Oracle® Enterprise Manager

Administrator's Guide for Software and Server Provisioning and Patching

11*g* Release 1 (11.1.0.1.0)

E16599-06

October 2012

ORACLE®

Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching, 11*g* Release 1 (11.1.0.1.0)

E16599-06

Copyright © 1996, 2012, Oracle and/or its affiliates. All rights reserved.

Primary Author: Aravind Jayaraaman

Contributing Author: Aparna Kamath, Pushpa Raghavachar

Contributor: Sudip Datta, Rajat Nigam, Akanksha Sheoran, Hari Prasanna Srinivasan, Madhup Gulati, Bharat Paliwal, Dan Hynes, Kamaldeep Khanuja, Pradeepkumar Itharaju, Premjith Rayaroth, Eunhei Jang, Ryan Lemos, Ranjit Madampath, Macks Ningombam, Shamik Ganguly, Suparna Gurunaga, Amandeep Barboza, Rahul Gupta, Srivatsan Vijayaraghavan, Sudhir Ravindramohan, Vijay Sriram, Vivek Dosi, Kapil Makhija, Raman Venkateswaran, Siddharth Paradkar, Sachin Goel, Saurabh Garg.

Contents

3 Infrastructure Requirements

Part II Database Provisioning

4 Simulating Database Provisioning Deployment Procedures to Check Prerequisites

5 Provisioning Oracle Database

6 Provisioning Oracle Grid Infrastructure for Standalone Servers

7 Provisioning Oracle Grid Infrastructure and Oracle RAC for Clustered Environments

8   Extending Oracle Real Application Clusters

9   Deleting or Scaling Down Oracle Real Application Clusters

10   Provisioning Oracle Database Replay Client

Part III   Provisioning Linux Operating System

11   Provisioning Linux Operating System

Part IV   Provisioning Virtualization Systems

12   Provisioning Virtualization Systems

Part V   Middleware Provisioning

13   Provisioning Oracle Fusion Middleware

14   Provisioning SOA Artifacts and Composites

15   Provisioning Coherence Nodes and Clusters

16   Provisioning Oracle Service Bus Resources

17 Provisioning Oracle BPEL Processes

18 Provisioning Oracle Application Server

Part VI Patching

19 Getting Started with Patching

20 Patching Using My Oracle Support

21 Simulating Patching Deployment Procedures to Check Prerequisites

22 Patching Oracle Database

23 Patching Oracle RAC

24 Patching Oracle Clusterware

25 Patching Oracle Automatic Storage Management (ASM)

26   Patching Primary and Standby Databases Configured with Oracle Data Guard

27   Patching Oracle Application Server

28   Patching Linux Hosts

29   Patching Solaris Hosts

30   Patching Microsoft Windows Hosts

Part VII   Customization

31   Customizing Deployment Procedures

Part VIII   Enterprise Manager Command Line Interface (EMCLI)

32   Using EMCLI to Run Deployment Procedures

Part IX   Additional Information

A   Troubleshooting Issues

B   Using emctl partool Utility

C   Using Oracle Software Library

D   Creating Templates for Oracle Database

E   Out-Of-Box RuntimeData Templates

F   Sample Property Files for the Out-of-Box RuntimeData Templates

G   PXE Booting and Kickstart Technology

H   About Patching Automation Reports

I   Converting Standalone Oracle Management Agent to Cluster Agent

J   Manually Downloading Prerequisites from My Oracle Support

K   Staging Gold Image to Software Library or Staging Location

L   Downloading Installation Binaries and Uploading to Software Library or Staging Location

M   Reviewing Log Files

Index

Preface

The Administrator's Guide for Software and Server Provisioning and Patching introduces you to the lifecycle management solutions offered by Enterprise Manager Grid Control, and helps you understand what Deployment Procedures are and how they can be used for meeting your provisioning and patching requirements.

The guide has been created with great care and attention using multiple sources of information, including valuable notes from My Oracle Support, white papers and marketing collaterals from product management, feedback from development teams and also from customers that came in the form of bugs, services requests, and so on.

Understandably, this guide describes one of the most critical offerings of Enterprise Manager Grid Control and since some of the concepts may be difficult to describe using textual descriptions, we have considered some of the best practices in the industry and added as many examples and graphics as possible to help you understand the concepts, processes, and features better.

Audience

This guide is primarily meant for administrators who want to use the Deployment Procedures offered by Enterprise Manager Grid Control to meet their lifecycle management challenges. As an administrator, you can be either a Designer, who performs the role of a system administrator and does critical data center operations, or an Operator, who runs the out-of-box, automated procedures to maintain the enterprise configuration.

Structure of the Document

This guide covers conceptual information, facts you must know before you start using the Deployment Procedures, and a comprehensive set of instructions that help you understand how you can run the Deployment Procedures for different areas.

This section describes how the entire guide is structured. In particular, this section describes the following:

- Different Parts of the Guide
- Structure of the Chapters

Different Parts of the Guide

This guide has the following parts:

Part	Description
Part I, "Overview and Setup Details"	Introduces you to the lifecycle management solutions offered by Enterprise Manager Grid Control, describes the concepts and terms related to these solutions, and the infrastructure requirements for running the Deployment Procedures.
Part II, "Database Provisioning"	Describes the Deployment Procedures offered for provisioning Oracle Databases, Oracle RAC, and Oracle Database Replay Clients. Also describes the Deployment Procedures offered for extending, scaling down, or deleting an existing Oracle RAC instance.
	Covers information about the supported releases, core components deployed, prerequisites to be met, and instructions to be followed for running the Deployment Procedures.
Part III, "Provisioning Linux Operating System"	Describes how you can provision Linux operating system on bare metal servers and virtual servers.
Part IV, "Provisioning Virtualization Systems"	Describes how you can provision Oracle Virtual Server (OVS) and Guest Virtual Machines, clone Guest Virtual Machines, and live-migrate Guest Virtual Machines.
Part V, "Middleware Provisioning"	Describes the Deployment Procedures offered for provisioning Oracle Application Server, Oracle Service Bus resources, and Oracle BPEL Processes.
	Covers information about the supported releases, core components deployed, prerequisites to be met, and instructions to be followed for running the Deployment Procedures.
Part VI, "Patching"	Describes the Deployment Procedure offered for patching Oracle Database, ASM, Oracle Application Server, Oracle Clusterware, Oracle RAC, and operating system such as Linux, Microsoft Windows, and Sun Solaris.
	Covers information about the supported releases, core components deployed, prerequisites to be met, and instructions to be followed for running the Deployment Procedures.
Part VII, "Customization"	Describes how you can customize Deployment Procedures. If you are already familiar with Deployment Procedures and have already used the out-of-box procedures for provisioning and patching, then this part will help you understand how you can create custom copies of those default procedures with changes specific to your environment. For example, you can enable e-mail notifications or even run the procedure as another user.
Part VIII, "Enterprise Manager Command Line Interface (EMCLI)"	Describes the EMCLI commands you can use to run Deployment Procedures.
Part IX, "Additional Information"	Provides additional information that are relevant and related to the details covered in the other chapters of this guide.

Structure of the Chapters

The chapters of part II, III, IV, and V cover multiple use cases and provide a comprehensive set of instructions. These chapters have the following structure:

- *Getting Started* section that helps you get started with that chapter. The section provides a high-level overview of the steps involved in performing a particular

provisioning or patching task. Consider this section to be a documentation map to understand the sequence of actions you must take to successfully run a Deployment Procedure. Every step has a set of reference links that you can click to reach the relevant sections within the chapter and view more information.

- *Deployment Procedure* section that describes the Deployment Procedure offered by Enterprise Manager Grid Control for a particular provisioning or patching requirement. Read this section to understand the Deployment Procedure, the use cases it covers, and so on.

- *Supported Releases* section that lists the versions and releases of the target that are supported by the Deployment Procedure for provisioning or patching. Read this section to understand the coverage of the Deployment Procedure.

- *<Use Case>* sections that describe the use cases covered in the chapter. Every use case has the following structure:

  - *Deployment Phases* section that describes the deployment phases involved in a particular provisioning or patching task. Read this section to understand the flow of interview screens involved in the Deployment Procedure.

  - *Prerequisites* section that lists the prerequisites you must meet before running the Deployment Procedure. Read this section and meet all the prerequisites before you move on to the instructions to run the Deployment Procedure.

  - *Provisioning or Patching Procedure* section that outlines the instructions for running the Deployment Procedure. Read this section and follow the steps to successfully run the Deployment Procedure for your provisioning or patching requirement.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information, see the following documents in the Enterprise Manager documentation library available on Oracle Technology Network (OTN):

- *Oracle Enterprise Manager Grid Control Release Notes*

- *Oracle Enterprise Manager Grid Control Basic Installation Guide*

- *Oracle Enterprise Manager Grid Control Advanced Installation Guide*

You can access the Enterprise Manager documentation library using the following URL:

http://www.oracle.com/technology/documentation/oem.html

Conventions

The following conventions are used in this document:

Convention	Meaning
boldface	Indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.
Other Graphics	Graphics have been used extensively in addition to the textual descriptions to ensure that certain concepts and processes are illustrated better.

Part I

Overview and Setup Details

This part provides an overview of the lifecycle management challenges you may face in your environment and introduces you to the solutions offered by Enterprise Manager Grid Control that can help you manage critical data center operations more easily and efficiently.

In particular, this part covers the following chapters:

- Chapter 1, "Overview of Lifecycle Management Solutions"
- Chapter 2, "Things to Know"
- Chapter 3, "Infrastructure Requirements"

1

Overview of Lifecycle Management Solutions

This chapter provides an overview of the lifecycle management solutions offered by Enterprise Manager Grid Control. In particular, this chapter covers the following:

- Lifecycle Management Challenges
- Enterprise Manager Grid Control's Lifecycle Management Solutions
- Benefits of Using Provisioning and Patching Features
- Success Stories

1.1 Lifecycle Management Challenges

In today's world, proliferation of low cost servers in our data centers and numerous software deployments on those servers have brought in a fresh set of lifecycle management challenges.

One of the well-acknowledged problems include the difficulty in managing consistency and compatibility across operating systems and software deployments, server drifts and security vulnerabilities that lead to lack of compliance. Besides that, other problems include the difficulty in deploying new software and maintaining them, difficulty in provisioning new servers with variety of configurations and maintaining them, and difficulty in adapting to changes in workload of the environment.

These lifecycle management challenges eventually result in system administrators and database administrators spending significant amount of their time in software and server provisioning and maintenance operations.

1.2 Enterprise Manager Grid Control's Lifecycle Management Solutions

Enterprise Manager Grid Control offers lifecycle management solutions in the form of Deployment Procedures that help you manage the entire lifecycle of software, applications, and servers. Using Deployment Procedures, you can automate not only the deployment of software, applications, and servers but also the maintenance of these deployments. This makes critical data center operations easy, efficient, and scalable resulting in lower operational risk and cost of ownership.

The provisioning and patching solutions are offered across the stack for:

- Operating Systems, with Bare metal provisioning on Linux and operating system patching

- Databases, with Real Application Clusters (RAC) and Oracle Grid Infrastructure provisioning, extension, and deletion and flexible patching for Oracle database and Oracle Real Application Clusters

- Middleware, with application server J2EE, BPEL, and SOA provisioning and patching

The following illustrates how Enterprise Manager Grid Control manages the software lifecycle using its Deployment Procedures:

Figure 1–1   Software Lifecycle Management Overview
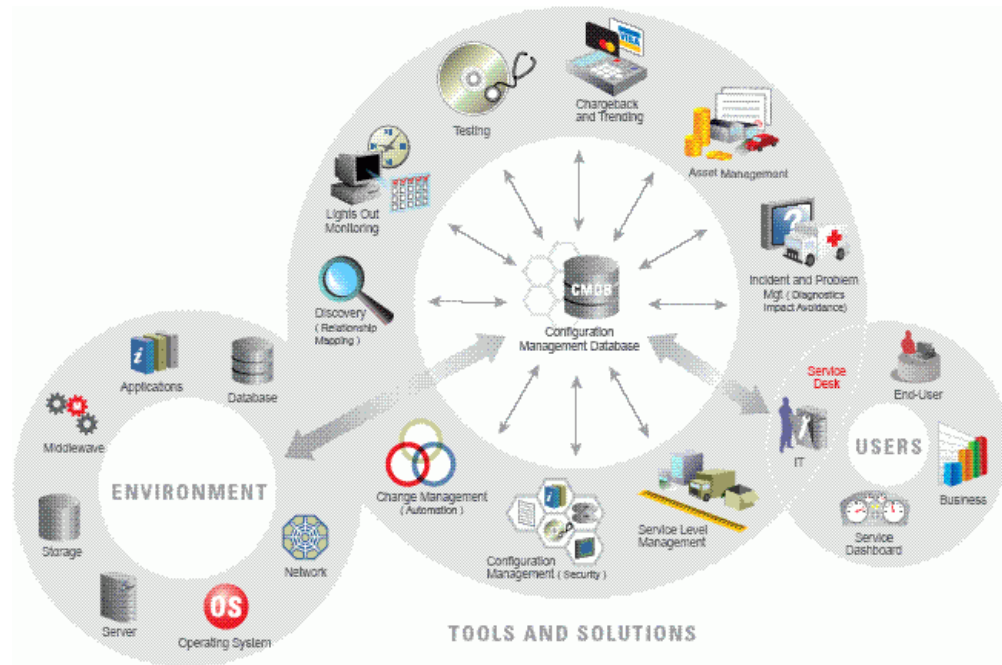
You can see in the illustration that Enterprise Manager Grid Control orchestrates the initial reference sandbox deployment and then the mass unattended deployment of gold images created from these reference deployments.

Besides deployment of software, Enterprise Manager Grid Control also automates the ongoing patch lifecycle management of these deployments - proactively informs you about the critical patches and vulnerabilities in the deployments, automatically downloads the required patches, applies them to the deployments, and also verifies whether the patches were applied successfully.

1.3 Benefits of Using Provisioning and Patching Features

The following are the advantages of using the provisioning and patching features offered by Enterprise Manager Grid Control:

- The provisioning and patching features provide a repeatable, reliable, and automated solution for performing mass, unattended, and schedulable deployment of

  - Software and servers based on Gold Images created using reference deployment or installation media

  - Software and operating system updates

- Complex and multi-tier software like Oracle Real Application Clusters (RAC) and Application Server Clusters

- Orchestrates not only provisioning of software but completely automates configuration of software and ensures zero-time for patching of mission critical systems by orchestrating rolling patching for complex multi-tier installation like Real Application Clusters (RAC) databases and Application Server clusters.

- Allows new resources to be provisioned at short notice based on compliant and tested gold images.

- Enterprise Manager Grid Control's Critical Patch Facility proactively and regularly queries Oracle Metalink for critical patches that have been released and notifies the administrators of only those patches applicable to them. It can also invoke the Patch application in context, and remedy the vulnerable installations. The Critical Patch Facility also supports an offline mode to serve the case of data centers that are not connected to the Internet.

- Completely automates the patching operations across the stack. For example, for database patching, it takes care of shutting down and starting up database instances as required by the patch.

- Allows multiple operations to be accommodated in a single change window.

- Enterprise Manager Grid Control also provides command-line interface support to all out-of-box provisioning and patching deployment procedures. These features can hence be invoked by custom scripts.

- Software provisioning and patching features support SUDO and PAM authentication.

- Single interface for multiple players. For example, component designers responsible for creating gold images based on corporate standards and the operators all use the same Enterprise Manager Grid Control console.

- Automation of repeatable installation and patching operations across the stack leads to substantial cost savings in terms of costs and man-hours.

1.4 Success Stories

NetApp Uses Automated Provisioning to Reduce Manual Effort Time from a Few Weeks to a Few Clicks

"With Oracle Enterprise Manager, we are now able to meet corporate standards 100% of the time. And, our pre-patched gold images, enabled by Oracle Enterprise Manager, save time and resources."

– Alok Arora, Director, Enterprise Integration and Architecture, NetApp, Inc.

NetApp, formerly Network Appliance, creates innovative storage and data management solutions that help its customers accelerate business breakthroughs and achieve outstanding cost efficiency. NetApp's dedication to principles of simplicity, innovation, and customer success has made it one of the fastest-growing storage and data management providers today.

The following were their lifecycle management challenges:

- Reduce two-week time frame needed for IT group to create a complex multitier service-oriented architecture (SOA) deployment with components spanning various hosts.

- Provision new deployments for a growing number of internal customers quickly, efficiently, and in a manner that complies with corporate standards.

- Improve day-to-day work life for IT staff by reducing amount of time needed to roll out patches developed during each patching cycle, a process that took between 8 and 10 weeks.

The following were the lifecycle management solutions offered by Oracle:

- Utilized the Provisioning and Patch Automation Pack for Oracle Middleware, and helped NetApp to expedite provisioning of new SOA deployments with gold images that enable administrators to provision deployments via a single management console.

- Created a standard process to provision and patch Fusion Middleware in a completely automated, and reliable manner — reducing manual effort from a few weeks to a few clicks.

- Created test benches based on production installations — allowing NetApp to apply patches on these test systems and then deploy them to multiple production systems.

- Enabled NetApp to deploy standard Fusion Middleware component images based on features licensed, and to have a standard process that leads to best-practice deployments with no configuration drifts from corporate standards.

- Enabled the IT staff to direct its attention to more strategic initiatives.

- Provided IT staff with processes and procedures that enable them to meet unpredictable demands from customers.

More Success Stories

To view more success stories and understand how Enterprise Manager Grid Control's lifecycle management solutions help organizations control IT management costs while improving the quality of their services, access the following URL:

http://www.oracle.com/customers/products/ent_manager.html

2

Things to Know

Enterprise Manager Grid Control offers complete lifecycle management solutions for Server and Software. This chapter describes the overall coverage of the provisioning and patching features, the coverage of Deployment Procedures used to orchestrate these tasks, and other aspects you must know before you begin using them.

Oracle strongly recommends you to read this chapter so that you understand the fundamentals and terms used in this context.

This chapter covers the following:

- Understanding Provisioning and Patching Solutions
- Understanding Deployment Procedures
- Understanding Provisioning Deployment Procedures
- Understanding Patching Deployment Procedures
- Understanding Oracle Software Library
- Understanding Linux Patching
- Understanding SOA Artifacts Provisioning
- Understanding Virtualization Management Provisioning
- Understanding Bare Metal Provisioning
- Understanding My Oracle Support
- Understanding Other Related Concepts

2.1 Understanding Provisioning and Patching Solutions

This section describes the following:

- Provisioning
- Patching
- Patching Versus Upgrading
- Patches

2.1.1 Provisioning

Provisioning is the first phase of the software lifecycle that deals with installation and configuration of software, applications, or servers across different platforms, environments, and locations.

Enterprise Manager Grid Control offers provisioning Deployment Procedures that help you seamlessly deploy software across your enterprise configuration.

2.1.2 Patching

Patching is one of the phases of the software lifecycle that helps you maintain the software over a period of time and always keep it updated with the latest features offered by the software vendor. It is a process that deals with migrating from one version of the software product to another, higher, much-improved version.

Enterprise Manager Grid Control offers patching Deployment Procedures that help you seamlessly patch Oracle software across your enterprise configuration.

Oracle releases patches periodically to help you migrate to a newer version of the software product and reap the benefits of an improved version. These patches contain critical enhancements and bug fixes that make the product better and offer more value to your organization.

For example, for the Oracle Database 10g Release 1 (10.2.0.1) base release, Oracle has released patches that help you migrate to either Oracle Database 10g Release 2 (10.2.0.2), Oracle Database 10g Release 3 (10.2.0.3), or Oracle Database 10g Release 4 (10.2.0.4); the last one being the most recent patch.

2.1.3 Patching Versus Upgrading

As described in the previous section, *Patching* refers to the process of migrating from one version to another, newer version of the software product within a particular release.

For example, if Oracle Database 10g is the release, then Patching refers to migrating from Oracle Database 10g Release 1 (10.2.0.1) to Oracle Database 10g Release 2 (10.2.0.2). Here, Release 2 (10.2.0.1), Release 2 (10.2.0.2), and so on are different versions of the 10g release.

Upgrading refers to the process of upgrading from one release to another, newer release of the software product.

For example, if Oracle Database 10g is one release and Oracle Database 11g is another release, then Upgrading refers to upgrading Oracle Database 10g to Oracle Database 11g. Here, 10g, 11g, and so on are different releases of the Oracle Database product.

2.1.4 Patches

Patches are periodic software updates released by Oracle. They contain bugs fixes only, and they must be applied on the software product to maintain a bug-free and improved version of the software.

Patches can be one-off patches, interim patches, or critical patch updates (CPU). Some of these may even be requirement-specific and released only to a few customers.

Patches can be stored as generic components in Oracle Software Library. For information about Oracle Software Library, see Section 2.5.1, "Oracle Software Library".

2.1.5 Patch Sets

Patch Sets are new releases of the software product released by Oracle. They contain a collection of patches, important enhancements, and bug fixes. They are much larger in

size, and are released for all customers. They must be applied on the base or existing release of the software product to maintain an up-to-date version of the software.

Patch Sets can be stored as generic components in Oracle Software Library. For information about Oracle Software Library, see Section 2.5.1, "Oracle Software Library".

2.2 Understanding Deployment Procedures

This section describes the following:

- Deployment Procedures
- Understanding Phases and Steps
- Activities Involved in Deployment Procedures
- Users of Deployment Procedures
- Viewing Details of Deployment Procedures

2.2.1 Deployment Procedures

A *Deployment Procedure* is a procedure that contains a hierarchal sequence of provisioning or patching steps, where each step may contain a sequence of other steps. In other words, the workflow of all tasks that need to be performed for a particular lifecycle management activity is encapsulated in a Deployment Procedure.

For example, imagine a Deployment Procedure to be the "Coffee" button on the coffee vending machine. So, every time you need a cup of coffee, you press the "Coffee" button and it automatically adds the required ingredients and serves coffee through its nozzle.

You select a Deployment Procedure in the Enterprise Manager Grid Control console to perform a particular provisioning or patching task. The Deployment Procedure contains a set of predefined steps that run one by one automatically to complete the entire task.

Enterprise Manager Grid Control comes with a set of default Deployment Procedures that help you accomplish common provisioning and patching-related tasks. Each Deployment Procedure is unique, and is designed to perform a particular operation according to the source being provisioned or target being patched. For example, the Deployment Procedure to patch a single instance database differs from the one to patch an Oracle RAC environment or an application server.

Understandably, your provisioning and patching-related requirements may vary from one environment to another, or from a test installation to a production installation. You can use the default Deployment Procedures and customize them according to your needs.

2.2.2 Understanding Phases and Steps

Deployment Procedures comprise various phases and steps that run serially or in parallel to perform a particular provisioning or patching operation.

A phase contains steps or more phases. The different types of phases are:

- *Rolling Phase*

  Rolling phase is a phase where steps are run serially across targets.

- *Parallel Phase*

  Parallel phase is a phase where steps are run in parallel across targets.

A step is an abstraction of a unit of work. For example, starting the database is a step. It is either part of a phase or independent. The different types of steps are:

- *Manual Step*

  Manual step is a task that requires user interaction and cannot be automated. Typically, Enterprise Manager Grid Control displays the instructions that need to be performed by you. After you perform those operations, you proceed to the next step. For example, logging in to a system and updating the kernel parameter, rebooting a system, providing special privileges to the user, and so on.

