2](#)

PowerBroker tool, [4-37](#)
 executing host command, [4-37](#)

preferred credential, hosts, [4-15](#)

privilege delegation setting, [4-15](#)

problems
 diagnosing on hosts, [4-12](#)

program resource utilization statistics, hosts
 monitoring, [4-18](#)

provision Linux, [2-1](#)
 getting started, [2-1](#)

provisioning
 provisioning linux operating system, [2-1](#)
 storage, [4-19](#)

provisioning operating systems, [2-18](#)

R

reference host, [2-4](#)

refresh, storage, [4-24](#)

Remote File Editor tool, [4-38](#)

routing configuration, network cards, [4-32](#)

RPM packages, [3-2](#), [3-4](#)

RPM repository, [2-5](#), [3-2](#), [3-3](#)
 overview, [2-5](#)
 setting up, [2-10](#), [3-3](#)

S

services
 in hosts, [4-31](#)

setting up
 environment to monitor hosts, [4-13](#)
 host credentials, [4-15](#)
 host monitoring, [4-15](#)
 monitoring credentials, [4-16](#)
 target for host administration, [4-16](#)

stage server, [2-4](#)
 overview, [2-4](#)

stage server (*continued*)

statistics
 hosts, [4-12](#)
 storage, [4-13](#)

storage
 file systems, [4-21](#)
 history, [4-23](#)
 hosts monitoring, [4-18](#)
 layers, [4-23](#)
 network file systems, [4-21](#)
 refresh, [4-24](#)
 statistics, [4-13](#)
 utilization, [4-19](#)
 vendor distribution, [4-23](#)
 volumes, [4-22](#)

sudo command, [4-37](#)
 executing host command, [4-37](#)

T

target setup
 host administration, [4-16](#)

targets
 host, [4-13](#)

tools
 hosts, [4-37](#)

U

ULN channels, [3-1](#), [3-4](#)

ULN configuration channel, [3-2](#)

ULN custom channel, [3-2](#)

Unbreakable Linux Network (ULN), [3-1](#)

user administration, [4-35](#), [4-36](#)

V

vendor distribution, storage, [4-23](#)

volumes, storage, [4-22](#)

W

web-based enterprise management (WBEM)
 fetchlet metrics, [4-44](#)

Y

YAST, installing on Linux hosts, [4-14](#)

yum patching tool, [3-3](#), [3-6](#)