- *Computational Step*

  Computational step is a task whose operations are performed within Enterprise Manager Grid Control and they do not require any user intervention. This step gathers additional information for running a procedure. This step cannot be manually inserted by you. For example, running an SQL query against a schema to gather more data for other steps to run, retrieving target properties from the repository and updating the runtime information, and so on.

- *Action Step*

  Action step is a task that performs some operations run on one or more targets. They must be enclosed within a phase. The Deployment Procedure maps the

Action Step and target pair to a job in the Enterprise Manager Job System. The Deployment Procedure can schedule, submit, and run a job per Action Step per target. For example, running a script, applying a patch, upgrading an Oracle home, and so on.

Also note that an Action Step is said to have completed successfully only when all its associated jobs have completed successfully. If a job fails, then the Action Step also fails. You may then manually restart the job in question or ignore and instruct the Deployment Procedure to proceed.

The different types of Action Steps include:

– *Job*

    Job Step is a special type of Action Step that executes a predefined job type on a target. This is used if you want to execute a job type as a part of a Deployment Procedure. You need to pass job parameters for a step.

– *Directive*

    Directive step is a special type of Action Step to deploy a directive alone. This is useful when users want to store their custom scripts in Oracle Software Library (Software Library) and reuse them in a Deployment Procedure.

– *Component* (Generic or Registered)

    Generic component step is a special type of Action Step to deploy a Software Library component and the associated Directive. Deployment Procedure Manager executes the directive with respect to the component. Components used for Generic Component Step generally have one directive associated with them. This association is done by selecting both the component and directive while creating the step. All directives that you associate with the component while uploading to the Software Library are ignored while running the step.

    – *Host Command*

       Host Command Step is a special type of Action Step that encapsulates simple host commands. This step allows the user to enter a command line or a script (multiple commands) to be executed on the target host.

2.2.3 Activities Involved in Deployment Procedures

The activities involved in Deployment Procedures are design-time activities and run-time activities.

- *Design-time* activities include creating components, directives, and images, and storing them in Oracle Software Library; and sometimes even customizing the default Deployment Procedures according to the needs of the organization.

- *Run-time* activities include using these Deployment Procedures to actually provision or patch software, applications, or servers.

2.2.4 Users of Deployment Procedures

In a typical data center, the main users of Deployment Procedures are Designers (Lead Administrators) and Operators.

Designers (Lead Administrators) are users who perform design-time activities described in Section 2.2.3, "Activities Involved in Deployment Procedures". Essentially, they are responsible for setting up the infrastructure for running Deployment Procedures.

Operators are users who perform run-time activities described in Section 2.2.3, "Activities Involved in Deployment Procedures". Essentially, they are responsible for running the Deployment Procedures to provision or patch a target.

2.2.5 Viewing Details of Deployment Procedures

To view the default configuration settings of a Deployment Procedure and the steps involved in it, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures tab, from the table, select the Deployment Procedure for which you want to view details, and click **View**.

   Enterprise Manager Grid Control displays the View Procedure page that shows the default configuration settings and steps involved in the selected Deployment Procedure.

2.3 Understanding Provisioning Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for provisioning:

Target Type	Provisioning Deployment Procedure	Chapter No.
Oracle Database (Standalone)	■ Oracle Database Provisioning	■ Chapter 5
	■ Oracle Database Replay Client Provisioning	■ Chapter 10
	■ Oracle Grid Infrastructure Provisioning for Standalone Servers	■ Chapter 6
Oracle RAC Database	■ Oracle Grid Infrastructure / RAC Provisioning	■ Chapter 7
	■ One Click Extend Cluster Database	■ Chapter 8
	■ Delete/Scale down Oracle Real Application Clusters	■ Chapter 9
Middleware	■ Application Server Deployment (myJ2EE) 10.1.2.0.2	■ Chapter 18
	■ Application Server Deployment 10.1.3	■ Chapter 18
	■ Application Server Deployment 10.1.3.xSOA	■ Chapter 18
	■ Oracle BPEL Process Manager Provisioning	■ Chapter 17
	■ Oracle Service Bus Provisioning	■ Chapter 16
	■ SOA Artifacts Provisioning	■ Chapter 14
	■ SOA Composites Provisioning	■ Chapter 14
Virtualization	■ Live Migration	Chapter 12
	■ Virtual Machine Cloning	
	■ Edit Virtual Machine	
	■ Create Shared Disk	
	■ Virtual Machine Provisioning Using PXE	
	■ Virtual Machine Provisioning Using ISO	
	■ Virtual Machine Provisioning Using Template	
	■ Import ISO Image	
	■ Import Template	
	■ Import Guest Virtual Machine (V2V)	
	■ Import Template From V2V	
	■ Import Template From P2V	
	■ Discover Shared Disks	

> **Note:** The Cloning Wizard is no longer supported in Enterprise Manager 11g Grid Control Release 1 (11.1.0.1.0).

2.4 Understanding Patching Deployment Procedures

This section describes the following:

- Patching Deployment Procedures
- Understanding Patching Modes
- Understanding Patching Methodologies

2.4.1 Patching Deployment Procedures

The following are the patching Deployment Procedures offered by Enterprise Manager Grid Control.

Target Type	Patching Deployment Procedure	Chapter No.
Operating System Patching	■ Patch Windows hosts ■ Patch Linux hosts ■ Patch Solaris hosts	■ Chapter 30 ■ Chapter 28 ■ Chapter 29
Oracle Database (Standalone)	Patch Oracle Database	Chapter 22
Oracle ASM	Patch Standalone Oracle ASM	Chapter 25
Oracle RAC	■ Patch Oracle RAC Database - Rolling ■ Patch Oracle RAC Database - All Nodes	Chapter 23
Oracle Clusterware	■ Patch Oracle Clusterware - Rolling ■ Patch Oracle Clusterware - All Nodes	Chapter 24
Middleware	Patch Application Server	Chapter 27

> **Note:**
>
> - The *Patch Oracle Database* Deployment Procedure supports 9.2.0.8 or higher database patch sets. For patch sets that are lower than 9.2.0.8, you cannot use Deployment Procedures, and therefore, you will have to manually apply them.
>
> - By default, for the "Patch Oracle Database" Deployment Procedure, the Stop CSS Daemon and Start CSS Daemon steps are disabled. However, if you are applying a 10.1.0.5 database patch set, then before running the Deployment Procedure, enable the Stop CSS Daemon and Start CSS Daemon steps, and set the Run Privilege setting of these two steps to Normal.

2.4.2 Understanding Patching Modes

Using Deployment Procedures, you can apply a patch or patch set in either online mode or offline mode.

Online Patching Mode helps you if the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support.

Using Enterprise Manager Grid Control, you can schedule a job that will run periodically and connect to My Oracle Support, check for the latest patches and patch sets, and automatically download them. This relieves you of the manual effort involved in searching the latest patches and patch sets, and downloading them whenever they are available.

Offline Patching Mode helps you if the host that runs Enterprise Manager Grid Control does not have an Internet connection to connect to My Oracle Support.

In this case, you must manually download the patch and patch sets from My Oracle Support using another host that has an Internet connection. Using Enterprise Manager Grid Control, you can manually upload these patches and patch sets to the Patch Cache so that they can be applied later.

2.4.3 Understanding Patching Methodologies

For Oracle RAC Database and Oracle Clusterware, you can patch all instances of the cluster either in a rolling or parallel mode. The parallel manner is called *All Nodes* in Enterprise Manager Grid Control.

- *Rolling* mode refers to the patching methodology where the nodes of the cluster are patched individually, that is, one by one.

  For example, if you are patching a clusterware that has five nodes, then the first node is shut down, patched, and restarted, and then the process is rolled over to the next node until all the nodes are patched successfully.

  This methodology is best suited for the following use cases:

  - When you are applying one-off patches that support this methodology.

    > **Note:** Although most one-off patches support this methodology, there are some one-off patches that can be applied only using the *All Nodes* methodology. You can identify the supported methodology while downloading the one-off patches from My Oracle Support.

  - When you want to maintain high availability of your targets, so when one node is being patched, the other nodes are available for service.

    > **Note:** You cannot use this methodology to patch shared Oracle homes. If you want to patch a shared Oracle home, then use the *All Nodes* methodology.

- *All Nodes* (or parallel) mode refers to the patching methodology where all the nodes are patched at a time, collectively. In this methodology, all the nodes are shut down and the patch is applied on all of them at the same time.

  This methodology is best suited for the following use cases:

  - When you are applying patch sets, or one-off patches that support this methodology.

  - When you want to patch a shared Oracle home clusterware. In case of a shared Oracle home, all the nodes are shut down, but the patch is applied only on one of the nodes.

    For example, if you are patching a shared Oracle home clusterware that has five nodes, the patch and the scripts within it are run only once on the shared

Oracle home, and the other nodes are ignored or skipped. This saves time as the patching operation is performed only once, that is, only on the shared Oracle home, and not on all the nodes.

> **Note:** Patching of shared Oracle homes is supported only for Oracle Clusterware 10g Release 2 (10.2.0.1) or higher.

2.5 Understanding Oracle Software Library

This section describes the following:

- Oracle Software Library

- Components

- Directives

- Images, Network Templates, Hardware Templates, Disk Layouts

- Assignments

2.5.1 Oracle Software Library

Oracle Software Library (Software Library) is a repository that stores certified software images (for example, Oracle Database, operating system, Oracle Real Application Clusters, third party software, and so on) and other related entities. These can then be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manger Grid Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

For example, imagine Software Library to be the container inside the coffee vending machine that has smaller compartments for storing milk, coffee powder, and sugar. So, before you press the "Coffee" button, you would ideally stock milk, coffee powder, and sugar in those smaller compartments of the container inside the vending machine.

The Software Library is part of and within Enterprise Manager Grid Control. It is not a separate, add-on product that needs to be installed separately. You can use the Enterprise Manager Grid Control console to access the Software Library and create components, images, directives, and store metadata and binary content.

For information about components, images, and directives, see Section 2.5.2, "Components", Section 2.5.3, "Directives", and Section 2.5.4, "Images, Network Templates, Hardware Templates, Disk Layouts", respectively.

For information about setting up the Software Library, see Section 3.1.2, "Setting Up Oracle Software Library". For information about using other options offered by the Software Library, see Appendix C, "Using Oracle Software Library".

2.5.2 Components

Components are entities in the Software Library that represent the primary building blocks for Deployment Procedures. A component can represent operating system software, Oracle software, or any third-party software and applications.

For example, imagine a Component to be an ingredient, such as milk, coffee powder, or sugar stored in the container inside a coffee vending machine. Without these, the vending machine cannot prepare a cup of coffee.

Components are individually maintained within the Software Library, and you can associate versions, states, and maturity levels with each component. They may also be combined with other components as needed, to specify the complete software configuration or image that is provisioned on target hosts.

2.5.3 Directives

Directives are entities in the Software Library that represent a set of instructions to be performed. These are constructs used to associate scripts with software components and images. These scripts contain directions on how to interpret and process the contents of a particular component or an image.

For example, imagine a Directive to be the program that runs inside the coffee vending machine to automate the process of preparing coffee. The program essentially instructs the vending machine to add the ingredients one by one, and then serve the drink through its nozzle.

Directives encapsulate the script, the command line used to invoke the script, and the script configuration properties. They capture everything required for invoking the script on a machine during a provisioning operation.

Directives are usually categorized based on the provisioning life cycle phases they are targeted for, or the actions they perform. Imagine Directives as set of executable instructions that run from a supported shell (for example, borne-again, Perl, Python), programming language (for example, Java), or execution framework or interpreter (such as "make" or "ant").

Directives are contained within a file stored in the Oracle Software Library and referenced from the software components that employ them. Versions, states and maturity levels can be associated with each directive.

2.5.4 Images, Network Templates, Hardware Templates, Disk Layouts

Images are entities that represent a collection of components along with the necessary directives that create a deployable configuration for a single or set of target hosts.

Network Templates, Hardware Templates, and Disk Layouts are used to define the network, hardware, and disk layout configuration of the target hosts, respectively.

For example, imagine an Image to be a cappuccino flavored drink readily available in a paper carton. In this case, you do not prepare the coffee; the drink already contains milk, cappuccino coffee powder, and sugar, and it's all mixed and readily available for consumption.

Images are stored in the Oracle Software Library and versions, states, and maturity levels can be associated with them.

2.5.5 Assignments

Assignment is the activity to select an image and provision them on hardware servers. Assignments allow specifying properties (such as image properties, file system, root password, network profiles, port, IP address, etc.) that will be used during the provisioning of hardware servers. Assignments also allow for setting of advanced parameters like NTP, NIS, NFS and kernel parameters.

2.6 Understanding Linux Patching

Linux Host Patching is a feature in Enterprise Manager Grid Control that helps in keeping the machines in an enterprise updated with security fixes and critical bug fixes, especially in a data centre or a server farm. This feature support in Enterprise Manager Grid Control enables you to:

- Set up Linux RPM Repository based in Unbreakable Linux Network(ULN) channels

- Download Advisories (Erratas) from ULN

- Set up 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 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:

Linux Host	A host target in Enterprise Manager Grid Control that is running the Linux operating system.
Linux Patching Group	A set of managed Linux hosts that are associated with a common RPM repository. 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 Enterprise Linux.
ULN Channel	A channel is a grouping of RPM packages on the ULN network. For example, el4_latest channel contains all packages for OEL 4.
RPM Repository	RPM repository is a directory that contains RPM packages and their metadata (extracted by running yum-arch and createrepo). The RPM repository is accessible via http or ftp. An RPM repository can be organized to contain packages from multiple channels.
	For example, /var/www/html/yum/Enterprise/EL4/latest might contain packages from the el4_latest 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.

2.7 Understanding SOA Artifacts Provisioning

SOA artifacts deployment proecdures support provisioning of SOA composites, Web Service policies, and policy and credential stores.

Following are some of the terms used in SOA artifacts provisioning:

SOA Composite

A SOA composite is a logical construct. Its components can run in a single process on a single computer or be distributed across multiple processes on multiple computers. A complete application might be constructed from just one composite, or it could combine several different composites. The components making up each composite might all use the same technology, or they might be built using different technologies.

SOA Infra Domain

The SOA Infrastructure domain is a WebLogic domain that contains soa-infra binaries, The SOA Infrastructure includes a set of service engines (BPEL process, human workflow, decision service, and Oracle mediator) that execute the business logic of their respective components within the SOA composite application (for example, a BPEL process).

Web Services

A Web service is a program that can be accessed remotely using different XML-based languages. What this program can do (that is, the functionality it implements) is described in a standard XML vocabulary called Web Services Description Language (WSDL). For example, a banking Web service may implement functions to check an account, print a statement, and deposit and withdraw funds. These functions are described in a WSDL file that any consumer can invoke to access the banking Web service. As a result, a consumer does not have to know anything more about a Web service than the WSDL file that describes what it can do.

A Web service consumer (such as, a desktop application or a Java Platform, Enterprise Edition client such as a portlet) invokes a Web service by submitting a request in the form of an XML document to a Web service provider. The Web service provider processes the request and returns the result to the Web service consumer in an XML document.

WS Policies and Assertions

Policies describe the capabilities and requirements of a Web service such as whether and how a message must be secured, whether and how a message must be delivered reliably, and so on. Policies belong to one of the following categories: Reliable Messaging, Management, WS-Addressing, Security, and MTOM.

Policies are comprised of one or more assertions. A policy assertion is the smallest unit of a policy that performs a specific action. Policy assertions are executed on the request message and the response message, and the same set of assertions is executed on both types of messages. The assertions are executed in the order in which they appear in the policy. Assertions, like policies, belong to one of the following categories: Reliable Messaging, Management, WS-Addressing, Security, and MTOM.

Policy Stores

The Policy Store is a repository of system and application-specific policies and roles. Application roles can include enterprise users and groups specific to the application (such as administrative roles). A policy can use any of these groups or users as principals. A policy store can be file-based or LDAP-based. A file-based policy store is

an XML file, and this store is the out-of-the-box policy store provider. An LDAP-based policy store can use either of the following LDAP servers: Oracle Internet Directory or Oracle Virtual Directory (with a local store adapter, or LSA).

Credential Stores

A Credential Store is a repository of security data (credentials) that certify the authority of users, Java components, and system components. A credential can hold user name and password combinations, tickets, or public key certificates. This data is used during authentication, when principals are populated in subjects, and, further, during authorization, when determining what actions the subject can perform.

2.8 Understanding Virtualization Management Provisioning

Virtualization management involves the monitoring, administration, and maintenance of virtual servers and guest virtual machines in your enterprise. Oracle provides the customer increased value by extending Enterprise Manager capabilities to monitor virtualized entities alongside the physical infrastructure and perform complete lifecycle management of virtual servers and software running on them.

This section describes the following:

- Terms Used in Virtualization

- Privileges Required for Administrative and Provisioning Operations

2.8.1 Terms Used in Virtualization

The following describes the terms used in virtualization:

Oracle VM

Oracle VM is server virtualization software. Installing Oracle VM on a bare metal box allows for the creation of multiple guest virtual machines on it. Just like Oracle databases and Oracle Fusion Middleware, Oracle VM is available for download on Oracle Technology Network (OTN).

Virtual Server

Virtual Server is a generic term used to describe a physical box running virtualization software (hypervisor) on it. A new virtual server can be provisioned by installing Oracle VM server software on a bare metal physical box. "Virtual Server" is a target type in Enterprise Manager that represents the Oracle VM targets.

Every Oracle VM server can perform one or more of the functions described below:

- **Master Server**

  The Master Server is the core of the server pool operations. It acts as the contact point of the server pool to the outside world, and also as the dispatcher to other servers within the server pool. The load balancing is implemented by the Master Server. For example, when you start a guest virtual machine, the Master Server will choose a Guest VM Server with the maximum resources available to run the guest virtual machine. There can be only one Master Server in a server pool. The state of a virtual server pool is equivalent to the state of the master virtual server.

- **Utility Server**

  The Utility Server is responsible for I/O intensive operations such as copying or moving files. Its function focuses on the creation and removal operations of guest virtual machines, virtual servers, and server pools. There can be one or more

Utility Servers in a server pool. When there are several Utility Servers, Server Pool Master chooses the Utility Server with the maximum CPU resources available to conduct the task.

- **Guest VM Server**

  The primary function of the Guest VM Server is to run guest virtual machines, thus acting as a hypervisor. There can be one or more Guest VM Servers in a server pool. When there are several Guest VM Servers, Master Server chooses the Guest VM Server with the maxhmum resources available (including memory and CPU) to start and run the virtual machine.

Monitoring Server

The monitoring server monitors the virtual server remotely. Multiple virtual servers are monitored by one agent. Enterprise Manager agent must not be installed on virtual servers. The Monitoring Server agent (Enterprise Manager agent that monitors Oracle VM servers) should be upgraded to version 10.2.0.5.0.

Virtual Server Pool

A Server Pool is a logical grouping of one or more virtual servers that share common storage. A virtual server can belong to only one virtual server pool at a time. Guest virtual machines and resources are also associated with server pools. Oracle VM Server Pool is an aggregate target type in Enterprise Manager to represent the server pool of Oracle VM Servers.

When the Oracle VM Server Pool is created, the user is asked to provide the details of the Master Server for that pool. By default, this Oracle VM Server also performs the functions of the Utility Server and Guest VM Server.

The user can later change the Utility Server and Guest VM Server functions using the Edit Virtual Server action.

Guest Virtual Machine

Guest Virtual Machine (also known as Guest VM) is the container running on top of a virtual server. Multiple guest virtual machines can run on a single virtual server. Guest virtual machines can be created from Oracle VM templates. Oracle VM templates provide pre-installed and pre-configured software images to deploy a fully configured software stack. Oracle VM templates can be downloaded from OTN at the following location:

http://www.oracle.com/technology/products/vm/templates.html

A guest virtual machine monitored by an Enterprise Manager agent is treated like any other Host target. This type of guest virtual machine is referred to as a "Managed Guest VM".

A guest virtual machine that is not monitored by any Enterprise Manager agent is referred to as an "Unmanaged Guest VM". Only limited monitoring and configuration information is available for such a guest virtual machine.

2.8.2 Privileges Required for Administrative and Provisioning Operations

Table 2–1 lists the privileges required to perform administrative and provisioning tasks on server pools, virtual servers, and guest virtual machines.

Table 2–1   Privileges for Provisioning and Administrative Operations in Virtualization Central

Task	Privileges
Create Virtual Server Pool	Add any Target system privilege
Remove Virtual Server Pool	Full privilege on the Server Pool Target and the member Virtual Servers
Manage Virtual Server Pool	Operator privilege on the Server Pool Target
Create Guest Virtual Machine	Operator privilege on the Server Pool Target
Discover Guest Virtual Machine	Operator privilege on the Server Pool Target
Clear Operation Error (Virtual Server Pool)	The Privilege check depends on the operation that has failed. - If Create Virtual Server Pool action has failed: User must be the job owner or must have Super Administrator privileges - If Remove Virtual Server Pool has failed: User must have Full privilege on the Server Pool Target and the member Virtual Servers as well - If Stop Virtual Server / Reboot Virtual Server operation on the Master Virtual Server had failed: User must have Operator privilege on the Master Virtual Server and the Server Pool Target
Convert P2V	Operator privilege on the Server Pool Target
Convert V2V	Operator privilege on the Server Pool Target
Register Virtual Server	Add any Target system privilege
De-register Virtual Server	Full privilege on the Virtual Server and the Server Pool Target
Edit Virtual Server	Operator privilege on the Virtual Server and the Server Pool Target
Enable Server Maintenance Mode	Operator privilege on the Virtual Server
Remove Server Maintenance Mode	Operator privilege on the Virtual Server
Live Migrate All Guest Virtual Machine	Operator privilege on the source Virtual Server and the Server Pool Target. Since the destination virtual server gets selected automatically and may be any virtual server in the Server Pool, Operator privilege on all member virtual servers of the Server Pool Target is essential. If preferred servers are set for the guest virtual machine, you will need Operator privilege on preferred virtual servers only.
Stop Virtual Server	Operator privilege on the Virtual Server and the Server Pool Target
Reboot Virtual Server	Operator privilege on the Virtual Server and the Server Pool Target

Table 2–1   (Cont.) Privileges for Provisioning and Administrative Operations in Virtualization Central

Task	Privileges
Clear Operation Error (Virtual Server)	Privilege check depends on the operation that has failed. ■ If Register Virtual Server has failed: User must be the job owner or must have Super Administrator privileges ■ If De-register Virtual Server has failed: User must have Full privilege on the Server Pool Target and the Virtual Servers as well ■ If Stop Virtual Server/Reboot Virtual Server operation has failed: User must have Operator privilege on the Master Virtual Server and the Server Pool Target
Live Migrate Guest Virtual Machine	Operator privilege on the source Virtual Server and the Server Pool Target. Since the destination virtual server gets selected automatically and may be any virtual server in the Server Pool, Operator privilege on all member virtual servers of the Server Pool Target is needed. If preferred servers are set for the guest virtual machine, you will need Operator privilege on preferred virtual servers only.
Clone Guest Virtual Machine	Operator privilege on the Server Pool Target
Save Guest Virtual Machine as Template	Operator privilege on the Server Pool Target
Edit Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target. An Edit Guest VM action may be performed on a Halted guest virtual machine as well. In this case privileges will be checked only on the Server Pool Target.
Update Preferred Server List	Operator privilege on the Virtual Server and the Server Pool Target. If the preferred server selection is set to Automatic, the guest virtual machine may come up on any virtual server within the server pool hence, you must have Operator privilege on all the members virtual servers in the server pool Target. If the preferred server selection is set to Manual, you must have Operator privilege on the specified list of member virtual servers.
Delete Guest Virtual Machine	Full privilege on the Server Pool Target
Stop Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target

Table 2–1 (Cont.) Privileges for Provisioning and Administrative Operations in Virtualization Central

Task	Privileges
Start Guest Virtual Machine	Operator privilege on the Server Pool Target.
	If the preferred server selection for that guest virtual machine is set to Automatic, the guest virtual machine may come up on any virtual server within the server pool. Hence, you must have Operator privilege on all the members virtual servers in the Server pool Target.
	If the preferred servers for that guest virtual machine is set to Manual, you must have Operator privilege on the specified list of member virtual servers.
Reboot Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target.
Pause Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target
Un-pause Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target
Suspend Guest Virtual Machine	Operator privilege on the Virtual Server and the Server Pool Target
Resume Guest Virtual Machine	Operator privilege on the Server Pool Target.
	If the preferred server selection for that guest virtual machine is set to Automatic, the guest virtual machine can come up on any virtual server within the server pool. Hence, you must have Operator privilege on all the members virtual servers in the Server pool Target.
	If the preferred server selection for that guest virtual machine is set to Manual, you must have Operator privilege on the specified list of member virtual servers.
VNC Console	Operator privilege on the Virtual Server and the Server Pool Target
Clear Operation Error (Guest Virtual Machines)	Privilege check depends on the operation that has failed.
	▪ If Delete Guest VM has failed:
	User must have Full privilege the Server Pool Target
	▪ If Start/Stop/Reboot/Pause/Un-pause/Suspend/Resume Guest VM has failed:
	User must have Operator privilege the Server Pool Target
	▪ If Create/Clone/Migrate/Edit/Save has failed:
	User must have Operator privilege the Server Pool Target
View Virtual Server/Guest Virtual Machine	View privilege on the Virtual Server Pool.
	If the User has Full/ Operator/View privilege on the Virtual Server Pool then the user gets View privileges on all the member Virtual Server targets automatically.

2.9 Understanding 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 the Software and Server Management challenge is its Bare Metal Provisioning Application, which is a part of Provisioning Pack (available at http://www.oracle.com/technology/products/oem/mgmt_ solutions/provisioning.html) of the Enterprise Manager Grid Control and Oracle Management Pack for Linux (available at http://www.oracle.com/technology/products/oem/omp_linux.html). Bare Metal Provisioning Application addresses the data center, server farm challenge to provision software and servers quickly, efficiently, and make them operational.

Bare Metal Provisioning Application provides server lifecycle management capabilities that enable one to build manage and optimize their server infrastructure. The application provides an automated, repeatable and reliable solution that:

- Automates deployment of consistent, certified Linux operating system images along with other software on a larger number of servers.

- Leads to faster, unattended deployment of software and operating systems.

- Allows provisioning of middleware, Clusterware, Real Application cluster (RAC) etc. on top of Linux stack.

- Supports heterogeneous hardware and network configuration.

- Automatically discovers bare metal and live target servers for provisioning.

- Especially for Oracle software the application encodes best practices out-of-the-box for deployment and patching.

- Results in 10x or more reduction in manual labor that leads to substantial cost savings.

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.

2.9.1 Overview of the Bare Metal Provisioning Environment

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 Enterprise Manager and Provisioning Application, the following need to be setup and configured before using the provisioning application.

Software Library and its Entities

For information about Software Library and its entities, see Section 2.5.1, "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 Enterprise Manager 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 Section 3.7.3, "Setting Up Boot Server" for setting up a boot server with DHCP/TFTP combination.

Also refer to section Section 3.7.12, "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. Refer to Appendix G, "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. 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 Enterprise Manager target. Refer to section Section 3.7.4, "Setting Up Stage Server" for setting up a stage server. Also refer to section Section 3.7.10, "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 Enterprise Manager 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. Refer to section Section 3.5.2, "Setting Up the RPM Repository" for setting up a RPM repository. Also refer to section Section 3.7.11, "Configuring RPM Repository".

2.9.2 Overview of the Bare Metal Provisioning Process

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

Setting Up Provisioning Environment (Section 3.7, "Setting Up Infrastructure for Bare Metal Provisioning"):

- Setting up Boot/DHCP server and Stage server, setting up RPM repository and Software Library

- Configuring above entities in Enterprise Manager

Provisioning Linux using Bare Metal Provisioning Application (Section 11.3, "Provisioning Bare Metal Boxes"):

- Creating Linux "Default Image" by creating an image and associating the image with a particular network subnet or MAC address

- Staging the image on the Stage server

■ Powering up the bare metal machine on the network to begin the PXE-based OS boot and install process

2.9.2.1 Setting Up Provisioning Environment

The user needs to perform a one-time activity of setting up a Boot server, Stage server, and RPM repository that are required by the provisioning application. Once this infrastructure is set up, it can be used repeatedly for provisioning Linux on several bare metal machines. See Section 3.7, "Setting Up Infrastructure for Bare Metal Provisioning" for more information.

2.9.2.2 Provisioning Linux Using Bare Metal Provisioning Application

The following sections outline the steps to provision Linux. See Chapter 11, "Provisioning Linux Operating System" for more information.

Figure 2–1   User Flow for Creating Deployable Images for Hardware Servers
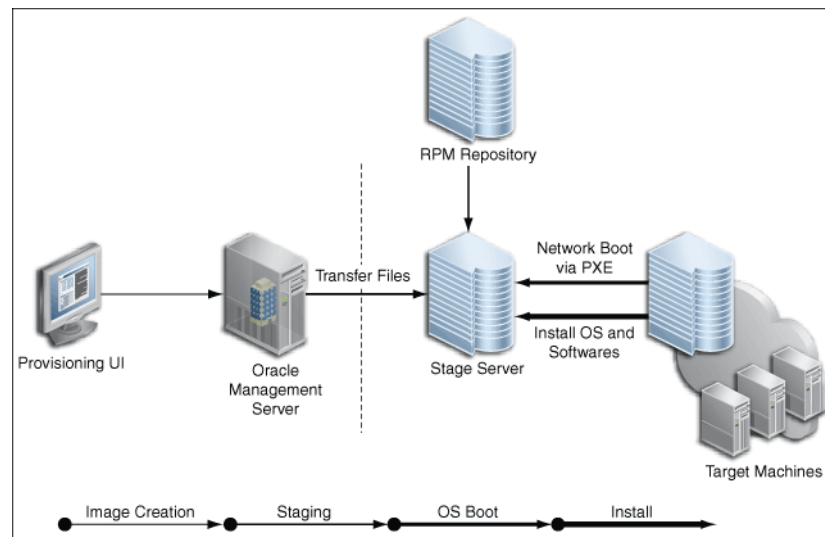
2.9.2.2.1   Creating Linux Default Image

Enterprise Manager allows the user to create a Linux "Default Image" from Linux installation on a reference machine. The default image consists of an OS binary component along with directives required to stage and provision the image. Once an image is created, it can be associated with a subnet prefix or a set of MAC addresses. Doing this ensures that the image gets provisioned on bare metal machines in the specified subnet or having specified MAC address. See Chapter 11.3.2, "Creating a Default Image" for more information.

2.9.2.2.2   Staging the Image on the Stage Server

The user is required to stage the contents of the image, namely the binaries associated with the components, directives and other templates, on the stage server. This caches the image in the staging area in the subnet for improved performance and prepares the image to be provisioned. See Chapter 11.3.4, "Staging the Default Image" for more information.

2.9.2.2.3   Powering up the Bare Metal Machine on the Network to Begin the PXE-based OS Boot and Install Process

Once the image is staged, it is ready to be provisioned on bare metal machines that are powered on in the subnet. When a bare metal machine is connected to the subnet, the Bare Metal Provisioning Application discovers the machine and boots it over the network using the PXE protocol. It then installs the operating system and the Enterprise Manager management agent on the server. Installing the agent makes the server a managed target in the Enterprise Manager console. See Chapter 11.3.5, "Ready to Go" for more information.

Figure 2–2 explains the sequence in the PXE process:

Figure 2–2   PXE Process
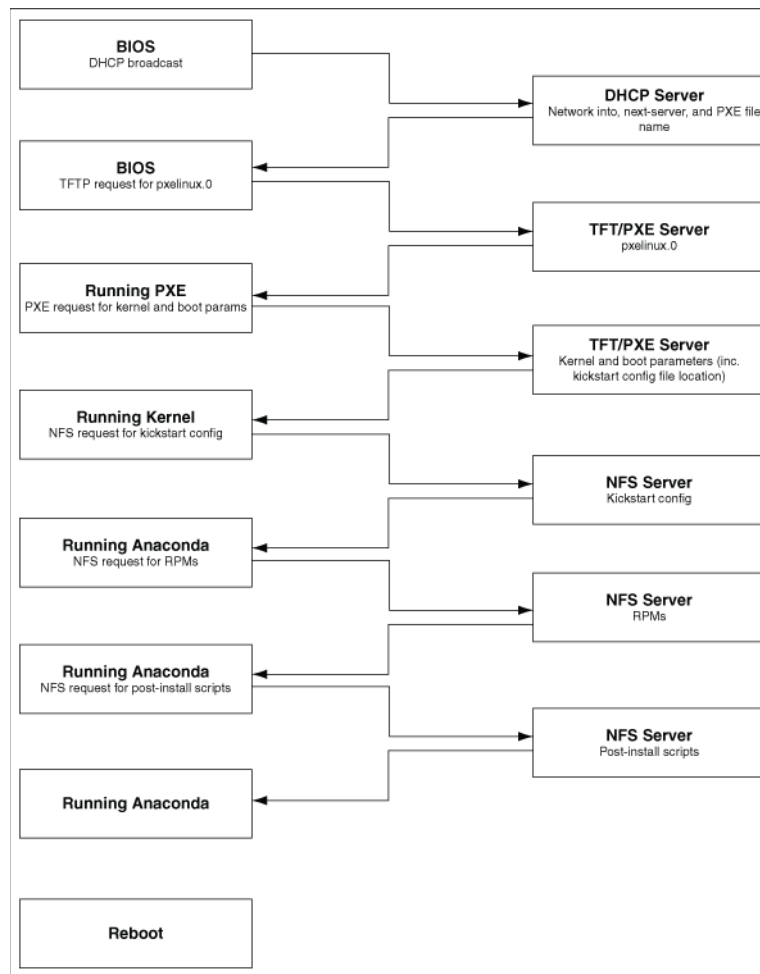
See Chapter G, "PXE Booting and Kickstart Technology" for a description of the PXE boot process.

2.10 Understanding My Oracle Support

My Oracle Support provides you with personalized, proactive, and collaborative support. You can benefit from integrated service request work flows and detailed, dynamic system configuration data, shared with Oracle Support in real-time to improve system stability and resolve problems more easily. Using the configuration manager, you can share configuration information with Oracle Support. When this is done, Oracle Support engineers can view your configuration information through My

Oracle Support to assist with diagnosing and resolving system-critical issues and help you to avoid problems before they occur. My Oracle Support functionality is fully integrated with Oracle Enterprise Manager.

This section describes the following:

- Patches and Updates
- Patch Plans
- Patch Recommendations
- Validation
- Merge Requests

2.10.1 Patches and Updates

From the Enterprise Manager Grid Control console, the Patches and Updates tab allows you to view and download recommended patches and updates for your Oracle products.

2.10.2 Patch Plans

A patch plan is a collection of patches which you might want to consider applying as a group. A patch plan could equate to a maintenance window (you can set a date for the plan), could contain patches for a single target, a collection of similarly configured targets (the same key patches applied to all of your databases, for example), or all of the patches for a system or group of systems, such as all patches for the current month for production databases.

You can include a patch for any target in a plan. The plan also supports validation of the patches against your environment by looking for conflicts with existing (installed) patches, missing pre-requisites for patches you want to install (a candidate patch), as well as looking for conflicts between patches within the plan itself. Currently, the validation process works for products such as Oracle Database, Fusion Middleware and Enterprise Manager. More Oracle products will be added in future release.

Patch plan is only available when using the configuration manager.

2.10.3 Patch Recommendations

Patch recommendations are proactive notifications of potential system issues and recommendations that help you improve system performance and avert outages. My Oracle Support's Patch Recommendations region provides a portal to all recommended patches. From the bar graph, you can drill down to a list of recommended patch, view details about those patches, download the patches, or add them to a patch plan. A bar graph summarizes the number of issues found (one issue = one recommendation for one target). The Recommended Patches region appears by default on the Patch and Updates page and can also be added to the Dashboard. You can edit this region to filter its contents.

2.10.4 Validation

When My Oracle Support validates a plan, it checks for conflicts using patch information from Oracle, the inventory of patches on your system (gathered by the configuration manager) and information from the candidate patches. In addition to checking patch merge conflicts between patches listed in the plan, the validation

function also checks for conflicts between the listed plan patches and patches that have already been applied.

Not all Oracle products and platforms support validation, but those that do (such as the Oracle Database, Oracle Fusion Middleware , and Oracle Enterprise Manager) will help you later when you want to deploy the patches. The validation process will catch a variety of problems which are easier to resolve now then when you are trying to deploy the patches. It catches patches that conflict and require a replacement patch (also referred to as a Merged Patch, or MLR), missing pre-requisites, patches not intended for a specific release, and other conditions that would cause problems at installation time.

2.10.5 Merge Requests

When patching complex heterogeneous IT environments, enevitably conflicts may arise between specific patch versions. Conflicts occur when two or more patches attempt to update/change the same files. Because individual files cannot be applied as one-offs, they must be "merged" together in a single patch containing all changes/updates to the file(s) in question so that the file(s) are updated only once with all the fixes.

Merge Requests are also known as Merge Label Requests (MLRs).

2.11 Understanding Other Related Concepts

This section describes the following:

- Provisioning Archive (PAR) Files

- Oracle Patch Advisories

- Staging Area

- Support for SUDO, PAM, and Privilege Delegation

- Tracking the Status of Deployment Procedures

2.11.1 Provisioning Archive (PAR) Files

Provisioning Archive files or PAR files contain Deployment Procedures and/or components and directives from the Software Library. These are default Deployment Procedures offered by Enterprise Manager Grid Control. When you import or export Deployment Procedures using the default PARDeploy utility, you technically import or export PAR files.

For information on PARDeploy utility and how you can use it to import or export PAR files, see Appendix B, "Using emctl partool Utility".

2.11.2 Oracle Patch Advisories

Oracle Patch Advisories describe critical software patches for Oracle products. To help ensure a secure and reliable configuration, all relevant and current critical Oracle patches must be applied.

To promote critical patch application, Enterprise Manager Grid Control performs an assessment of vulnerabilities by examining the host configurations collected for your enterprise to determine the Oracle homes that require one or more critical patches to be installed.

2.11.3 Staging Area

Staging Area is a temporary location on the target host where the generic component from the Software Library or the reference installation used for cloning will be staged for a short period, and then extracted and used for deployment.

The generic component or the reference installation may be in the form of a ZIP file. After it is placed in the staging area, once the Deployment Procedure is ready for deployment, the contents of the ZIP file are extracted and then deployed. You must always ensure that the staging area has write permission so that the Deployment Procedure can stage the ZIP file.

2.11.4 Support for SUDO, PAM, and Privilege Delegation

All the Deployment Procedures offered by Enterprise Manager Grid Control require administrator privileges to run. While most steps within a Deployment Procedure can be run as a normal user, there are some steps that require special permissions and privileges, and unless you provide the administrator's credentials, you cannot proceed with the deployment.

Under such circumstances, you can do one of the following:

■   Customize the Deployment Procedure to disable the steps that require special privileges, run the other steps as a normal user, and have the administrator run the disabled steps later.

■   Customize the Deployment Procedure to use authentication utilities for running some steps within the Deployment Procedure with the privileges of another user. The authentication utilities supported by Enterprise Manager Grid Control are SUDO, PowerBroker, and Privilege Delegation.

For information on how you can customize Deployment Procedures and how Privilege Delegation scores over other two utilities, see Chapter 31.7.1, "Using Privilege Delegation".

2.11.5 Tracking the Status of Deployment Procedures

After you have submitted a Deployment Procedure, you can track its status from the Procedure Completion Status page. To access this page, follow these steps:

1.   In Grid Control, click **Deployments**.

2.   On the Deployments page, from the Deployment Procedure Manager section, click **Procedure Completion Status**.

3.   On the Deployment Procedure Manager page, click the **Procedure Completion Status** tab.

4.   On the Procedure Completion Status page, you can view the status of the Deployment Procedures.

Table 2–2   Deployment Procedure Status

Status	Description
Scheduled	implies that the Deployment Procedure is scheduled to be executed at the date and time that you have specified.
Running	implies that the Deployment Procedure is currently being executed.
Action Required	Implies that the Deployment Procedure has stopped running as user interaction is required.

Table 2–2   (Cont.) Deployment Procedure Status

Status	Description
Suspended	Implies that the Deployment Procedure has been temporarily stopped from execution. You can resume from the paused step later.
Failed	Implies that the Deployment Procedure has failed and cannot execute the further steps. You can also ignore the error and proceed further.
Completed	Implies that the Deployment Procedure has successfully completed its execution.
Stopped	Implies that the Deployment Procedure has been permanently stopped from execution by the user.
Saved	Implies that the Deployment Procedure has not been given for execution, and has been saved.
Completed with Errors	Indicates that the Deployment Procedure completed, but completed with errors possibly because some of the steps within it might have failed and the steps might have the *Skip target/Continue on Error* flag enabled.

You can also perform the following actions:

a. Search for a particular Deployment Procedure using the **Search** section. Click **Search** to refine your search criteria.

b. View the status of all Deployment Procedures. You can also manually refresh the page and view the updated status by clicking **Refresh**.

c. View real-time status information based on a particular refresh period such as 30 seconds, 1 minute, or 5 minutes.

d. Stop or suspend any Deployment Procedure by selecting them and clicking **Stop** or **Suspend**, respectively. You can resume at any point by clicking **Resume** or **Retry**.

e. Delete any Deployment Procedure by selecting them and clicking **Delete**.

3

Infrastructure Requirements

Before you run any of the Deployment Procedures, meet the infrastructure requirements described in this chapter. This chapter is essentially for administrators or designers who create the infrastructure for provisioning and patching. The requirements described in this chapter have to be performed just once.

This chapter covers the following:

- Mandatory Infrastructure Requirements for Provisioning and Patching
- Optional Infrastructure Requirements
- Setting Up Infrastructure for Online Patching
- Setting Up Infrastructure for Offline Patching
- Setting Up Infrastructure for Linux Patching
- Setting Up Infrastructure for Virtualization Systems
- Setting Up Infrastructure for Bare Metal Provisioning

3.1 Mandatory Infrastructure Requirements for Provisioning and Patching

This section describes the mandatory setup requirements you must meet before running any Deployment Procedure. In particular, this section describes the following:

- Creating Designer and Operator User Accounts
- Setting Up Oracle Software Library
- Applying Patches
- Discovering and Monitoring Hosts and Targets

3.1.1 Creating Designer and Operator User Accounts

The main users of Deployment Procedures are *Designers* (Lead Administrators) and *Operators*. *Designers* are users who perform design-time activities and *Operators* are users who perform run-time activities. For more information about these user accounts, see Section 2.2.4, "Users of Deployment Procedures".

Before you begin using Deployment Procedures, Oracle recommends you to create these user accounts in Enterprise Manager Grid Control so that *Designers* have super administrator privileges to perform advanced tasks, and *Operators* have basic privileges to run Deployment Procedures.

This section describes the following:

- Creating Designer User Account
- Creating Operator User Account

> **Note:** Even if you do not create the designer or operator user account, you will be allowed to run the Deployment Procedure as long as you have super administrator privileges.

3.1.1.1 Creating Designer User Account

To create a *Designer* user account, follow these steps:

1. In Grid Control, click **Setup**.

2. On the Overview of Setup page, from the vertical menu on the left, click **Administrators**.

3. On the Administrators page, click **Create**.

4. In the Create Administrator wizard, do the following:

   a. On the Properties page, specify the name *Designer*, provide a password, check **Super Administrator**, and click **Next**.

   b. On the Review page, review the information you have provided for this user account, and click **Finish**.

3.1.1.2 Creating Operator User Account

To create a *Operator* user account, follow these steps:

1. In Grid Control, click **Setup**.

2. On the Overview of Setup page, from the vertical menu on the left, click **Administrators**.

3. On the Administrators page, click **Create**.

4. In the Create Administrator wizard, do the following:

   a. On the Properties page, specify the name *Operator* and provide a password. Leave the other fields blank and click **Next**.

b. On the Roles page, select **Public**, and click **Next**.

c. On the System Privileges page, select the system privileges that you want to grant to this operator user account.

d. On the Target Privileges page, select the targets that the operator user account will access.

e. On the Review page, review the information you have provided for this user account, and click **Finish**.

3.1.2 Setting Up Oracle Software Library

Oracle Software Library (Software Library) can be configured on any mounted file system that is readable and writable from Oracle Management Service (OMS).

- If you have a single OMS, then you can configure the Software Library either on the host where the OMS is running or in a shared location.

  > **Note:** If you have plans to have multiple OMSes in the future, then Oracle strongly recommends you to configure the Software Library in a shared location.

  If you configure the Software Library in a shared location, then use the Check Availability feature in the Software Library section in the Administration page to verify that it is accessible by the OMS. For more details, see Appendix C, "Using Oracle Software Library".

  Ensure that there is enough space available for the storage of software binaries, and associated scripts for the entities that you want to create and store.

- If you have multiple OMSes, then you can configure the Software Library either on one of the hosts where the OMS is running or in a shared location, but you must ensure that the Software Library is accessible by all the OMSes from the configured location.

  You can use the Check Accessibility feature in the Software Library section in the Administration page to verify that the Software Library location is accessible by all the OMSes. For more details, see Appendix C, "Using Oracle Software Library".

Ensure that there is at least 1 GB to 2 GB of space available for successfully configuring the Software Library. Understandably, this space requirement increases over a period of time as you create more components and gold images. For example, if you are creating a gold image of a database home of 1 GB, then ensure that you have at least 2 GB of space.

If you want to increase the space available in the Software Library to store entity files or binaries, then do one of the following:

– **Method A:** Manually add more space to that folder or to that shared location where you have configured the Software Library. Ensure that you do not change the Software Library location. After increasing the space manually, go to the Provisioning page, click the **Administration** tab, and in the Software Library Configuration section, click **Refresh**.

– **Method B:** Add another Software Library location by clicking **Add** so that all all the new images, components, and so on are stored in this new location.

3.1.2.1 Configuring Oracle Software Library

To configure the Software Library, follow these steps:

1. In Grid Control, click **Deployments**, and then the **Provisioning** subtab.

    There are a number of tabs here for creating components, directives and other entities. You can access some or all of the tabs depending on the privileges assigned to you.

2. Click the **Administration** tab. This requires super administrator privileges similar to Designer user account/SYSMAN user.

    For information about creating the super administrator, see Section 3.7.2, "Creating Super Administrator for Enterprise Manager".

3. In the Software Library Configuration section of the **Administration** tab, click **Add**.

4. On the Add Software Library Location page, enter the directory location and then click **OK**.

> **Note:** When you add a Software Library location for the first time, the configuration might take some time. While the location is being configured, Enterprise Manager Grid Control displays a progress page that shows information about the PAR files being uploaded. However, this behavior is expected only when the location is being configured for the first time. Subsequently, adding other Software Library locations will be quicker.

When the Software Library is configured, out-of-box Provisioning Archive files (PAR files) will be deployed. These files contain pre-build entities such as components, directives, and so on for various applications such as bare metal provisioning and patching.

For information about PAR files, see Section 2.11.1, "Provisioning Archive (PAR) Files". For more information about Software Library, see Appendix C, "Using Oracle Software Library".

3.1.2.2 Limitations and Known Issues

The following are the limitations and known issues related to the Software Library:

- Relocation of the Software Library is not supported from the Enterprise Manager Grid Control console.

- The directory `/tmp` on Linux and Unix, or `C:\TEMP` on Microsoft Windows must not be used as the Software Library location.

- (For Microsoft Windows Only) In order to have the Software Library configured across multiple OMSes, you must have a disk partition with the same name for all the OMS hosts and have a cluster file system running on the partition. This can be accomplished by installing OCFS via an Oracle CRS version 10.2.0.1 installation.

  The OMS service must be configured to login as a domain user account, and the domain user account must have the required permissions for accessing the shared storage locations. All the OMS services must login using the same domain user account.

3.1.3 Applying Patches

Apply the patches required for provisioning and patching as described in My Oracle Support note 427577.1.

To view this note, access My Oracle Support at the following URL. On the main page (based on the classic metalink view), from the **Quick Find** list, select **Document ID (Knowledge Base, Forum, or Bug)**, and quote the note number and click **Go**.

http://support.oracle.com

For more information about these patches, view the ReadMe file associated with the patch.

3.1.4 Discovering and Monitoring Hosts and Targets

Before running a provisioning-related Deployment Procedure, ensure that the target hosts are discovered and monitored in Enterprise Manager Grid Control. Similarly, before running a patching-related Deployment Procedure, ensure that the targets you want to patch are discovered and monitored in Enterprise Manager Grid Control.

For this purpose, you require **Oracle Management Agent 11g Release 1 (or higher)** to be installed on the target hosts where you want to run the procedure. For information about installing Oracle Management Agent 11g Release 1 (or higher), see *Oracle Enterprise Manager Grid Control Basic Installation Guide* or *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide* available at:

http://www.oracle.com/technology/documentation/oem.html

In the case of Oracle RAC and Oracle Clusterware, before you run the Deployment Procedures to provision or extend them, ensure that you have cluster-based Oracle Management Agent (Management Agent) on every node of the cluster.

- To install cluster-based Management Agent on every node of the cluster, see the *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide* available at:

  http://www.oracle.com/technology/documentation/oem.html

- If you have standalone Management Agent installed on every node of the cluster, then convert it to cluster-based Management Agent. For instructions to convert, see Appendix I, "Converting Standalone Oracle Management Agent to Cluster Agent".

3.1.5 Verifying Communication between OMS and Management Agent

Ensure that the OMS and the Management Agent are able to communicate with each other. To verify this, check the status of the Management Agent:

- Navigate to the Oracle home of the OMS and run the following command to see a message that confirms that OMS is up and running.

  ```
  $<ORACLE_HOME>/bin/emctl status oms
  ```

- Navigate to the Oracle home of the Management Agent and run the following command to see a message that confirms that the Management Agent is up and running:

  ```
  $<ORACLE_HOME>/bin/emctl status agent
  ```

- In Grid Control, click **All Targets**. On the All Targets page, search and click the Management Agent, and on the Home page of the Management Agent, check its status.

3.2 Optional Infrastructure Requirements

This section describes the optional setup requirements you may meet before running any Deployment Procedure. In particular, this section describes the following:

- Setting Up SUDO, PAM, Privilege Delegation Settings
- Setting Up E-mail Notifications
- Setting Up Preferred Credentials
- Refreshing Host Configuration

3.2.1 Setting Up SUDO, PAM, Privilege Delegation Settings

Enterprise Manager Grid Control allows you to run Deployment Procedures using authentication utilities such as SUDO, PowerBroker, and so on. This support is offered using the Privilege Delegation mechanism available in Enterprise Manager Grid Control. Privilege Delegation is a framework that allows you to use either SUDO or PowerBroker to perform an activity with the privileges of another user (locked accounts).

For information on how you can customize Deployment Procedures and use SUDO, PowerBroker, or Privilege Delegation settings, see Chapter 31.7, "Customizing Steps to Run as Another User".

> **Note:** The certified SUDO versions are 1.6.7 to 1.6.9. The certified PBRUN versions are 4.0.8 and 5.x. Higher versions of these utilities may continue to work unless some fundamental changes have been introduced to their behavior.

3.2.2 Setting Up E-mail Notifications

Enterprise Manager Grid Control can send e-mail notification every time you run a Deployment Procedure. However, by default, Deployment Procedures do not have this feature enabled. To configure them to send e-mail notifications, you must customize the Deployment Procedure.

For information on how you can customize Deployment Procedures and set up e-mail notifications, see Chapter 31.10, "Setting Up E-Mail Notifications".

3.2.3 Setting Up Preferred Credentials

Before you run any Deployment Procedure, Oracle recommends you to register the credentials of the host on which you want to provision or the Oracle home that you want to patch, as preferred credentials in Enterprise Manager Grid Control. This section explains how you can set up preferred credentials. In particular, this section explains the following:

- Setting Up Preferred Credentials for Provisioning

- Setting Up Preferred Credentials for Patching

3.2.3.1 Setting Up Preferred Credentials for Provisioning

To register the credentials of the host (on which you want to provision) as preferred credentials in Enterprise Manager Grid Control, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Preferences**.

2. On the General page, from the vertical menu pane to the left, click **Preferred Credentials**.

3. On the Preferred Credentials page, from the table, for the Host target type, click the Set Credentials icon that is shown in the column **Set Credentials**.

   For example, if you want to provision a database on a host named xyz.server.com, then for Host target type, click the Set Credentials icon.

   Enterprise Manager Grid Control displays the <Target Type> Preferred Credentials page.

   > **Note:** For setting up preferred credentials for virtual server targets, select Oracle VM Server as the target type and click the Set Credentials icon.

4. On the <Target Type> Preferred Credentials page, in the Target Credentials section, for the Host target on which you want to provision, specify the user name and password. If you want to use default credentials, then specify the default credentials in the Default Credentials section.

5. Click **Apply**.

3.2.3.2 Setting Up Preferred Credentials for Patching

To register the credentials of the Oracle home (that you want to patch) as preferred credentials in Enterprise Manager Grid Control, run the `set_credential` verb using the Enterprise Manager Command Line Interface (EM CLI).

> **Note:** If the credentials for all the databases and/or hosts are the same, then as an alternative to using the EMCLI verbs, you can register the common credentials of those targets as preferred credentials. This approach is easier than running the EMCLI verbs.
>
> To do so, click **Preferences** from the top-right corner of the console, and from the left menu bar, click **Preferred Credentials**. On the Preferred Credentials page, register the credentials as preferred credentials.

```
set_credential
```

```
                       -target_type="ttype"
                       [-target_name="tname"]
                       -credential_set="cred_set"
                       [-user="user"]
                       -columns="col1:newval1;col2:newval2;..."
                       [-input_file="tag1:file_path1;tag2:file_path2;..."]
                       [-oracle_homes="home1;home2"]
```

[ ] denotes that the parameter is optional.

Example 1:

```
emcli set_credential
      -target_type=oracle_database
      -target_name=myDB
      -credential_set=DBCredsNormal
      -user=admin1
      -column="username:joe;password:newPass;role:newRole"
```

Example 2:

```
emcli set_credential
      -target_type=host
      -target_name=host.foo.com
      -credential_set=OHCreds
      -user=admin1
      -column="OHUsername:joe;OHPassword:newPass"
      -oracle_homes="database1;mydb"
```

> **Note:** The preferred credentials set through EM CLI are specific to
> the Enterprise Manager Grid Control user with which the EM CLI is
> set up. Therefore, Oracle recommends you to use the central
> Enterprise Manager Grid Control accounts for this purpose, so the
> Database Administrators across board or a group can use this.
>
> For more information about EM CLI, see Enterprise Manager
> Command Line Interface Guide available at:
>
> http://www.oracle.com/technology/documentation/oem.h
> tml

The following describes the options available for this verb:

Table 3–1   EM CLI Options for set_credential Verb

Option	Description
target_type	Type of target. This must be "host" in case the -oracle_homes parameter is specified.
target_name	Name of the target. Omit this argument to set enterprise preferred credentials. This must be the host name in case the -oracle_homes parameter is specified.
credential_set	Credential set affected.
user	Enterprise Manager user whose credentials are affected. If omitted, the current user's credentials are affected.

Table 3–1   (Cont.) EM CLI Options for set_credential Verb

Option	Description
columns	Name and new value of the column(s) to set. Every column of the credential set must be specified. Alternatively, a tag from the -input_file argument can be used so that the credential values are not seen on the command line. You can specify this argument more than once.
input_file	Path of file that has the -columns argument(s). This option is used to hide passwords. Each path must be accompanied by a tag referenced in the -columns argument. You can specify this argument more than once.
oracle_homes	Name of oracle homes on the target host. Credentials will be added/updated for all specified homes.

3.2.4 Refreshing Host Configuration

Before you run any Deployment Procedure, Oracle recommends you to refresh the configuration of the hosts. To do so, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployments page, in the Configuration section, click **Refresh Host Configuration**.

3. On the Refresh Host Configuration: Select Hosts page, from the Available Hosts pane, select the hosts, which the Deployment Procedure will use, and move them to the Selected Hosts pane.

4. Click **Refresh Hosts**.

3.3 Setting Up Infrastructure for Online Patching

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

- Enabling Online Mode

- Setting Up My Oracle Support Credentials and Proxy Connection Settings

- (Optional) Setting Up Online Settings

3.3.1 Enabling Online Mode

To patch the targets in online mode, you must set the connection setting in Enterprise Manager Grid Control to Online mode.

> **Note:**   This is the default mode for patching in Enterprise Manager Grid Control.

To set the connection setting to Online mode, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Setup**. Enterprise Manager Grid Control displays the Overview of Setup page.

2. On the Overview of Setup page, from the vertical menu pane to the left, click **Patching Setup**.

Enterprise Manager Grid Control displays the Patching Setup page with multiple tabs, where My Oracle Support and Proxy Connection tab is selected by default.

3. On the Patching Setup page, select **Online and Offline Settings** tab.

4. On the Online and Offline page, in the Settings section, select **Online**.

5. Click **Apply**.

3.3.2 Setting Up My Oracle Support Credentials and Proxy Connection Settings

In an online patching mode, a *Refresh From My Oracle Support* job is run periodically (either daily or every 24 hours) to automatically download XML files that are required for generating Patch Advisories. Also, the Deployment Procedures automatically download the required patches and patch sets from My Oracle Support every time they are run. For these purposes, Enterprise Manager Grid Control requires credentials to log in to My Oracle Support.

This section provides instructions that help you set up My Oracle Support credentials and proxy server details, if there are any, in Enterprise Manager Grid Control so that it can use these settings to connect to My Oracle Support.

To set up My Oracle Support credentials and proxy connection settings, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Setup**. Enterprise Manager Grid Control displays the Overview of Setup page.

2. On the Overview of Setup page, from the vertical menu pane to the left, click **Patching Setup**. Enterprise Manager Grid Control displays the Patching Setup page with multiple tabs, where My Oracle Support and Proxy Connection tab is selected by default.

3. On the My Oracle Support and Proxy Connection tab, do the following:

   a. In the My Oracle Support section, by default, **Patch Search URL** displays the URL where the patches will be searched. If you want to verify the connectivity to My Oracle Support, then click **Test**.

   If the credentials are valid and if Enterprise Manager Grid Control is able to connect to My Oracle Support successfully, then a status message is displayed to this effect.

   b. In the My Oracle Support Connection Setting section, do one of the following:

   If the host where Enterprise Manager Grid Control is running has a direct connection to Internet, that is, without any proxy server, then select **No Proxy**.

   If the host where Enterprise Manager Grid Control is running has a connection to Internet through a proxy server, then select **Manual Proxy Configuration**. And, depending on the protocol to be used (HTTP or HTTPS), specify the name of the proxy server, port, realm (if credentials are configured for realm), and credentials (if credentials are required for authentication).

   The Connection Configuration subsection shows the default values set for **Timeout** and **Number of Retries**. Timeout indicates the time in milliseconds after which it should automatically time out if the connection takes unusually long time to connect to My Oracle Support. Number of Retries indicates the number of times it should retry connecting to My Oracle Support after timing out. Change the default values to higher values only if required.

   c. In the Agent Connection Setting section, select the settings to be considered for all Management Agent-related communications.

If you want to use the settings specified in the My Oracle Support Connection String, then select **Use My Oracle Support Connection Settings**.

If the host where Management Agent is running has a direct connection to the Internet, that is, without any proxy server, then select **No Proxy**. The Connection Configuration subsection shows the default values set for **Timeout** and **Number of Retries**. Change the default values to higher values only if required.

If the host where Management Agent is running has a connection to Internet through a proxy server, then select **Manual Proxy Configuration**. And, depending on the protocol to be used (HTTP or HTTPS), specify the name of the proxy server, port, sites for which proxy must not be used, realm (if credentials are configured for realm), and credentials (if credentials are required for authentication).

   d. In the Test URL section, to verify whether the Management Agent is reachable, specify a valid Management Agent URL and click **Test**. If the URL is valid and if Enterprise Manager Grid Control is able to reach the Management Agent successfully, then a status message is displayed to this effect.

   e. Click **Apply**.

4. From the top-right corner of the page, click **Preferences**. Enterprise Manager Grid Control displays the General page.

5. On the General page, from the vertical menu pane to the left, click **Preferred Credentials**. Enterprise Manager Grid Control displays the Preferred Credentials page.

6. On the Preferred Credentials page, in the My Oracle Support Preferred Credentials section, click **Set Credentials**.

7. On the My Oracle Support Preferred Credentials page, set the My Oracle Support credentials and click **Apply**.

3.3.3 (Optional) Setting Up Online Settings

In an online patching mode, a *Refresh From My Oracle Support* job is run periodically to automatically download XML files that are required for Patch Advisories. Also, the Deployment Procedures automatically download the required patches and patch sets from My Oracle Support every time they are run. For these purposes, Enterprise Manager Grid Control needs to know where the XML files should be stored and the size limit of patch cache up to which it can store the downloaded patches and patch sets.

This section provides instructions that help you set up these online settings.

To set up online settings, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Setup**. Enterprise Manager Grid Control displays the Overview of Setup page.

1.

2. On the Overview of Setup page, from the vertical menu pane to the left, click **Patching Setup**.

Enterprise Manager Grid Control displays the Patching Setup page with multiple tabs, where My Oracle Support and Proxy Connection tab is selected by default.

3. On the Patching Setup page, select **Online and Offline Settings** tab.

4. On the Online and Offline page, do the following:

   a. In the Settings section, retain the default connection setting, that is, **Online**, and set the patch cache size.

   Patch Cache is a repository where all downloaded patches are stored. By default, the maximum size or capacity of patch cache is set to 700 MB, which means it can store patches only up to the limit of 700 MB. Once the patch cache reaches this size, all older patches will be automatically deleted to free up some space.

   b. In the Metadata Dump Directory section, retain the default dump directory selection, that is, **No**.

   Dump Directory is a directory where My Oracle Support metadata files can be stored. By default, the dump directory is a temporary storage location on the host where Oracle Management Service is running. The metadata files are XML files required for patch advisories and they are downloaded automatically in an online patching mode.

5. Click **Apply**.

3.4 Setting Up Infrastructure for Offline Patching

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

- Enabling Offline Mode
- Enabling Patch Recommendations
- Uploading OPatch Patches to Oracle Software Library
- Uploading Patches to Oracle Software Library

3.4.1 Enabling Offline Mode

To patch the targets in offline mode, you must set the connection setting in Enterprise Manager Grid Control to *Offline* mode.

To set the connection setting to *Offline* mode, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Setup**. Enterprise Manager Grid Control displays the Overview of Setup page.

2. On the Overview of Setup page, from the vertical menu pane to the left, click **Patching Setup**.

   Enterprise Manager Grid Control displays the Patching Setup page with multiple tabs, where My Oracle Support and Proxy Connection tab is selected by default.

3. On the Patching Setup page, select **Online and Offline Settings** tab.

4. On the Online and Offline page, in the Settings section, select **Offline**.

5. Click **Apply**.

3.4.2 Enabling Patch Recommendations

This section explains what you must do to receive patch recommendations automatically. In particular, this section covers the following:

- Identifying Affected Targets and Required Patches

- Uploading Metadata XML Files

- Creating "Refresh From My Oracle Support" Job

3.4.2.1 Identifying Affected Targets and Required Patches

The information about the targets affected by the latest patches and the patch you have to manually download is available in metadata XML files. In an offline mode, you must use another host that has internet connection and manually download these XML files from My Oracle Support.

The following shows the metadata XML files you must download:

Table 3–2   Downloading Metadata XML Files

Metadata XML File	Location To Download From
aru_product.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_products
aru_releases.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_releases
aru_platforms.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_platforms
aru_languages.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_languages
aru_product_groups.xm	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_product_groups
aru_product_releases.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_product_releases
aru_component_releases.xml	https://updates.oracle.com/ARULink/XMLAPI/download_seed_data?table=aru_component_releases
patch_recommendations.xml	https://updates.oracle.com/Orion/Services/search?language_id=0&classification_group=recommended&patch_type=all&group_by=target_type&group_by=release&group_by=platform

3.4.2.2 Uploading Metadata XML Files

To upload the metadata XML files, follow these steps:

1. In Grid Control, from the top-right corner of the page, click **Setup**. Enterprise Manager Grid Control displays the Overview of Setup page.

2. On the Overview of Setup page, from the vertical menu pane to the left, click **Patching Setup**.

   Enterprise Manager Grid Control displays the Patching Setup page with multiple tabs, where My Oracle Support and Proxy Connection tab is selected by default.

3. On the Patching Setup page, select **Online and Offline Settings** tab.

4. On the Online and Offline Settings page, in the Metadata Cache section, upload the metadata XML files to the Management Repository:

5. Click **Apply**.

3.4.2.3 Creating "Refresh From My Oracle Support" Job

To extract information from the metadata XML files and display them in Grid Control, you must schedule a job, that is, a "Refresh From My Oracle Support" job.

> **Note:**
>
> - In offline mode, the job does not actually connect to My Oracle Support. Instead, it looks for the metadata XML files you manually downloaded and stored in the Management Repository to compute the patch recommendations.
>
> - In online mode, the job actually connects to My Oracle Support and automatically downloads the metadata XML files, and then computes the patch recommendations based on the inventory information collected.

To create "Refresh From My Oracle Support" job, follow these steps:

1. In Grid Control, click **Jobs**. Enterprise Manager Grid Control displays the Job Activity page.

2. On the Job Activity page, from the **Create Job** list, select **Refresh From My Oracle Support** and click **Go**. Enterprise Manager Grid Control displays the Create 'Refresh From My Oracle Support' Job page.

3. On the Create 'Refresh From My Oracle Support' Job page, specify a name for this job, schedule it to run at a particular time, and grant access to roles that must have access to this job.

4. Click **Submit**.

3.4.3 Uploading OPatch Patches to Oracle Software Library

From My Oracle Support, manually download patch 6880880 for the required platform and the version of the target you are patching, and upload it to Oracle Software Library (Software Library). This is a platform-specific patch, so you will have to carefully select the platform while downloading this patch.

For example, if you are patching Oracle Database 11g Release 1 (11.1) that is running on Linux x86, then download patch 6880880 for the Linux x86 platform and for the 11.1.0.0.0 release.

To upload patches to the Software Library, follow these steps:

Ensure that you upload it under the product family "System Management Products and Product - Universal Installer" with the appropriate release and platform details.

1. In Grid Control, click **Deployments**. Enterprise Manager Grid Control displays the Deployments page.

2. On the Deployments page, from the **Patching** section, click **View/Upload Patch**. Enterprise Manager Grid Control displays the Patch Cache page.

3. On the Patch Cache page, for **Patch File**, click **Browse** to upload the patch. In the Patch Attributes section, specify details about the patch. Ensure that you upload it under the product family "System Management Products and Product - Universal Installer".

3.4.4 Uploading Patches to Oracle Software Library

For patching targets in offline mode, you must have already stored the patches in the Software Library so that the Deployment Procedure can retrieve them from there and apply them to the affected targets. For this purpose, you must upload the patches to the Software Library.

To upload a patch to the Software Library:

1. In Grid Control, click **Deployments**. Enterprise Manager Grid Control displays the Deployments page.

2. On the Deployments page, from the Patching section, click **View/Upload Patch**.

3. On the View Patches page, click **Upload Patch**.

4. On the Upload Patch page, in the Upload Criteria section, do the following:

   a. From the **Product Family** list, select the product family for which you are uploading the patch. For example, if you want to patch Enterprise Manager Database Control, then select the product family it belongs to, that is, Oracle System Management Products.

   b. From the **Product** list, select the product for which you are uploading the patch. For example, Enterprise Manager Database Control.

   c. From the **Select Type of Patch**, select either **Patch** or **Patch Set** depending on the type of patch you are uploading.

5. In the Upload Method section, select the source from where you are uploading the patches.

6. In the Upload Patch section, click **Browse** and select the patch you want to upload. You can upload multiple patches simultaneously. The other fields in this section will be prefilled based on the information available in the metadata XML files.

7. Click **Upload**.

3.5 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

- Setting Up Linux Patching Group for Compliance Reporting

3.5.1 Prerequisites for Using the Linux Patching Feature

To use the Linux Patching feature, perform these setup activities:

1. Configure the Software Library as described in Section 3.1.2, "Setting Up Oracle Software Library".

2. Deploy the PAR files from the Oracle home of the OMS:

   ```
   $ cd $ORACLE_HOME/sysman/prov/paf
   emctl partool deploy -parDir . -force
   ```

3. The patch user credentials to be used for patching must have *write* access under Oracle home directory of the Management Agent.

4. Operating system credentials used to create groups and set up repository must have SUDO as root privilege.

5. Following commands must be enabled through SUDO:

   - /bin/cp

   - /bin/rm

- /bin/chmod

- /sbin/chkconfig

- yum

- up2date

- sed

- rpm

3.5.2 Setting Up the RPM Repository

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

3.5.2.1 Prerequisites for Setting Up the RPM Repository

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

1. Identify a Redhat or OEL host, install a Management Agent and point to the OMS. This host must have "sudo" package installed.

2. Obtain a valid CSI number.

3. Download the up2date packages corresponding to the host version and release from `https://linux.oracle.com/switch.html`.

4. Upload the up2date packages to the Software Library.

    - First compress up2date and up2date-gname into a zip file and name it as up2date_comp.zip.

    - Upload the zip file to Deployments > Provisioning > Components > Oracle Components > Stage Server Up2date Component > 10.2.0.4.0 > Linux > UP2DATE_RPM. The Software Library by default has packages for Red Hat Enterprise Linux 4, i386 uploaded.

5. Ensure that ULN staging host is able to communicate with the ULN network. If proxy is required, up2date from the host needs to be configured as well. The connectivity with ULN will be detrimental for `up2date -register -nox` command.

6. Patch user (OS credentials used to setup the Staging server) must have write permission under the agent home. Patch user must also have SUDO privilege.

3.5.2.2 Setting Up the RPM Repository

To set up an RPM Repository that downloads latest RPM packages and Advisories from ULN, follow these steps.

1. In Grid Control, from the top-right corner of the page, click **Setup** and select **Patching Setup**.

2. In the **Linux Patching Setup** tab, click the Setup RPM Repository link.

3. In the Setup Linux RPM Repository section, select the RPM Repository server by clicking the search icon. Select the machine assigned for subscribing to ULN.

4. In the Credentials section, specify the user name and password to use. Click **Apply**.

5. In the Deployment Procedure submission confirmation, click the link "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) Set the refresh mode to "30 seconds refresh".

7. Wait till the step "Installing Up2date" is completed or skipped.

8. In the Phase Status page, click the link "Register with ULN"and do the following:

   a. Log in to the RPM Repository server machine.

   b. Configure up2date to use a proxy server, if any, by following the instructions at:

      ```
      https://linux.oracle.com/uln_faq.html - 9
      ```

   c. Register the machine 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

9. After registering the host, select the target and click **Confirm**, and then click **Done** to go to the main flow.

10. Click the link "Subscribe to ULN channels". Do the following:

    a. When you register a server, it will be subscribed to a channel that has the latest Enterprise Linux packages for the appropriate architecture. To subscribe to additional channels, log in to http://linux.oracle.com after you register your system. Click on the **Systems** tab to manage subscriptions for each subscribed server.

    b. Subscribe either to el*_addon channel (this channel contains createrepo) or manually install the createrepo package.

11. Type the command `up2date -nox -show-channels` to verify the list of subscribed channels.

12. Once the deployment procedure completes successfully, go to **Setup**, select **Patching Setup**, and then select **Manage RPM Repository** to verify if the ULN channels are displayed in the Enterprise Manager Console.

13. In the Manage Repository Home page, check if all the subscribed channels are listed and if all the packages are downloaded.

3.5.3 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 Section 3.5.2, "Setting Up the RPM Repository".

In particular, this section describes the following:

- Prerequisites for Setting Up Linux Patching Group
- Setting Up a Linux Patching Group

3.5.3.1 Prerequisites for Setting Up Linux Patching Group

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

- RPM Repository server must be set up or a custom RPM Repository must be set as a channel in Enterprise Manager.

- Yum or up2date should be installed in the target machines.

- Sudo must be installed on the target machines.

- You must have Operator privileges on the hosts that you want to add to the Linux host patching group.

- Patch user must have write access under the agent home. Patch user must have sudo privilege.

3.5.3.2 Setting Up a Linux Patching Group

To set up a Linux patching group, do the following:

1. In Grid Control, from the top-right corner of the page, click **Setup** and select **Patching Setup**.

2. In the **Linux Patching Setup** tab, click the **Setup Groups** link.

3. In the Setup Groups page, click **Create**.

4. In the Create Group: Properties page, specify a unique name for the group. Select the maturity level, Linux distribution, and Linux hosts to be added to the group. Click **Next**.

5. In the Create Group: Package Repositories page, select the RPM Repositories to be associated with the group (click the search icon to select repository).

   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 select package repository.

   Under the "Package Compliance" section, choose whether to include "Rogue" packages in compliance reporting or not.

   Click **Next**.

6. In the Create Group: Credentials page, specify the host credentials or choose to use preferred credentials. Click **Next**.

7. If **Automatically Update Hosts** is not selected, the steps, "Patching Scripts" and "Schedule" pages will be skipped.

8. In the Patching script page, you can specify any pre/post patching operations to be done. This is not a mandatory step.

9. In the Schedule Page, select the schedule for the "Update Job". Click **Next**.

10. Validate all the parameters in the Review page. Click **Finish**.

11. On clicking **Finish**, three jobs will be submitted if you have not selected the **Automatically Update Hosts** option. If the option is selected, then four jobs will be submitted. Table 3–3 explains the jobs submitted. Follow the jobs submitted by clicking the job's link.

12. Go to Deployments > Linux Patching page. Verify the compliance report generated. The group created will have at least one out-of-date package.

Table 3–3   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 yum and up2date tool.
Compliance Collection	Compares the packages already installed in the machine with the packages versions in the selected RPM Repositories and generates Compliance Reports.
Package Information	Collects metadata information from the selected RPM Repositories.
Packages Update	Updates non-compliant packages.

3.6 Setting Up Infrastructure for Virtualization Systems

To start using guest virtual machines, you will need to perform the following setup activities:

- Configuring Virtual Servers
- Creating Server Pools
- Registering Virtual Servers

3.6.1 Configuring Virtual Servers

The virtual server will need to be configured before it can be registered with the server pool to enable Enterprise Manager to perform monitoring, administration, and provisioning.

1. Depending on the Oracle VM version you are using, do one of the following:

    - For Oracle VM version 2.2 or higher, create a storage repository on the master virtual server as follows:

      ```
      # /opt/ovs-agent-2.3/utils/repos.py -n example.com:/test
      ```

    - For Oracle VM version 2.1.2 or 2.1.5, create /OVS mount points for shared storage on the Oracle VM server.

2. No more than 24 Oracle VM servers must be registered and managed with one Oracle Management Agent.

3. Enterprise Manager supports only Oracle VM server version 2.1.2 or above. All Oracle VM servers version 2.1.2 or above managed by Enterprise Manager should be updated to have OVS agent of version 2.2-70 or higher.

4. The following ports must be open:

   ```
   8002,8003,7777,8899,5900-5950,5353,631
   ```

   To verify which ports are open, use the following command:

   ```
   service iptables status
   ```

5. If you are migrating from Oracle VM Manager to Enterprise Manager, see Chapter 12.10, "Migrating from Oracle VM Manager to Enterprise Manager".

6. If you want to enable High Availability mode, see Chapter 12.8.4, "Configuring Oracle VM Servers and Guest Virtual Machines for High Availability".

7. If a user other than root is used to connect to the virtual server, `/etc/sudoers` file must be updated to give "sudo all" privilege to this user. To provide sudo privilege for non-root user, add the following entry:

    ```
    <non-root user> ALL=(ALL) ALL
    ```

    In addition, comment out the following line:

    ```
    Defaults    requiretty
    ```

3.6.2 Creating Server Pools

It is essential that you create a server pool to manage all Oracle VM servers. Server pools can be created with the high-availability option enabled or not. A server pool must have a minimum of one virtual server registered with it.

To create a Virtual Server Pool, do the following:

1. In Grid Control, click **Targets**, then click **Virtual Servers**.

2. In the Virtualization Central page, select Action as Virtual Server Pool Create Virtual Server Pool and click **Go**.

3. In the Create Virtual Server Pool page, provide the values specified in Table 3–4. Click **OK** to submit the server pool creation job.

Virtualization Central >
Create Virtual Server Pool
Provide details of the virtual server pool and the master virtual server you want to register with Enterprise Manager.

* Virtual Server Pool Name   My_virtual_server_pool
Shared Directory   /OVS
High Availability Mode   ☑ Enable

Master Server Virtual IP   198.68.0.1
Virtual IP to connect to the master virtual server and support the master server failover. This input is required only for Oracle VM ver:

Master Virtual Server Details
* Virtual Server Host   vs_host
Name of master Virtual Server
* Monitoring Server Agent   stabd70.us.oracle.com:3872
EM Agent that will monitor the Virtual Server
* Monitoring Server User   ms_user
Operating system user that owns the EM Agent home
* Monitoring Server Password   ••••
Credentials for operating system user that owns the EM Agent home
* SSH Username   ssh_user
SSH Username to connect to Virtual Server
* SSH Password   ••••••
SSH Password to connect to Virtual Server
* Oracle VM Agent Password   ••••••
OVM Agent Password on Virtual Server
* Staging Directory   \tmp
Directory location on Virtual Server that can be used by EM to stage management scripts

Navigate to Virtualization Central and verify that the status of the server pool you created is `Registering..` After the job is completed, the Server Pool must appear in Virtualization Central.

After the job is completed, the server pool must appear in Virtualization Central and status will be Up.

Table 3–4   Virtual Server Pool Configuration Details

Element	Description
Virtual Server Pool Name	Specifythe name of the virtual server pool that you are creating.

Table 3–4   (Cont.) Virtual Server Pool Configuration Details

Element	Description
Shared Directory	Specifythe directory mounted as the shared directory. This shared directory is the same for all virtual servers within a virtual server pool. For example, `/OVS`. All guest virtual machine templates and running guest virtual machines are maintained in this location.
High Availability Mode	Select **Enable** if you want the virtual server to participate in high availability. For more information about high availability, see Chapter 12.8, "High Availability and Fault Tolerance".
Master Server Virtual IP	If you have chosen to enable high availability mode, optionally specify an unused IP address to support master server failover, and to be used as virtual IP for master server communication.
Virtual Server Host	Specify either FQDN name or IP address of the master virtual server for the server pool.
Monitoring Server Agent	Specify the Management Agent that will monitor the virtual server.
Monitoring Server User	Specify the operating system user that owns the Management Agent home/Management Agent install, and has read and write permission on the Management Agent home directory.
Monitoring Server Password	Specify the password credentials for the monitoring server user.
SSH Username	Specify the user on the virtual server who has sudo privileges.
SSH Password	Specify the password credentials for the SSH user.
Oracle VM Agent Password	Specify the password of the Oracle VM Agent running on the virtual server.
OVS Proxy Location	Specify a directory on the virtual server where the scripts required for monitoring and administration are stored when the virtual server is registered with Enterprise Manager Grid Control.

3.6.3 Registering Virtual Servers

For Enterprise Manager Grid Control to monitor and manage a virtual server, you must register it with a server pool. Once you create a master virtual server, you can also register other virtual servers to the server pool.

To register virtual servers, follow these steps:

1. In the Virtualization Central page, select a Server Pool, select Action as Virtual Server - Register Virtual Server and click **Go**.

2. In the Register Virtual Server page, specify details of the virtual servers to be registered as shown in Table 3–5. Click **Add Another Row** to specify multiple virtual servers details.

    You can batch assign details to different virtual servers and click **Assign** to specify common details for multiple virtual servers.

You can add a virtual server from file by clicking **Add From File**. Select the file from a local machine or a host machine in Enterprise Manager. The file should contain details of the virtual servers in the following format:

```
<virtualServerName>;<Monitoring Server Name>;<Monitoring Server User>;<SSH
User>;<Proxy Location>;<ServerPoolName>
```

For example:

```
vs1.xyz.com;abc.xyz.com:4889;oracle;orauser;/tmp;ServerPool1.xyz.com
```


3. Click **OK**.

After registering the virtual server, navigate to Virtualization Central and verify that that the virtual server appears under the specific server pool.

Table 3–5   Register Virtual Server Details

Element	Description
Virtual Server Host	Specify the name of the master virtual server for the server pool.
Monitoring Server Agent	Specify the monitoring agent.
Monitoring Server User	Specify the Operating system user that owns the Management Agent home/Management Agent install, and has read and write permission on agent home directory.
Monitoring Server Password	Specify the password credentials for the monitoring server user.
SSH Username	Specify the user on the virtual server who has sudo privileges.
SSH Password	Specify the password credentials for the SSH user.
Oracle VM Agent Password	Specify the password of the Oracle VM Agent running on the virtual server.
OVS Proxy Location	Specify a directory on the virtual server where the scripts required for monitoring and administration are stored when the virtual server is registered with Enterprise Manager.

Table 3–5 (Cont.) Register Virtual Server Details

Element	Description
Virtual Server Pool	Specify the virtual server pool to which the virtual server will belong.

3.7 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:

- Prerequisites for Using Bare Metal Provisioning
- Creating Super Administrator for Enterprise Manager
- Setting Up Boot Server
- Setting Up Stage Server
- Setting Up RPM Repository
- Setting Up Software Library
- Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host
- Setting Up Preferred Credentials for Bare Metal Provisioning in Enterprise Manager
- Accessing the Bare Metal Provisioning Application
- Configuring Stage Server
- Configuring RPM Repository
- Configuring Boot Server
- Configuring the DHCP Server
- Configuring Software Library

3.7.1 Prerequisites for Using Bare Metal Provisioning

The following steps should be followed for evaluating or using the Bare Metal Provisioning application:

Install or upgrade to the latest Enterprise Manager 11g Grid Control Release 1 (11.1.0.1).

- If you are using Enterprise Manager 10g Grid Control version 10.2.0.4 or earlier, then first upgrade it to version 11.1.0.1. Refer to the Enterprise Manager Grid Control 11g Release 1(11.1.0.1) notes available at
  `http://www.oracle.com/technology/documentation/oem.html`

Download the agent rpm kit from the following location:

`http://www.oracle.com/technology/software/products/oem/htdocs/provisioning_agent.html`

3.7.2 Creating Super Administrator for Enterprise Manager

Only a user who is a Super Administrator for Enterprise Manager can configure various elements like stage server, boot server etc. for use with the Provisioning Application. Not only that, it is only these users who can actually create assignments

for actually provisioning target machines with any image. For more information about assignments, see Section 2.5.5, "Assignments".

Follow the steps below to create a user who is a super administrator.

1. Log into the Enterprise Manager and click on the **Setup** link on top right hand corner as shown in the picture below.

2. On the Setup page, click on the **Administrators** link on the left hand side column.

3. On the **Administrators** page, click **Create** as shown below.

4. On the Create step that comes up fill up the necessary details as shown and select the **Super Administrator** check box as shown below. Click **Next**.

5. On the **Review** page, click **Finish** to complete the user creation.

3.7.3 Setting Up Boot Server

This section explains how to set up the boot server. Following are the prerequisites:

■   2 GB RAM for boot server, stage server, and RPM repository server.

■   Boot server and stage server must be on the same physical machine.

■   If boot server and stage server reside on different machines, then the boot install directory (/tftpboot/linux-install) should be mounted on the stage server.

If you have the required boot server, stage server, and RPM repository already created, then set up the preferred credentials. For more information on setting up preferred credentials, see Section 3.7.8, "Setting Up Preferred Credentials for Bare Metal Provisioning in Enterprise Manager".

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 (**/tftpboot/linux-install** in the given examples).

3. Configure DHCP Server:

   Edit the **dhcpd.conf** (*/etc/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 hostname 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:

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

6. Install Oracle Management Agent.

3.7.4 Setting Up Stage Server

Stage server must meet the following requirements:

- **Large Storage**

  The files related to components and directives of an image are first copied to the stage server in preparation for the network installation, and are kept there for future use. The stage server thus acts as a huge cache of files, which requires a large storage.

  The stage server can also host the staging storage on **Network Attached Storage(NAS)**. Multiple stage servers can use the same **NAS**.

- **High Memory**

  The stage directives associated with the components and images are directives that are executed during staging phase of a component or Image. They contain commands to unpack and layout the files in order to facilitate the network installation. Depending on the size of the components and images, these commands place high memory requirements on part of the stage server.

- **Sufficient Bandwidth**

  Staging process could be very time consuming if the network between the Stage server and software library (on Oracle Management Service or OMS Server) does not have sufficient bandwidth to enable fast transfer of files. Similarly, the link between the stage server and hardware servers should have high bandwidth to make the installation process faster.

- **NFS or HTTP Support**

  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 uses NAS for staging storage, the NAS server should have the NFS support.

  If the stage server cannot have NFS support, it must be accessible by HTTP.

Follow the instructions listed below to set up a Linux machine as the stage server:

1. Create a top-level directory.

   - Create a top-level directory on the stage server where all the files will be stored. In the following steps, **STAGE_TOP_LEVEL_DIRECTORY** refers to the absolute path of this top-level directory.

   - You can specify the HTTP or NFS location of agent rpms when you set up a default image in the Provisioning application. Alternatively, you can copy agent rpms to STAGE_TOP_LEVEL_DIRECTORY so that they are picked up automatically. If 11.1 agent rpms are available in STAGE_TOP_LEVEL_DIRECTORY, they will be picked up. Otherwise, depending on the agent rpm versions available in STAGE_TOP_LEVEL_DIRECTORY, the 10.2.0.5 or earlier versions of agent rpms will be picked up. The 32-bit and 64-bit agent rpms

will be picked up depending on whether the Linux operating system being provisioning is 32-bit or 64-bit.

To copy agent RPMs to STAGE_TOP_LEVEL_DIRECTORY, do the following:

a. Download Agent RPMs zip file from the following location:

```
http://www.oracle.com/technology/software/products/oem/htdocs/provision
ing_agent.html
```

b. Unzip the agent RPMs zip file into a temporary directory and then copy the agent RPMs from the emagent directory to STAGE_TOP_LEVEL_ DIRECTORY.

2. Configure NFS services.

Perform the following steps on the stage server.

> **Note:** If the stage server uses **NAS** for staging storage, the following steps need to be performed on the **NAS server** as well.

1. Ensure the NFS service is running. One can check this by running **service nfs status**.

Modify the "/etc/exports" file to have the following entry:

```
{Directory path} {host_name_or_ip_prefix}* (ro,sync)
```

For example, /STAGE_TOP_LEVEL_DIRECTORY 10.152.* (ro,sync), if the hardware servers to be provisioned have the IP prefix 10.152.

Or, /STAGE_TOP_LEVEL_DIRECTORY provision-host* (ro,sync), if the hardware servers to be provisioned have names starting with provision-host.

2. After the modification is made, run the **service nfs restart** command to make the changes visible to nfs daemons.

3. Install Enterprise Manager Agent.

Refer to the "Installing Management Agent" chapter in *Oracle Enterprise Manager Grid Control Installation Guide* to install a 10.2.0.3 or higher version of Management agent on the Stage Server.

> **Note:** Ensure that the preferred credentials set for the staging server host has "write" access to the staging storage.

4. From Grid Control Console, set the privileged preferred credentials for the stage server. For information about setting preferred credentials in Enterprise Manager, see Section 3.7.8, "Setting Up Preferred Credentials for Bare Metal Provisioning in Enterprise Manager".

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 Enterprise Manager. This preferred credential should also be a valid ORACLE_HOME credential (belonging to ORACLE_HOME owner's group).

3.7.5 Setting Up RPM Repository

> **Note:** It is recommended that you use RAM of 2 GB.

3.7.5.1 Setting UP RHEL 4 RPM Repository

RPM Repository is used as the source of Linux and application packages that need to be installed on the newly provisioned bare metal box. For example, an RPM Repository may be created to contain all the 32-bit Linux rpms and another repository may be created to contain Linux x86-64 bit rpms. Two separate Linux images can then be created each based on one of the repositories.

RHEL RPM repository to be used should have the following Red Hat Install tree structure:

There are multiple ways to create a RPM repository. If Red Hat Enterprise Linux CDs are available, do the following:

1. Copy all the contents of the first CD to a directory say RPM_REPOS.

2. Copy all rpms from other CDs to <RPM_REPOS>/Redhat/RPMS. Change directory to the RPMS directory:

   ```
   cd <RPM_REPOS>/Redhat/RPMS
   ```

3. Add custom RPMs to the repository.

   a. If there are custom RPMs installed on the reference host that need to be provisioned on the bare metal machine, make sure to copy them to the following repository location:

      ```
      <RPM_REPOS>/Redhat/RPMS
      ```

   b. Install anaconda-runtime RPM on the machine hosting the RPM repository. This might require other dependent packages to be installed.

   c. Run the following commands:

      ```
      cd /usr/lib/anaconda-runtime
      ./genhdlist --productpath=RedHat --withnumbers --hdlist <RPM_
      REPOS>/RedHat/base/hdlist <RPM_REPOS>
      ```

4. Run `yum-arch` :

   This should create a **headers** directory. Make sure this directory contains a **header.info** file.

   If yum is not installed then download it from the Linux Vendor's Web site.

5. Create a symbolic link in `/var/www/html` to <RPM_REPOS> directory.

The repository should now be available through http if an apache server is running.

3.7.5.2 Setting Up OEL 4 RPM Repository

Oracle Enterprise Linux (OEL) RPM repository should have the Install tree structure shown below:

You can set up Oracle Enterprise Linux (OEL) Repository by using the OEL installation media as follows:

1. Download Oracle Enterprise Linux from http://edelivery.oracle.com/linux.

2. Copy all the contents of the first CD to a directory say **RPM_REPOS**.

3. Copy all rpms from other CDs to <RPM_REPOS>/Enterprise/RPMS. Change directory to the RPMS directory:

   ```
   cd <RPM_REPOS>/Enterprise /RPMS
   ```

4. Add custom RPMs to the repository.

   a. If there are custom RPMs installed on the reference host that need to be provisioned on the bare metal machine, make sure to copy them to the following repository location:

      ```
      <RPM_REPOS>/Enterprise/RPMS
      ```

   b. Install anaconda-runtime RPM on the machine hosting the RPM repository. This might require other dependent packages to be installed.

   c. Run the following commands:

      ```
      cd /usr/lib/anaconda-runtime
      ./genhdlist --productpath=Enterprise --withnumbers --hdlist <RPM_
      REPOS>/Enterprise/base/hdlist <RPM_REPOS>
      ```

5. Run `yum-arch` :

   This should create a **headers** directory. Make sure this directory contains a **header.info** file.

6. Create a symbolic link in /var/www/html to <RPM_REPOS> directory.

The repository should now be available through http if an apache server is running.

3.7.5.3 Setting Up RHEL 5/OEL 5 RPM Repository

Oracle Enterprise Linux (OEL) RPM repository should have the Install tree structure shown below:

You can set up Oracle Enterprise Linux (OEL) Repository by using the OEL installation media as follows:
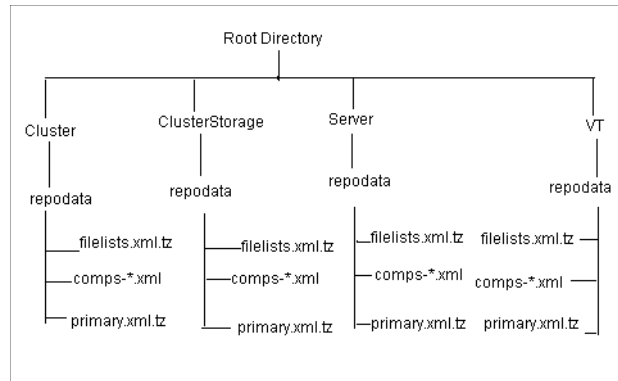
1. Download Oracle Enterprise Linux from http://edelivery.oracle.com/linux.

2. Copy all the contents of the first CD to a directory say `Root Directory`.

3. Copy all contents from the Cluster, ClusterStorage, Server, and VT directories in the other CD to the respective directories.

   For setting up RPM repository for OVS provisioning, copy RPMs from Oracle, Server, and VT directories to the respective directories.

4. Run createrepo for all four directories. For example:

   ```
   createrepo <Root Directory>/cluster
   ```

5. Add custom RPMs to the repository as follows:

   a. If there are custom RPMs 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, such as Cluster, VT, ClusterStorage, and Server.

   b. Run the createrepo command on this directory. For example:

   ```
   createrepo ClusterStorage
   ```

6. Create a symbolic link in `/var/www/html` to `<Root Directory>` directory.

The repository should now be available through http if an apache server is running.

3.7.6 Setting Up Software Library

Software Library should be located in a directory accessible by all OMSes. If there is one OMS the directory can be local. For multiple OMS environments, the directory can be on a Network File Server or a Netapp filer that is accessible from all the OMSes. One has to ensure that there is enough space available on the shared storage to store files that hold the binary data for one's components.

Software components that are generated as part of the default or single-server image creation during the bare metal provisioning process are stored in the Software Library. They are accessible under the **Components** tab in the Provisioning Application user interface.

Ensure that the shared storage is accessible through NFS mount points to all OMS servers in the environment.

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

Ensure that the following criteria are met before provisioning:

Table 3–6   Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host

Resource Name	Checklist
Boot Server	DHCP server is up and running.
	The `next_server` entry in `/etc/dhcpd.conf` file points to this boot server.
	TFTP is up and running.
	Boot Server is co located on the same machine as Stage Server. If not then Network Install Directory (/tftpboot/linux-install) is exposed to the Staging server for mounting.
	Boot Server is present in the same subnet where the target machines to be provisioned are present or will be added.
	Enterprise Manager agent is installed.
	Boot server machine is visible as a managed target in Enterprise Manager.
	A brand new PXE box actually detects the boot server and starts to boot it (even if no image is installed yet)
Stage Server	Stage server is as close as possible to the target servers.
	Large storage, High Memory and Sufficient Memory.
	If NAS server is used for storage then it should have NFS support.
	Enterprise Manager agent is installed.
	Boot server machine is visible as a managed target in Enterprise Manager.
	The required agent rpm is staged for installing agents on targets.
	Preferred Credentials are set.
	Stage server is reachable from the box to be populated (or the same subnet)
RPM Repository	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	Agent is installed on local disk and not on NFS mounted directory.
	Preferred Credentials are set.
Software Library	Shared storage used for the software library is accessible through NFS mount points to all OMS servers.

3.7.8 Setting Up Preferred Credentials for Bare Metal Provisioning in Enterprise Manager

Preferred credentials simplify access to managed targets by storing target login credentials in the Management Repository. With preferred credentials set, users can access an Enterprise Manager target that recognizes those credentials without being prompted to log into the target. Preferred credentials are set on a per user basis, thus ensuring the security of the managed enterprise environment.

Enterprise Manager supports two types of preferred credentials:

- **Normal Credentials:** Are used by Enterprise Manager functions that need operating system access, but do not require administrator privileges.

- **Privileged Credentials:** Are used by functions that need administrator privileges. Credentials for users that have sudo access on the target machine can be used as privileged credentials.

The Provisioning application requires preferred credentials to be setup for machines, which are part of the application.The preferred credentials need to be set for the following machines:

- **Referenced Installation Host:** Privileged credentials are needed to execute the command to get all the available RPMs form this machine. The credentials should also be valid ORACLE_HOME credentials (belonging to ORACLE_HOME owner's group)

- **Stage Server:** You must set the privileged preferred credentials for the stage server. Oracle recommends the stage server to 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 Enterprise Manager. This preferred credential should also be a valid ORACLE_HOME credential (belonging to ORACLE_HOME owner's group).

- **Provisioning Targets:** In case you are planning to provision existing target machines, ensure the privileged credentials are setup. These credentials are required to clear the boot-sector and reboot the machine.

For details of how to set preferred credentials in Enterprise Manager, see Section 3.2.3, "Setting Up Preferred Credentials".

> **Caution:**   It is recommended that you do not use root as the preferred credentials.

3.7.9 Accessing the Bare Metal Provisioning Application

To access the Bare Metal Provisioning Application, do the following:

1. Log in to Oracle Enterprise Manager.

2. The Provisioning application can be accessed by going to the Deployments Tab and then to Provisioning sub tab, as shown below.

The graphical user interface of the provisioning application shows various tabs for Components, Directives, and Images etc. A user can access all or some tabs shown above depending upon the privileges assigned to him. For example, in the figure below, the Administration and Assignments tabs are disabled for the user. Refer to Section 3.7.2, "Creating Super Administrator for Enterprise Manager" for creating users that can access the Administration tab.

In this section, we will assume that the user has super user privileges and can thus access the administration tab. This tab contains different sections for configuring different elements in the environments.

3.7.10 Configuring Stage Server

In this section, it is assumed that the stage server has been created and the necessary setup has been done.

Click **Add Server** in the **Staging Server Configuration** section. The following page comes up.

In the above step, enter the following details:

Staging Server Hostname refers to the host name of the stage server. It is recommended that you specify the IP address of the stage server, to avoid DNS resolution problems.

Full Directory Path refers to the top-level directory on the stage server which contains the agent rpm. All the staged files will be stored in this location.

Base URL is used to expose the top-level stage directory to hardware servers being provisioned via NFS.

NOTE: If Stage Server uses **NAS** then the **Base URL** should be of the form:

```
<file>://<NAS_Hostname>/< NAS_DIRECTORY>
```

Provide the necessary values for all the parameters and click **OK** to add the stage server.

3.7.11 Configuring RPM Repository

To configure the RPM repository, the RPM repository must have been created.

Click **Add** in the RPM Repository Configuration section. The following page is displayed:

Repository Name Assign a name to the RPM repository that is going to be added.

Complete URL refers to the URL pointing to the directory where the required RPMs are located. For example,
http://sample.oracle.com/yum/EnterpriseLinux/EL4/oracle/i386

Provide the necessary values for all the parameters shown in the picture above and click **OK** to add the RPM repository.

3.7.12 Configuring Boot Server

To configure the Boot Server, the Boot Server must have been created.

Click **Add** in the Boot Server Configuration section. The following page is displayed:

Boot Server Hostname refers to the host on which the boot server is setup.

Network Install Directory refers to the directory on the boot server where the pxelinux.cfg file is located. It is generally **/tftpboot/linux-install** directory on linux systems.

Provide the necessary values for all the parameters shown in the picture above and click **OK** to add the Boot Server.

3.7.13 Configuring the DHCP Server

To configure the DHCP server, ensure that the DHCP server is set up as mentioned in Section 3.7.3, "Setting Up Boot Server".

To configure the DHCP server, do the following:

1. In the Administration tab, in the DHCP Server Configuration section, click **Add**.

2. In the Add DHCP Server page, select the host on which the DHCP server resides and select the DHCP configuration path as `/etc/dhcpd.conf`. Click **OK**.

The DHCP Server is configured and will appear in the Administration page.

3.7.14 Configuring Software Library

To set up and configure the Software Library, see Section 3.1.2, "Setting Up Oracle Software Library".

Part II

Database Provisioning

This part describes the Deployment Procedures you can run for provisioning database targets. In particular, this part describes the Deployment Procedures for:

> **Note:** The Cloning Wizard is no longer supported in Enterprise Manager 11g Grid Control Release 1 (11.1.0.1.0).

4

Simulating Database Provisioning Deployment Procedures to Check Prerequisites

Enterprise Manager Grid Control offers database provisioning Deployment Procedures that help you provision database targets such as Oracle Database, Oracle Grid Infrastructure for standalone servers, Oracle Grid Infrastructure and Oracle RAC for clustered environments, and so on. However, before you run the database provisioning Deployment Procedures, Oracle strongly recommends you to simulate the Deployment Procedures so that you check the prerequisites and ensure that you meet all the necessary requirements for database provisioning, and also resolve all operational conflicts and issues beforehand.

Table 4–1 lists the database provisioning Deployment Procedures that support the simulation mode, and describes how you must run them in that mode.

Table 4–1   Simulating Database Provisioning Deployment Procedures for Checking Prerequisites

Deployment Procedure Name	Support or Not Supported	Instructions to Simulate
Oracle Database Provisioning	Yes	Run the Deployment Procedure as described in Chapter 5, "Provisioning Oracle Database", but on the Review page (final page of the wizard), click **Analyze Prerequisites**.
Oracle Grid Infrastructure Provisioning for Standalone Servers	Yes	Run the Deployment Procedure as described in Chapter 6, "Provisioning Oracle Grid Infrastructure for Standalone Servers", but on the Review page (final page of the wizard), click **Analyze Prerequisites**.
Oracle Grid Infrastructure / RAC Provisioning	Yes	Run the Deployment Procedure as described in Chapter 7, "Provisioning Oracle Grid Infrastructure and Oracle RAC for Clustered Environments", but on the Review page (final page of the wizard), click **Run Prerequisites Only**.
One Click Extend Cluster Database	Yes	Run the Deployment Procedure as described in Chapter 8, "Extending Oracle Real Application Clusters", but on the Extend Real Application Clusters page, click **Run Prerequisites Only**.
Delete/Scale down Oracle Real Application Clusters	No	-

Table 4–1   (Cont.) Simulating Database Provisioning Deployment Procedures for Checking Prerequisites

Deployment Procedure Name	Support or Not Supported	Instructions to Simulate
Oracle Database Replay Client Provisioning	No	-

5

Provisioning Oracle Database

This chapter explains how you can mass-deploy Oracle Database in an unattended, repeatable, and reliable manner. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedure
- Supported Releases
- Provisioning Oracle Database Using Existing Installation
- Provisioning Oracle Database Using Gold Image
- Provisioning Standalone Oracle Database Using Installation Binaries

> **Note:** Before you run the Deployment Procedure as outlined in this chapter to provision an Oracle Database, Oracle strongly recommends you to check the prerequisites to ensure that you meet all the necessary requirements for provisioning a database. To do so, simulate the Deployment Procedure as described in Chapter 4, "Simulating Database Provisioning Deployment Procedures to Check Prerequisites", and resolve all operational conflicts and issues beforehand.

5.1 Getting Started

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

Table 5–1 Getting Started with Provisioning Oracle Database

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Database.	To learn about the Deployment Procedure, see Section 5.2, "Deployment Procedure".

Table 5–1   (Cont.) Getting Started with Provisioning Oracle Database

Step	Description	Reference Links
Step 2	**Knowing About the Supported Releases** Know what releases of Oracle Database can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 5.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle Database. Select the use case that best matches your requirement.	■ To learn about cloning an existing Oracle Database, see Section 5.4, "Provisioning Oracle Database Using Existing Installation". ■ To learn about provisioning Oracle Database using a gold image, see Section 5.5, "Provisioning Oracle Database Using Gold Image". ■ To learn about provisioning a standalone Oracle Database using its software binaries from an installation medium, see Section 5.6, "Provisioning Standalone Oracle Database Using Installation Binaries".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for cloning an existing Oracle Database, see Section 5.4.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Database using a gold image, seeSection 5.5.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Database using its software binaries from an installation medium, see Section 5.6.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Database.	■ To clone an existing Oracle Database, follow the steps explained in Section 5.4.2, "Provisioning Procedure". ■ To provision Oracle Database using a gold image, follow the steps explained in Section 5.5.2, "Provisioning Procedure". ■ To provision Oracle Database using its software binaries from an installation medium, follow the steps explained in Section 5.6.2, "Provisioning Procedure".

5.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for provisioning Oracle Database:

■ *Oracle Database Provisioning*

Using this Deployment Procedure, you can perform the following operations across releases and platforms:

- Deploy Oracle Database software using certified gold images, live reference installations (existing database installations), and installation media

- Optionally create multiple databases, one on each host, in parallel

- Use staging location for fast file transfer

For example, if you want to have a fresh installation of Oracle Database, the recommended option is to archive the software binaries of the database from an installation medium, upload the archived binaries to Oracle Software Library (Software Library), and then provision the database using the archived binaries from the Software Library.

> **IMPORTANT:** This Deployment Procedure creates Oracle homes and databases, but does not clone data files. Therefore, when you use this Deployment Procedure to clone an existing, running instance of Oracle Database, you essentially create a new database that has only the structure and format copied from the existing one, but not the data.
>
> If you want to clone data files and maintain a copy that is identical to the existing, running instance, then use the Clone Database Wizard. To access this wizard, on the Deployments pages, from the Cloning section, select **Clone Database**.

5.3 Supported Releases

Using these Deployment Procedures, you can provision the following releases of Oracle Database across multiple hosts:

- Oracle Database 10g Release 1 (10.1.x.x)

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

- Oracle Database 11g Release 2 (11.2.x.x)

> **Note:** To view an updated list of supported database releases, refer to *My Oracle Support* note 412431.1.

> **Note:** For provisioning Oracle Databases 10g Release 1 (10.1.x.x) or lower, use the Clone Oracle Home Wizard as well as the Deployment Procedures in Enterprise Manager Grid Control. To access this Clone Oracle Home Wizard, in Grid Control, click **Deployments**. In the Cloning section, click **Clone Oracle Home**.

5.4 Provisioning Oracle Database Using Existing Installation

This section describes how you can provision Oracle Database using an existing standalone Oracle Database installation that is on a host monitored by Enterprise Manager Grid Control.

In particular, this section covers the following:

- Prerequisites

■ Provisioning Procedure

5.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

■ Ensure that you meet the prerequisites described in the following sections:

– Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

– Section 3.2, "Optional Infrastructure Requirements"

■ Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different (for example, some target hosts might have the software of Oracle Grid Infrastructure and some might not), then contact your system administrator and fix the inconsistencies before running the Deployment Procedure. Either all or none of the target hosts must have the software of Oracle Grid Infrastructure.

To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

■ If you want to use standard response files while provisioning the database, then edit the dbca.rsp file available on the installation medium, and store it as a generic component in the Software Library. Alternatively, you can also maintain this file on any host.

To understand how response files work, see the *Oracle Database Installation Guide for Linux* available at one of these URLs:

For Oracle Database 11g Release 2

http://www.oracle.com/pls/db112/homepage

For Oracle Database 11g Release 1

http://www.oracle.com/pls/db111/homepage

For Oracle Database 10g Release 2

http://www.oracle.com/pls/db102/homepage

http://www.oracle.com/pls/db10g/portal.portal_demo3

To understand how you can upload the response file to the Software Library, see Section 5.5.1.1, "Uploading Database Response File to Oracle Software Library".

■ If you want to use a custom template to create a structure for the database, then create a template (a .dbt file), and store it in a location accessible from the target hosts. The file may be on the target host or on a shared location.

To understand how a template can be created and used for creating databases, see Appendix D, "Creating Templates for Oracle Database".

To understand how you can upload the template to the Software Library, see Section 5.5.1.2, "Uploading Database Template File to Oracle Software Library".

■ Ensure that the operating system user *oracle* is available on all the nodes of the cluster.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- If you want to stage the database software binaries to a shared stage location, that is, while provisioning from an existing installation, then ensure that you use an operating system user that has *write* permission on the staging location. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

  Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

- By default, the Deployment Procedure displays the operating system groups for Database Administrator (OSDBA) and Database Operator (OSOPER). Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in *Oracle Database Quick Installation Guide* available at:

  http://www.oracle.com/pls/db112/homepage

  The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

  For more information, see *Oracle Database 2 Day DBA Guide* available at:

  http://www.oracle.com/pls/db112/homepage

5.4.2 Provisioning Procedure

To provision a database using an existing installation, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Provisioning**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source and Destination page of the Oracle Database Provisioning Deployment Procedure.

4. On the Select Source and Destination page, do the following:

a. In the Select Source section, select **Existing Database Installation**. Then click the torch icon for **Source Host** and select the host on which the existing Oracle Database is running.

In the Source Host Details section, by default, **Oracle Home**, **Working Directory**, and **Files to exclude** are prefilled.

Oracle Home shows where the existing instance is installed, but this is a non-editable field. For **Working Directory**, specify the full path to a directory on source host where the files related to cloning can be staged temporarily. For **Files to exclude**, specify file names that must not be cloned to the source host. Use a comma to separate the file name, and use the wildcard (*) to indicate all files with the same extension. For example, *.trc. Note that any file or folder corresponding to the regular expressions provided here will be excluded from the clone archive that is created

In the Source Host Credentials section, select **Use Preferred Credentials** to use the credentials stored in the Management Repository.

b. In the Specify Destination Host Settings section, click **Add** and select the target host on which you want to provision the existing instance of Oracle Database.

Note: On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

Note: You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

Note: If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced section, specify any additional parameters you want to run while installing the single instance Oracle Database. For example, -debug.

▼ Advanced

Additional Parameters   -debug

Eg: -debug

You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. Using these parameters, you can even change the installation type of the database. For example, `INSTALL_TYPE=SE` (to install standard edition of database), `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

While provisioning software binaries from an existing Oracle Database installation, if you want to also stage them to a shared location, then select **Stage to Shared Location** and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. This stage location can also act as a source for future deployments. Ensure that this stage location has adequate space for storing the software binaries. For example, if the software binaries you are provisioning amount to 1 GB, then maintain at least 2 GB of space in the stage location.

☑ Stage to Shared Location

Shared Location   /dev

The location should be shared across all destibation hosts.

Note: The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select **Staging Location** as the source when you run the Deployment Procedure next time.

---

d. In the Group Details section, specify the Database Administrator (OSDBA) group and the Database Operator (OSOPER) group you want to use for operating system authentication. Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at:

http://www.oracle.com/pls/db112/homepage

The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

For more information, see Oracle Database 2 Day DBA Guide available at:

http://www.oracle.com/pls/db112/homepage

e. (Optional) In the Create Database section, select **Create and Configure Database** if you want to create a new database and configure it after installing the standalone Oracle Database.

> **IMPORTANT:** If you do not select **Create and Configure Database**, only the software binaries of Oracle Database are copied to the target host. You will have to manually configure the provisioned Oracle Database on each target host using Database Configuration Assistant (DBCA).

f. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

There may be occasions when the Management Agent on that target host is down and as a result, OMS is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when OMS should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select **End After** and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

g. Click **Next**.

Enterprise Manager Grid Control displays either the Database Configuration page or the Review page depending on the selection made in Step 4 (e).

In Step 4(e), if you had selected **Create and Configure Database**, then follow the instructions outlined in Step 5. Otherwise, go to Step 6.

5. On the Database Configuration page, do the following:

a. In the Database Identifiers section, specify basic configuration details such as the SID, the global database name, and the database credentials.

> **Note:**
>
> - SID must be unique for a database on a host. This means, the SID assigned to one database on a host cannot be reused on another database on the same host, but can be reused on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their SIDs need to be unique. However, if you install the third database on another host (host2), then its SID can be db1 or db2.
>
> - Global database name must be unique for a database on a host and also unique for databases across different hosts. This means, the global database name assigned to one database on a host can neither be reused on another database on the same host nor on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their global database names need to be unique. And if you install the third database on another host (host2), the global database name of even this database must be unique and different from all other names registered with Enterprise Manager Grid Control.
>
> - The database credentials you specify here will be used on all the destination hosts. However, after provisioning, if you want to change the password for any database, then you must change it manually.

b. In the Database Storage Options section, select the location to store your data files.

f you want to use a file system, then select **File System** and specify the full path to the location where the data file is present. For example, %ORACLE_BASE%/oradata /u01/product/db/oradata

If you want to use ASM, then select **Automatic Storage Management**, and click the torch icon to select the disk group name and specify ASMSNMP password. The Disk Group Name List window appears and displays the disk groups that are common on all the destination hosts.

For example, if the destination hosts are A and B, and if host A has three disk groups, mainly DATA1, DATA2, DATA3, and if host B has four disk groups, mainly DATA1, DATA3, DATA4, DATA5, then the Disk Group Name List windows shows DATA1 and DATA3 because these two are common across the two hosts.

> **Note:** This option will be enabled only if all the destination hosts have Oracle ASM 11g Release 2 (11.2) or higher configured.

c. (Optional) In the Advanced Database Parameters section, specify any additional DBCA parameters that you want to pass while configuring the database. For example, `-characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

d. (Optional) In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you want to use a custom response file that has options enabled. If you have stored the response file as a generic component in

the Software Library, then select **Software Library** and click the torch icon for Component to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

> **Note:** While creating the database, the values specified in the page for mandatory parameters, such as SID and so on, will take precedence over the values specified for the same parameters in the response file.

Oracle Database Response File

☑ Use Custom Database Response File

○ Software Library   ● Destination Host

Response File Location   /initrd   🖊

e. In the Oracle Database Template File subsection, select **Use Custom Database Template File** and use the template file that is being used in the existing instance of Oracle Database that you are cloning. If you have stored the template file as a generic component in the Software Library, then select **Software Library** and click the torch icon for **Component** to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

> **IMPORTANT:** If you do not select Use Custom Database Template File, the provisioned Oracle Database will be configured but will not have the same structure as the source. You will have to manually structure the provisioned Oracle Database.

Oracle Database Template File

☑ Use Custom Database Template File

○ Software Library   ○ Installed Home   ● Destination Host

Oracle Database Template File Location   /initrd   🖊

f. In the Listener Port section, specify a custom port that can be used for configuring the listener. If you leave this field blank, then the port 1521 or any other available port is honored.

Note that if the destination host has Oracle Grid Infrastructure installed, then the field is disabled because the database will use the listener configured in the Oracle home of the Oracle Grid Infrastructure.

g. Click **Next**.

6. On the Review page, do the following:

a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

> **Note:** Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

> **Note:** After submitting the Deployment Procedure, you should see a link on top of the page. Click the link to view its progress. Alternatively, from the Deployments page, in the Deployment Procedure Manager section, click **Procedure Completion Status** to view the status of the Deployment Procedure.

7. After the Deployment Procedure ends successfully, instrument the database to collect configuration information. For information about instrumenting the database, see Section 5.6.2.2, "Instrumenting the Database for Configuration Collections".

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If you tried provisioning Oracle Database 11g Release 2 (11.2.x.x) and if the Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

5.5 Provisioning Oracle Database Using Gold Image

This section describes how you can provision Oracle Database using a gold image from the Software Library.

This option is best suited when you have a copy of a stable, well-tested, and patched Oracle Database instance in the Software Library as a gold image. This option scores over a fresh installation because you save time in patching and testing a fresh instance.

In particular, this section covers the following:

- [Prerequisites](#)

- [Provisioning Procedure](#)

5.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that the gold image is available either in the Software Library or in a shared, staging location. If it is available in the Software Library, then ensure that it is in *Ready* status.

  To stage a gold image to the Software Library or to a shared location, follow the instructions outlined in Appendix K, "Staging Gold Image to Software Library or Staging Location".

- If you want to use standard response files while provisioning the database, then edit the `dbca.rsp` file available on the installation medium, and store it as a generic component in the Software Library. Alternatively, you can also maintain this file on any host.

  To understand how response files work, see the *Oracle Database Installation Guide for Linux* available at one of these URLs:

  For Oracle Database 11g Release 2

  http://www.oracle.com/pls/db112/homepage

  For Oracle Database 11g Release 1

  http://www.oracle.com/pls/db111/homepage

  For Oracle Database 10g Release 2

  http://www.oracle.com/pls/db102/homepage

  http://www.oracle.com/pls/db10g/portal.portal_demo3

  To understand how you can upload the response file to the Software Library, see Section 5.5.1.1, "Uploading Database Response File to Oracle Software Library".

- If you want to use a custom template to create a structure for the database, then create a template (a `.dbt` file), and store it in a location accessible from the target hosts. The file may be on the target host or on a shared location.

  To understand how a template can be created and used for creating databases, see Appendix D, "Creating Templates for Oracle Database".

  To understand how you can upload the template to the Software Library, see Section 5.5.1.2, "Uploading Database Template File to Oracle Software Library".

- Ensure that the operating system user *oracle* is available on all the nodes of the cluster.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Database.

  Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

  Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

- By default, the Deployment Procedure displays the operating system groups for Database Administrator (OSDBA) and Database Operator (OSOPER). Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in *Oracle Database Quick Installation Guide* available at:

  `http://www.oracle.com/pls/db112/homepage`

  The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

  For more information, see *Oracle Database 2 Day DBA Guide* available at:

  `http://www.oracle.com/pls/db112/homepage`

5.5.1.1 Uploading Database Response File to Oracle Software Library

To upload a database response file, follow these steps:

1. In Grid Control, click **Deployments**, and then click the **Provisioning** subtab.

2. On the Components page, from the table, expand **Oracle Components**, and then **RAC Provisioning**.

3. Under RAC Provisioning, expand **Database Response File**.

4. Under Database Template File, select **Oracle Database Response File** and click **Edit**.

5. On the Edit Component page, click the **Upload File** tab.

6. On the Upload File page, select one of the following options and upload the database response file:

    a. **Upload from Agent Machine**, if the template file is available on the host where the Management Agent is running.

    b. **Upload from Local Machine**, if the template file is available on the local host, that is, the host from where you are accessing the Grid Control console.

7. Click **Finish**. Enterprise Manager Grid Control uploads the template and returns to the Component page.

8. On the Components page, from the table, expand **Oracle Components**, and then **RAC Provisioning**.

9. Under RAC Provisioning, expand **Database Template File**. You should see **Oracle Database Template File** under it. Ensure that status of this component is *Ready*.

5.5.1.2 Uploading Database Template File to Oracle Software Library

To upload a database template file, follow these steps:

1. In Grid Control, click **Deployments**, and then **Provisioning**.

2. On the Components page, from the table, expand **Oracle Components**, and then **RAC Provisioning**.

3. Under RAC Provisioning, expand **Database Template File**.

4. Under Database Template File, select **Oracle Database Template File** and click **Edit**.

5. On the Edit Component page, click the **Upload File** tab.

6. On the Upload File page, select one of the following options and upload the database response file:

    a. **Upload from Agent Machine**, if the template file is available on the host where the Management Agent is running.

    b. **Upload from Local Machine**, if the template file is available on the local host, that is, the host from where you are accessing the Grid Control console.

7. Click **Finish**. Enterprise Manager Grid Control uploads the template and returns to the Component page.

8. On the Components page, from the table, expand **Oracle Components**, and then **RAC Provisioning**.

9. Under RAC Provisioning, expand **Database Template File**. You should see **Oracle Database Template File** under it. Ensure that status of this component is *Ready*.

5.5.2 Provisioning Procedure

To provision Oracle Database using a gold image from the Software Library, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Provisioning**. Then click **Schedule Deployment**.

Enterprise Manager Grid Control displays the Select Source and Destination page the Oracle Database Provisioning Deployment Procedure.

4. On the Select Source and Destination page, do the following:

a. In the Select Source section, do one of the following:

If the gold image is stored as a component in the Software Library, then select **Software Library**. Then, click the torch icon for **Component** and select the gold image that is in *Ready* status.

Note: If you do not see the required component in the Software Library, then follow the workaround described in Appendix A, "Troubleshooting Issues".

If the gold image was stored as an image in a staging location while provisioning a database in the past, then select **External Staging Server** and then **Gold Image**. Click the torch icon for **Select Host** and select the host where the gold image is stored. Then click the torch icon for **Stage Location** and select the location on the host where the gold image is available. For Product version, specify the version of the product you are provisionoing.

b. In the Specify Destination Host Settings section, click **Add** and select the target hosts on which you want to provision the gold image of Oracle Database.

Note: On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

Note: You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

Note: If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

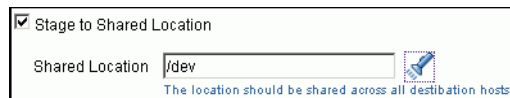If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced section, specify any additional parameters you want to run while installing the single instance Oracle Database. For example, -debug.

You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. Using these parameters, you can even change the installation type of the database. For example, `INSTALL_TYPE=SE` (to install standard edition of database), `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

If you want to stage the gold image to a shared location, then select **Stage to Shared Location** and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. Ensure that this stage location has adequate space for storing the gold image. For example, if the gold image you are provisioning is 1 GB, then maintain at least 2 GB of space in the stage location.

> **Note:** The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select **Staging Location** as the source when you run the Deployment Procedure next time.

d. In the Group Details section, specify the Database Administrator (OSDBA) group and the Database Operator (OSOPER) group you want to use for operating system authentication. Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at:

http://www.oracle.com/pls/db112/homepage

The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

For more information, see Oracle Database 2 Day DBA Guide available at:

http://www.oracle.com/pls/db112/homepage

e. (Optional) In the Create Database section, select **Create and Configure Database** if you want to create a new database and configure it after installing the standalone Oracle Database.

> **IMPORTANT:** If you do not select Create and Configure Database, only the software binaries of Oracle Database are copied to the target host. You will have to manually configure the provisioned Oracle Database using Database Configuration Assistant (DBCA).

f. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

There may be occasions when the Management Agent on that target host is down and as a result, OMS is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when OMS should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select **End After** and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

g. Click **Next**.

Enterprise Manager Grid Control displays either the Database Configuration page or the Review page depending on the selection made in Step 4 (e).

In Step 4 (e), if you had selected **Create and Configure Database**, then follow the instructions outlined in Step 5. Otherwise, go to Step 6.

5. On the Database Configuration page, do the following:

a. In the Database Identifiers section, specify basic configuration details such as the unique SID, the global database name, and the database credentials.

> **Note:**
>
> - SID must be unique for a database on a host. This means, the SID assigned to one database on a host cannot be reused on another database on the same host, but can be reused on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their SIDs need to be unique. However, if you install the third database on another host (host2), then its SID can be db1 or db2.
>
> - Global database name must be unique for a database on a host and also unique for databases across different hosts. This means, the global database name assigned to one database on a host can neither be reused on another database on the same host nor on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their global database names need to be unique. And if you install the third database on another host (host2), the global database name of even this database must be unique and different from all other names registered with Enterprise Manager Grid Control.
>
> - The database credentials you specify here will be used on all the destination hosts. However, after provisioning, if you want to change the password for any database, then you must change it manually.

   b. In the Database Storage Options section, select the location to store your data files.

If you want to use a file system, then select **File System** and specify the full path to the location where the data file is present. For example, %ORACLE_BASE%/oradata /u01/product/db/oradata

If you want to use ASM, then select **Automatic Storage Management**, and click the torch icon to select the disk group name and specify ASMSNMP password. The Disk Group Name List window appears and displays the disk groups that are common on all the destination hosts.

For example, if the destination hosts are A and B, and if host A has three disk groups, mainly DATA1, DATA2, DATA3, and if host B has four disk groups, mainly DATA1, DATA3, DATA4, DATA5, then the Disk Group Name List windows shows DATA1 and DATA3 because these two are common across the two hosts.

> **Note:** This option will be enabled only if all the destination hosts have Oracle ASM 11g Release 2 (11.2) or higher configured.

   c. (Optional) In the Advanced Database Parameters section, specify any additional DBCA parameters that you want to pass while configuring the database. For example, `-characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

   d. (Optional) In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you want to use a custom response file that has options enabled. If you have stored the response file as a generic component in

the Software Library, then select **Software Library** and click the torch icon for Component to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

Note: While creating the database, the values specified in the page for mandatory parameters, such as SID and so on, will take precedence over the values specified for the same parameters in the response file.

---

e. In the Oracle Database Template File subsection, select **Use Custom Database Template File** and use the template file that is being used in the existing instance of Oracle Database that you are cloning. If you have stored the template file as a generic component in the Software Library, then select **Software Library** and click the torch icon for **Component** to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

IMPORTANT: If you do not select Use Custom Database Template File, the provisioned Oracle Database will be configured but will not have the same structure as the source. You will have to manually structure the provisioned Oracle Database.

---

f. In the Listener Port section, specify a custom port that can be used for configuring the listener. If you leave this field blank, then the port 1521 or any other available port is honored.

Note that if the destination host has Oracle Grid Infrastructure installed, then the field is disabled because the database will use the listener configured in the Oracle home of the Oracle Grid Infrastructure.

g. Click **Next**.

6. On the Review page, do the following:

a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

> **Note:** Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

> **Note:** After submitting the Deployment Procedure, you should see a link on top of the page. Click the link to view its progress. Alternatively, from the Deployments page, in the Deployment Procedure Manager section, click **Procedure Completion Status** to view the status of the Deployment Procedure.

7. After the Deployment Procedure ends successfully, instrument the database to collect configuration information. For information about instrumenting the database, see Section 5.6.2.2, "Instrumenting the Database for Configuration Collections".

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If you tried provisioning Oracle Database 11g Release 2 (11.2.x.x) and if the Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

5.6 Provisioning Standalone Oracle Database Using Installation Binaries

This section describes how you can provision one or more standalone Oracle Database using the installation binaries uploaded to the Software Library.

This option is best suited when you want a completely new installation to be provisioned across multiple hosts. Of course, understandably, this is a fresh installation and you will have to update it with all the latest patches that have been released so far.

In particular, this section covers the following:

- Prerequisites

- Provisioning Procedure

5.6.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that the installation binaries are downloaded, and archived and uploaded either as a component in the Software Library or as an archived file in a staging location.

    For information about downloading the installation binaries, and archiving and uploading them to the Software library or to a staging location, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location".

- If you want to use standard response files while provisioning the database, then edit the dbca.rsp file available on the installation medium, and store it as a generic component in the Software Library. Alternatively, you can also maintain this file on any host.

    To understand how response files work, see the *Oracle Database Installation Guide for Linux* available at one of these URLs:

    For Oracle Database 11g Release 2

    http://www.oracle.com/pls/db112/homepage

    For Oracle Database 11g Release 1

    http://www.oracle.com/pls/db111/homepage

    For Oracle Database 10g Release 2

    http://www.oracle.com/pls/db102/homepage

    http://www.oracle.com/pls/db10g/portal.portal_demo3

    To understand how you can upload the response file to the Software Library, see Section 5.5.1.1, "Uploading Database Response File to Oracle Software Library".

- If you want to use a custom template to create a structure for the database, then create a template (a .dbt file), and store it in a location accessible from the target hosts. The file may be on the target host or on a shared location.

    To understand how a template can be created and used for creating databases, see Appendix D, "Creating Templates for Oracle Database".

    To understand how you can upload the template to the Software Library, see Section 5.5.1.2, "Uploading Database Template File to Oracle Software Library".

- Ensure that the operating system user *oracle* is available on all the nodes of the cluster.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Database.

  Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

  Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

- By default, the Deployment Procedure displays the operating system groups for Database Administrator (OSDBA) and Database Operator (OSOPER). Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in *Oracle Database Quick Installation Guide* available at:

  http://www.oracle.com/pls/db112/homepage

  The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

  For more information, see *Oracle Database 2 Day DBA Guide* available at:

  http://www.oracle.com/pls/db112/homepage

5.6.2 Provisioning Procedure

To provision a fresh standalone Oracle Database, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Provisioning**. Then click **Schedule Deployment**.

Enterprise Manager Grid Control displays the Select Source and Destination page.

4. On the Select Source and Destination page, do the following:

   a. In the Select Source section, do one of the following:

   If the software binaries are stored as a component in the Software Library, then select **Software Library**. Then, click the torch icon for **Component** and select the component that has the archived software binaries. Ensure that you select only components that are in "Ready" status. When you select a component from the Software Library, Enterprise Manager Grid Control automatically populates the component location.

   > **Note:** If you do not see the required component in the Software Library, then follow the workaround described in Appendix A, "Troubleshooting Issues".

   If the software binaries are stored as an archived file in a staging location, then select **External Staging Server** and then **Shiphome**. Click the torch icon for **Select Host** and select the host where the archived file is stored. Then click the torch icon for **Stage Location** and select the location on the host where the archived file is available. For Product version, specify the version of the product you are provisionoing.

   b. In the Specify Destination Host Settings section, click **Add** and select the target hosts on which you want to install the standalone Oracle Database.

   > **Note:** On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

Note: You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

Note: If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced section, specify any additional parameters you want to run while installing the database. For example, `-debug`.

You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. Using these parameters, you can even change the installation type of the database. For example, `INSTALL_TYPE=SE` (to install standard edition of database), `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

While installing software binaries from the Software Library, if you want to also stage them to a shared location, then select **Stage to Shared Location** and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. This staged location can also act as a source for future deployments.

Note: The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select **Staging Location** as the source when you run the Deployment Procedure next time.

---

d. In the Group Details section, specify the Database Administrator (OSDBA) group and the Database Operator (OSOPER) group you want to use for operating system authentication. Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at:

http://www.oracle.com/pls/db112/homepage

The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

For more information, see Oracle Database 2 Day DBA Guide available at:

http://www.oracle.com/pls/db112/homepage

e. (Optional) In the Create Database section, select **Create and Configure Database** if you want to create a new database and configure it after installing the standalone Oracle Database.

IMPORTANT: If you do not select **Create and Configure Database**, only the software binaries of Oracle Database are copied to the target host. You will have to manually configure the provisioned Oracle Database using Database Configuration Assistant (DBCA).

---

f. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to

run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

There may be occasions when the Management Agent on that target host is down and as a result, OMS is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when OMS should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select **End After** and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

g. Click **Next**.

Enterprise Manager Grid Control displays either the Database Configuration page or the Review page depending on the selection made in Step 4 (e).

In Step 4 (e), if you had selected **Create and Configure Database**, then follow the instructions outlined in Step 5. Otherwise, go to Step 6.

5. On the Database Configuration page, do the following:

a. In the Database Identifiers section, specify basic configuration details such as the SID, the global database name, and the database credentials.

Database Identifiers		
Host Name	**SID**	**Global Database Name**
dadvma0122.us.oracle.com	orcl	orcl.us.oracle.com

* Database Password	••••••••••••
* Confirm Database Password	••••••••••••

> **Note:**
>
> - SID must be unique for a database on a host. This means, the SID assigned to one database on a host cannot be reused on another database on the same host, but can be reused on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their SIDs need to be unique. However, if you install the third database on another host (host2), then its SID can be db1 or db2.
>
> - Global database name must be unique for a database on a host and also unique for databases across different hosts. This means, the global database name assigned to one database on a host can neither be reused on another database on the same host nor on another database on a different host. For example, if you have two databases (db1 and db2) on a host (host1), then their global database names need to be unique. And if you install the third database on another host (host2), the global database name of even this database must be unique and different from all other names registered with Enterprise Manager Grid Control.
>
> - The database credentials you specify here will be used on all the destination hosts. However, after provisioning, if you want to change the password for any database, then you must change it manually.

b. In the Database Storage Options section, select the location to store your data files.

If you want to use a file system, then select **File System** and specify the full path to the location where the data file is present. For example, %ORACLE_BASE%/oradata /u01/product/db/oradata

If you want to use ASM, then select **Automatic Storage Management**, and click the torch icon to select the disk group name and specify ASMSNMP password. The Disk Group Name List window appears and displays the disk groups that are common on all the destination hosts.

For example, if the destination hosts are A and B, and if host A has three disk groups, mainly DATA1, DATA2, DATA3, and if host B has four disk groups, mainly DATA1, DATA3, DATA4, DATA5, then the Disk Group Name List windows shows DATA1 and DATA3 because these two are common across the two hosts.

> **Note:** This option will be enabled only if all the destination hosts have Oracle ASM 11g Release 2 (11.2) or higher configured.

c. (Optional) In the Advanced Database Parameters section, specify any additional DBCA parameters that you want to pass while configuring the database. For example, `-characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

d. (Optional) In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you want to use a custom response file that has options enabled. If you have stored the response file as a generic component in

the Software Library, then select **Software Library** and click the torch icon for Component to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

Note: While creating the database, the values specified in the page for mandatory parameters, such as SID and so on, will take precedence over the values specified for the same parameters in the response file.

Oracle Database Response File
☑ Use Custom Database Response File

○ Software Library ⦿ Destination Host

Response File Location /initrd

e. In the Oracle Database Template File subsection, select **Use Custom Database Template File** and use the template file that is being used in the existing instance of Oracle Database that you are cloning. If you have stored the template file as a generic component in the Software Library, then select **Software Library** and click the torch icon for **Component** to select the component. Otherwise, select **Destination Host** and click the torch icon to select the location where the file resides.

IMPORTANT: If you do not select Use Custom Database Template File, the provisioned Oracle Database will be configured but will not have the same structure as the source. You will have to manually structure the provisioned Oracle Database.

Oracle Database Template File
☑ Use Custom Database Template File

○ Software Library ○ Installed Home ⦿ Destination Host

Oracle Database Template File Location /initrd

f. In the Listener Port section, specify a custom port that can be used for configuring the listener. If you leave this field blank, then the port 1521 or any other available port is honored.

Note that if the destination host has Oracle Grid Infrastructure installed, then the field is disabled because the database will use the listener configured in the Oracle home of the Oracle Grid Infrastructure.

g. Click **Next**.

6. On the Review page, do the following:

a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

> **Note:** Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

> **Note:** After submitting the Deployment Procedure, you should see a link on top of the page. Click the link to view its progress. Alternatively, from the Deployments page, in the Deployment Procedure Manager section, click **Procedure Completion Status** to view the status of the Deployment Procedure.

7. After the Deployment Procedure ends successfully, instrument the database to collect configuration information. For information about instrumenting the database, see Section 5.6.2.2, "Instrumenting the Database for Configuration Collections".

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If you tried provisioning Oracle Database 11g Release 2 (11.2.x.x) and if the Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

5.6.2.1 Manually Collecting and Uploading Configuration Information to My Oracle Support

To manually collect the configuration information, follow these steps:

1. On the host where the Oracle home was deployed or patched, navigate to the following location. Here, `<OracleHome>` is the Oracle home directory of the Oracle home you deployed or patched.

```
<OracleHome>/ccr/bin
```

2. Collect configuration information by running the following command:

   ```
   <OracleHome>/ccr/bin/emCCR collect
   ```

   For Oracle Configuration Manager 10.2.7 and higher, the collected configuration information is stored in the `<OracleHome>/ccr/hosts/<hostname>/state/upload/ocmconfig.jar` file. For lower versions of Oracle Configuration Manager, the collected configuration information is stored in the `<OracleHome>/ccr/state/upload/ocmconfig.jar` file. When you run the same command next time, the `ocmconfig.jar` file gets overwritten with fresh data. Therefore, at any point, you will see only one `ocmconfig.jar` file.

3. Upload the ocmconfig.jar file to a Service Request on My Oracle Support at the following URL:

   http://support.oracle.com

5.6.2.2 Instrumenting the Database for Configuration Collections

After the Deployment Procedure ends, once the Oracle Configuration Manager is installed and configured either in connected or disconnected mode, you must instrument the database to collect configuration information.

Oracle Configuration Manager Installed in Connected Mode

If you had installed and configured Oracle Configuration Manager in a connected mode, then to instrument the database, see the steps outlined in the Instrumenting the Database for Configuration Collections chapter of the Oracle Configuration Manager Quick Start Guide available at the following URL. The steps provided in this guide must be repeated for every database instance being patched or deployed.

http://www.oracle.com/technology/documentation/ocm.html

Oracle Configuration Manager Installed in Disconnected Mode

If you had installed and configured Oracle Configuration Manager in a disconnected mode, then to instrument the database, see the steps (except for the last step) outlined in the Instrumenting the Database for Configuration Collections chapter of the Oracle Configuration Manager Quick Start Guide available at the following URL:

http://www.oracle.com/technology/documentation/ocm.html

As a final step, after you run the emCCR collect command and collect the configuration information, manually upload the ocmconfig.jar file to a Service Request on My Oracle Support at the following URL. The steps provided in this guide must be repeated for every database instance being patched or deployed.

http://support.oracle.com/

6

Provisioning Oracle Grid Infrastructure for Standalone Servers

This chapter explains how you can mass-deploy Oracle Grid Infrastructure for standalone servers in an unattended, repeatable, and reliable manner. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedure

- Supported Releases

- Provisioning Oracle Grid Infrastructure Using an Existing Installation

- Provisioning Oracle Grid Infrastructure Using Certified Gold Image

- Provisioning Oracle Grid Infrastructure Using Installation Binaries

> **Note:** Before you run the Deployment Procedure as outlined in this chapter to provision an Oracle Grid Infrastructure for a standalone server, Oracle strongly recommends you to check the prerequisites to ensure that you meet all the necessary requirements for provisioning an Oracle Grid Infrastructure. To do so, simulate the Deployment Procedure as described in Chapter 4, "Simulating Database Provisioning Deployment Procedures to Check Prerequisites", and resolve all operational conflicts and issues beforehand.

6.1 Getting Started

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

Table 6–1   Getting Started with Provisioning Oracle Grid Infrastructure

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Grid Infrastructure. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 6.2, "Deployment Procedure".
Step 2	**Knowing About the Supported Releases** Know what releases of Oracle Grid Infrastructure can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 6.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle Grid Infrastructure. Select the use case that best matches your requirement.	■ To learn about cloning an existing Oracle Grid Infrastructure, see Section 6.4, "Provisioning Oracle Grid Infrastructure Using an Existing Installation". ■ To learn about provisioning Oracle Grid Infrastructure using a gold image, see Section 6.5, "Provisioning Oracle Grid Infrastructure Using Certified Gold Image". ■ To learn about provisioning a standalone Oracle Grid Infrastructure using its software binaries, see Section 6.6, "Provisioning Oracle Grid Infrastructure Using Installation Binaries".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for cloning an existing Oracle Grid Infrastructure, see Section 6.4.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Grid Infrastructure using a gold image, seeSection 6.5.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Grid Infrastructure using its software binaries, see Section 6.6.1, "Prerequisites".

Table 6–1 (Cont.) Getting Started with Provisioning Oracle Grid Infrastructure

Step	Description	Reference Links
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Grid Infrastructure.	■ To clone an existing Oracle Grid Infrastructure, follow the steps explained in Section 6.4.2, "Provisioning Procedure". ■ To provision Oracle Grid Infrastructure using a gold image, follow the steps explained in Section 6.5.2, "Provisioning Procedure". ■ To provision Oracle Grid Infrastructure using its software binaries, follow the steps explained in Section 6.6.2, "Provisioning Procedure".

6.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for provisioning Oracle Grid Infrastructure on standalone servers:

■ *Oracle Grid Infrastructure Provisioning for Standalone Servers*

Using this Deployment Procedure, you can perform the following operations:

■ Deploy Oracle Grid Infrastructure software using certified gold images, live reference installations, and installation media

■ Optionally configure Oracle Automatic Storage Management (Oracle ASM)

■ Use staging location for fast file transfer

Therefore, if you have a standalone server and want to use high availability features, then use this Deployment Procedure to provision Oracle Grid Infrastructure on the server before provisioing a standalone database using the Deployment Procedure as described in Chapter 5, "Provisioning Oracle Database".

The high available feature monitored in Enterprise Manager Grid Control appears as *Oracle High Availability Service*.

6.3 Supported Releases

Using this Deployment Procedure, you can provision the following release of Oracle Grid Infrastructure across standalone servers:

■ Oracle Grid Infrastructure 11g Release 2 (11.2)

6.4 Provisioning Oracle Grid Infrastructure Using an Existing Installation

This section describes how you can provision Oracle Grid Infrastructure using an existing installation that is running on a host monitored by Enterprise Manager Grid Control.

This option is best suited when you have a running instance of Oracle Grid Infrastructure that is stable and has all the latest patches applied, and you want to make identical copies of that software on multiple standalone servers.

In particular, this section covers the following:

- [Prerequisites](#)

- [Provisioning Procedure](#)

6.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

    To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- Meet the memory and disk space requirements for Oracle Grid Infrastructure installation on the target hosts. For information about the memory and disk space requirements for Oracle Grid Infrastructure installation, refer to the *Oracle Database Installation Guide 11g Release 2 (11.2)* available at:

    http://www.oracle.com/technology/documentation/database.html

- Ensure that the following groups exist. If you do not belong to these groups, then the Deployment Procedure automatically adds you to the groups.

    For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)* available at:

    http://www.oracle.com/technology/documentation/database.html

    - ASM Database Administrator (OSDBA) group

    - ASM Instance Operator (OSOPER) group

    - ASM Instance Administrator (OSASM) group

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

    If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

    For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

    For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Grid Infrastructure. Also ensure that the staging area is accessible from the Oracle Management Service (OMS) and the target host.

  Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the operating system groups you specify for OSDBA, OSOPER, and OSASM already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.

6.4.2 Provisioning Procedure

To provision Oracle Grid Infrastructure using an existing installation, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select O**racle Grid Infrastructure Provisioning for Standalone Servers** and click **Schedule Deployment**.

4. On the Select Source and Destination page, in the Select Source section, select **Existing Installation**. Then, click the torch icon for **Source Host** and select the host that has the existing installation. When you select the source, in the Source Host Details subsection, Enterprise Manager Grid Control automatically populates the fields, such as **Oracle home**, **Working Directory**, and **Files to exclude**.

   > **Note:** Enterprise Manager Grid Control creates a ZIP file of the Oracle home in the working directory of the source host, and then copies this ZIP file to the working directory of the destination host.

5. In the Source Host Credentials subsection, specify the operating system credentials of the selected host where Oracle Grid Infrastructure is installed. You can choose to use the preferred credentials stored in the Management Repository or override them with a new set of credentials.

6. In the Specify Destination Host Settings section, do the following:

   a. Click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure. Once you add a host, you will see default values for Oracle Base, Oracle Home, and other fields. You can either retain the default values or specify another custom value.

      > **Note:** If you decide to specify an Oracle home location of your choice, then Oracle recommends you to maintain the Oracle home *outside* the installing user's home, and NOT inside. For example, if your home is `/home/john`, then you must not provision Oracle Grid Infrastructure in `/home/john/gi11g`.

b. If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts. If you want to use different paths for each host, then select **Different for each host**.

> **Note:** If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

c. From the **Credentials** list, select one of the following to provide the operating system credentials of the destination host:

Preferred, if you want to use the preferred credentials stored in the Management Repository. For this, you must have already registered the preferred credentials.

Override Preferred Credentials, if you want to override the preferred credentials and use another set of credentials for this operation. If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation. Also ensure that the Deployment Procedure has been customized to support locked environments.

For information about customizing Deployment Procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

d. Optionally, click **Customize Host Settings** if you want to provide additional, custom information about the destionation host and customize its settings.

For example, you may want to use an alternative host name instead of using the first host name identified on the system. Similarly, you may want to specify a custom name for the Oracle home. Also, if the Oracle Management Agent home credentials are different from the Oracle home credentials, then you may want to specify the agent home credentials. Under such circumstances, you can use this option to customize the host settings.

e. In the Advanced Installation Parameters subsection, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle Grid Infrastructure. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

f. If you want to stage the software binaries to a shared location, then select **Stage to Shared Location** and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. The software binaries will be staged only once to the staging location. Once they are staged, they will be copied to the working directory of each of the destination hosts.

> **Note:** The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select **Staging Location** as the source when you run the Deployment Procedure next time.

7. In the Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), and ASM Instance Administrator (OSASM) group you want to use for operating system authentication.

Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at

http://www.oracle.com/pls/db112/homepage

The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

For more information, see Oracle Database 2 Day DBA Guide available at:

http://www.oracle.com/pls/db112/homepage

> **Note:** This section appears only if the source you selected is for a UNIX platform.

8. In the Listener Port section, specify a custom listener port that can be used. Note that if you leave this field blank, then the port 1521 or any other available port is honored.

9. In the Configure ASM section, select **Configure ASM** if you want to configure ASM as part of this Deployment Procedure.

> **Note:** If you do not select **Configure ASM**, then the software binaries of Oracle Grid Infrastructure are copied to the destination host and a high-available listener is configured.
>
> If you want to configure Oracle ASM after provisioning, then you will have to manually configure it using Auomatic Storage Management Configuration Assistant (ASMCA).
>
> The high-available listener is configured in both cases, that is, when you select **Configure ASM** on this page and also when you manually configure Oracle ASM after provisioning.

10. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select End After and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

11. Click **Next**.

> **Note:** If you have Oracle ASM or Oracle Grid Infrastructure already installed on the destination host, then you are prompted to either deconfigure them or specify another set of destination hosts that do not have them.

12. If you chose to configure ASM, then on the Configure Oracle ASM page, add a new disk group and specify its credentials. Otherwise, go to Step (14).

    a. In the ASM Storage section, click **Add Disk Group** to add disk groups.

    On the Disk Group Configuration window that appears, by default, you will not see any disk groups detected. Specify a disk discovery string and search for disks on the destination hosts. For example, if the disks are /dev/raw/raw1, /dev/raw/raw2, and /dev/raw/raw3 on the destination hosts, then specify /dev/raw/* for **Disk Discovery String** and click **Search**. Ensure that the disk paths are the same across all destination hosts. Once the disks are detected, select the ones you are interested in.

    Alternatively, if you know the exact path to the disks, then you can manually specify a comma-separated list.

    After selecting or specifying the disks, specify a name for the disk group.

    b. In the ASM Password section, specify the ASM password for SYS and ASMSNMP user accounts. Ensure that you meet the requirements described in Section 6.4.2.1, "ASM Password Requirements".

13. Click **Next**.

14. On the Review page, do the following:

    a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

    b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

    > **Note:** Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

    > **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

6.4.2.1 ASM Password Requirements

While specifying the ASM password, follow these guidelines:

Minimum Password Requirements

- Password must not exceed 30 characters.

- Password must not contain invalid characters: ! @ % ^ & * ( ) + = \ | ` ~ [{] } ; : ' " , < > ?

- Password must not be `change_on_install` or `asmsnmp`.

- A user name must not be a password.

- An empty password cannot be accepted.

Password Recommendations

- contains at least one lowercase letter.

- contains at least one uppercase letter.

- contains at least one digit.

- is at least 8 characters in length.

- uses the database character set which can include the underscore (_), dollar ($), and pound sign (#) characters.

- if contains special characters, including beginning the password with a number or symbol, then enclose the password with double-quotation marks.

- should not be an actual word.

6.5 Provisioning Oracle Grid Infrastructure Using Certified Gold Image

This section describes how you can provision Oracle Grid Infrastructure using a certified gold image.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

6.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- Ensure that the gold image is available either in the Software Library or in a staging location.

  To stage a gold image to the Software Library or to a shared location, follow the instructions outlined in Appendix K, "Staging Gold Image to Software Library or Staging Location".

- Meet the memory and disk space requirements for Oracle Grid Infrastructure installation on the target hosts. For information about the memory and disk space requirements for Oracle Grid Infrastructure installation, refer to the *Oracle Database Installation Guide 11g Release 2 (11.2)* available at:

  http://www.oracle.com/technology/documentation/database.html

- Ensure that the following groups exist. If you do not belong to these groups, then the Deployment Procedure automatically adds you to the groups.

  For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)* available at:

  http://www.oracle.com/technology/documentation/database.html

  - ASM Database Administrator (OSDBA) group

  - ASM Instance Operator (OSOPER) group

  - ASM Instance Administrator (OSASM) group

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

■ Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Grid Infrastructure. Also ensure that the staging area is accessible from the Oracle Management Service (OMS) and the target host.

Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

■ Ensure that the operating system groups you specify for OSDBA, OSOPER, and OSASM already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.

6.5.2 Provisioning Procedure

To provision Oracle Grid Infrastructure using a certified gold image, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select **Oracle Grid Infrastructure Provisioning for Standalone Servers** and click **Schedule Deployment**.

4. On the Select Source and Destination page, in the Select Source section, do one of the following:

   a. If the gold image is stored as a component in Oracle Software Library (Software Library), then select **Software Library**. Then, click the torch icon for **Component** and select the component that has the gold image.

   b. If the gold image was stored as an image in a staging location while provisioning a database in the past, then select **Staging Location** and then **Gold Image**. Click the torch icon for **Select Host** and select the host where the gold image is stored. Then click the torch icon for **Stage Location** and select the location on the host where the gold image is available.

   > **Note:** This must be a gold image that was staged earlier by running the same Deployment Procedure either for provisioning the software or for staging the gold image. You must not select a gold image that was manually created and staged. For information about running the Deployment Procedure only to stage the gold image in a staging location, see Appendix K, "Staging Gold Image to Software Library or Staging Location".

5. In the Specify Destination Host Settings section, do the following:

   a. Click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure. Once you add a host, you will see default values for Oracle Base, Oracle Home, and other fields. You can either retain the default values or specify another custom value.

> **Note:** If you decide to specify an Oracle home location of your choice, then Oracle recommends you to maintain the Oracle home *outside* the installing user's home, and NOT inside.

b. If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts. If you want to use different paths for each host, then select **Different for each host**.

> **Note:** If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

c. From the **Credentials** list, select one of the following to provide the operating system credentials of the destination host:

Preferred, if you want to use the preferred credentials stored in the Management Repository. For this, you must have already registered the preferred credentials.

Override Preferred Credentials, if you want to override the preferred credentials and use another set of credentials for this operation. If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation. Also ensure that the Deployment Procedure has been customized to support locked environments.

For information about customizing Deployment Procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

d. Optionally, click **Customize Host Settings** if you want to provide additional, custom information about the destination host and customize its settings.

For example, you may want to use an alternative host name instead of using the first host name identified on the system. Similarly, you may want to specify a custom name for the Oracle home. Also, if the Oracle Management Agent home credentials are different from the Oracle home credentials, then you may want to specify the agent home credentials. Under such circumstances, you can use this option to customize the host settings.

e. In the Advanced Installation Parameters subsection, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle Grid Infrastructure. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

f. If you want to stage the software binaries to a shared location, then select Stage to Shared Location and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. The software binaries will be staged only once to the staging location. Once they are staged, they will be copied to the working directory of each of the destination hosts.

> **Note:** The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select Staging Location as the source when you run the Deployment Procedure next time.

6. In the Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), and ASM Instance Administrator (OSASM) group you want to use for operating system authentication.

   Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at

   http://www.oracle.com/pls/db112/homepage

   The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

   For more information, see Oracle Database 2 Day DBA Guide available at:

   http://www.oracle.com/pls/db112/homepage

   > **Note:** This section appears only if the source you selected is for a UNIX platform.

7. In the Listener Port section, specify a custom listener port that can be used. Note that if you leave this field blank, then the port 1521 or any other available port is honored.

8. In the Configure ASM section, select **Configure ASM** if you want to configure ASM as part of this Deployment Procedure.

   > **Note:** If you do not select **Configure ASM**, then the software binaries of Oracle Grid Infrastructure are copied to the destination host and a high-available listener is configured.
   >
   > If you want to configure Oracle ASM after provisioning, then you will have to manually configure it using Auomatic Storage Management Configuration Assistant (ASMCA).
   >
   > The high-available listener is configured in both cases, that is, when you select **Configure ASM** on this page and also when you manually configure Oracle ASM after provisioning.

9. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

   There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is

indefinite. If you want to have a defninte period, then select End After and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

10. Click **Next**.

> **Note:** If you have Oracle ASM or Oracle Grid Infrastructure already installed on the destination host, then you are prompted to either deconfigure them or specify another set of destination hosts that do not have them.

11. If you chose to configure ASM, then on the Configure Oracle ASM page, add a new disk group and specify its credentials. Otherwise, go to Step (13).

   a. In the ASM Storage section, click **Add Disk Group** to add disk groups. From the Disk Group Configuration window that appears, you can either select one of the disks discovered by Enterprise Manager Grid Control or manually specify the disk you want to use.

   b. In the ASM Password section, specify the ASM password for SYS and ASMSNMP user accounts. Ensure that you meet the requirements described in Section 6.4.2.1, "ASM Password Requirements".

12. Click **Next**.

13. On the Review page, do the following:

   a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

   b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

   > **Note:** Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

   > **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

6.6 Provisioning Oracle Grid Infrastructure Using Installation Binaries

This section describes how you can provision Oracle Grid Infrastructure using the archived installation binaries.

This option is best suited when you want a completely new installation to be provisioned across multiple standalone servers. Of course, understandably, this is a fresh installation and you will have to update it with all the latest patches that have been released so far.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

6.6.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers
- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.
- Ensure that the installation binaries are downloaded, and archived and uploaded either as a component in the Software Library or as an archived file in a staging location.

  To understand how you can download the installation binaries, and archive and uload them to the Software Library or to a shared location, see Appendix L,

"Downloading Installation Binaries and Uploading to Software Library or Staging Location".

■ Meet the memory and disk space requirements for Oracle Grid Infrastructure installation on the target hosts. For information about the memory and disk space requirements for Oracle Grid Infrastructure installation, refer to the *Oracle Database Installation Guide 11g Release 2 (11.2)* available at:

http://www.oracle.com/technology/documentation/database.html

■ Ensure that the following groups exist. If you do not belong to these groups, then the Deployment Procedure automatically adds you to the groups.

For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)* available at:

http://www.oracle.com/technology/documentation/database.html

– ASM Database Administrator (OSDBA) group

– ASM Instance Operator (OSOPER) group

– ASM Instance Administrator (OSASM) group

Prerequisites for Operators

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

■ Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Grid Infrastructure. Also ensure that the staging area is accessible from the Oracle Management Service (OMS) and the target host.

Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

■ Ensure that the operating system groups you specify for OSDBA, OSOPER, and OSASM already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.

6.6.2 Provisioning Procedure

To provision Oracle Grid Infrastructure using the archived software binaries, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select **Oracle Grid Infrastructure Provisioning for Standalone Servers** and click **Schedule Deployment**.

4. On the Select Source and Destination page, in the Select Source section, do one of the following:

    a. If the software binaries are stored as a component in Oracle Software Library (Software Library), then select **Software Library**. Then, click the torch icon for **Component** and select the component that has the archived software binaries. Ensure that you select only components that are in "Ready" status. When you select a component from the Software Library, Enterprise Manager Grid Control automatically populates the component location.

    b. If the software binaries are stored as an archived file in a staging location, then select **Staging Location** and then **Shiphome**. Click the torch icon for **Select Host** and select the host where the archived file is stored. Then click the torch icon for **Stage Location** and select the location on the host where the archived file is available.

    For information about archiving the software binaries from an installation medium and posting them to the Software Library or a staging location, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location".

5. In the Specify Destination Host Settings section, do the following:

    1. Click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure. Once you add a host, you will see default values for Oracle Base, Oracle Home, and other fields. You can either retain the default values or specify another custom value.

> **Note:** If you decide to specify an Oracle home location of your choice, then Oracle recommends you to maintain the Oracle home *outside* the installing user's home, and NOT inside.

    2. If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts. If you want to use different paths for each host, then select **Different for each host**.

> **Note:** If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

    3. From the **Credentials** list, select one of the following to provide the operating system credentials of the destination host:

    Preferred, if you want to use the preferred credentials stored in the Management Repository. For this, you must have already registered the preferred credentials.

    Override Preferred Credentials, if you want to override the preferred credentials and use another set of credentials for this operation. If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the

locked account for performing the provisioning operation. Also ensure that the Deployment Procedure has been customized to support locked environments.

For information about customizing Deployment Procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

4. Optionally, click **Customize Host Settings** if you want to provide additional, custom information about the destination host and customize its settings.

For example, you may want to use an alternative host name instead of using the first host name identified on the system. Similarly, you may want to specify a custom name for the Oracle home. Also, if the Oracle Management Agent home credentials are different from the Oracle home credentials, then you may want to specify the agent home credentials. Under such circumstances, you can use this option to customize the host settings.

5. In the Advanced Installation Parameters subsection, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle Grid Infrastructure. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

6. If you want to stage the software binaries to a shared location, then select Stage to Shared Location and specify a shared, NFS-mounted location that is accessible by Oracle Management Service and all the destination hosts. The software binaries will be staged only once to the staging location. Once they are staged, they will be copied to the working directory of each of the destination hosts.

Note: The staged location can also act as a source for future deployments, that is, once the software binaries are available in the staged location, you can select Staging Location as the source when you run the Deployment Procedure next time.

---

6. In the Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), and ASM Instance Administrator (OSASM) group you want to use for operating system authentication.

Ensure that these groups already exist on the hosts you select for provisioning. If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in Oracle Database Quick Installation Guide available at

`http://www.oracle.com/pls/db112/homepage`

The new groups you create or the alternative groups you specify automatically get SYSDBA and SYSOPER privileges after the database is configured.

For more information, see Oracle Database 2 Day DBA Guide available at:

`http://www.oracle.com/pls/db112/homepage`

Note: This section appears only if the source you selected is for a UNIX platform.

---

7. In the Listener Port section, specify a custom listener port that can be used. Note that if you leave this field blank, then the port 1521 or any other available port is honored.

8. In the Configure ASM section, select **Configure ASM** if you want to configure ASM as part of this Deployment Procedure.

> **Note:** If you do not select **Configure ASM**, then the software binaries of Oracle Grid Infrastructure are copied to the destination host and a high-available listener is configured.
>
> If you want to configure Oracle ASM after provisioning, then you will have to manually configure it using Auomatic Storage Management Configuration Assistant (ASMCA).
>
> The high-available listener is configured in both cases, that is, when you select **Configure ASM** on this page and also when you manually configure Oracle ASM after provisioning.

9. In the Schedule section, retain the default selection, that is, **One Time (Immediately)** if you want to run the procedure immediately. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

   There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select End After and specify the end time, in hours and minutes, when it should terminate.

   > **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

10. Click **Next**.

    > **Note:** If you have Oracle ASM or Oracle Grid Infrastructure already installed on the destination host, then you are prompted to either deconfigure them or specify another set of destination hosts that do not have them.

11. If you chose to configure ASM, then on the Configure Oracle ASM page, add a new disk group and specify its credentials. Otherwise, go to Step (13).

    1. In the ASM Storage section, click **Add Disk Group** to add disk groups. From the Disk Group Configuration window that appears, you can either select one of the disks discovered by Enterprise Manager Grid Control or manually specify the disk you want to use.

    2. In the ASM Password section, specify the ASM password for SYS and ASMSNMP user accounts. Ensure that you meet the requirements described Section 6.4.2.1, "ASM Password Requirements".

12. Click **Next**.

13. On the Review page, do the following:

a. In the Procedure Instance Name section, for **Instance Name**, specify a unique custom name for the Deployment Procedure instance so that you can track its progress/status on the Procedure Completion Status page. If you do not specify a custom name, then a default name is computed and assigned to the procedure.

b. Review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Finish** to run the Deployment Procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Note: Click **Analyze Prerequisites** if you want to simulate the Deployment Procedure to check for prerequisites and to ensure that all the necessary requirements for provisioning are met. This helps you identify and resolve all operational conflicts and issues beforehand.

Note: if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

Note: If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.

If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

7

Provisioning Oracle Grid Infrastructure and Oracle RAC for Clustered Environments

This chapter explains how you can mass-deploy Oracle Grid Infrastructure and Oracle Real Application Cluster (Oracle RAC) for clustered environments in an unattended, repeatable, and reliable manner. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedure

- Supported Releases

- Cloning an Existing Oracle Grid Infrastructure and Oracle RAC

- Provisioning Oracle Grid Infrastructure and Oracle RAC Using Certified Gold Images

- Provisioning Oracle Grid Infrastructure and Oracle RAC Using Installation Binaries

> **Note:** Before you run the Deployment Procedure as outlined in this chapter to provision an Oracle Grid Infrastructure and Oracle RAC for a clustered environment, Oracle strongly recommends you to check the prerequisites to ensure that you meet all the necessary requirements for provisioning an Oracle Grid Infrastructure. To do so, simulate the Deployment Procedure as described in Chapter 4, "Simulating Database Provisioning Deployment Procedures to Check Prerequisites", and resolve all operational conflicts and issues beforehand.

> **Note:** To view an online demonstration of this feature, access the following URL:
>
> `http://www.oracle.com/technology/obe/demos/admin/demos.html`

7.1 Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Oracle Grid Infrastructure and Oracle RAC for clustered environments. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision Oracle

Grid Infrastructure and Oracle RAC. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 7–1   Getting Started with Provisioning Oracle Grid Infrastructure and Oracle RAC for Clustered Environments

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Grid Infrastructure and Oracle RAC for clustered environments. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 7.2, "Deployment Procedure".
Step 2	**Knowing About the Supported Releases** Know what releases of Oracle Grid Infrastructure and Oracle RAC can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 7.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle Grid Infrastructure and Oracle RAC. Select the use case that best matches your requirement.	■ To learn about cloning an existing Oracle Grid Infrastructure and and Oracle RAC, see Section 7.4, "Cloning an Existing Oracle Grid Infrastructure and Oracle RAC". ■ To learn about provisioning Oracle Grid Infrastructure and Oracle RAC using gold images, see Section 7.5, "Provisioning Oracle Grid Infrastructure and Oracle RAC Using Certified Gold Images". ■ To learn about provisioning Oracle Grid Infrastructure and and Oracle RAC using its software binaries, see Section 7.6, "Provisioning Oracle Grid Infrastructure and Oracle RAC Using Installation Binaries".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for cloning existing Oracle Grid Infrastructure and Oracle RAC, see Section 7.4.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Grid Infrastructure and Oracle RAC using gold images, seeSection 7.5.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Grid Infrastructure and Oracle RAC using their software binaries, see Section 7.6.1, "Prerequisites".

Table 7–1   (Cont.) Getting Started with Provisioning Oracle Grid Infrastructure and Oracle RAC for Clustered Environments

Step	Description	Reference Links
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Grid Infrastructure and Oracle RAC.	■ To clone an existing Oracle Grid Infrastructure and Oracle RAC, follow the steps explained in Section 7.4.2, "Cloning Procedure". ■ To provision Oracle Grid Infrastructure and Oracle RAC using gold images, follow the steps explained in Section 7.5.2, "Deployment Procedure". ■ To provision Oracle Grid Infrastructure and Oracle RAC using their software binaries, follow the steps explained in Section 7.6.2, "Deployment Procedure".

7.2 Deployment Procedure

Oracle RAC enables you to create a cluster of Oracle Databases so that they operate as a single entity. In order to create a cluster, Oracle RAC uses Oracle Clusterware, which is a part of Oracle Grid Infrastrcuture.

Oracle Grid Infrastrcuture provides the infrastructure for an enterprise grid architecture, and for a cluster, it offers Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

Enterprise Manager Grid Control provides the following Deployment Procedure for provisioning Oracle Grid Infrastructure and Oracle RAC for clustered environments:

■ *Oracle Grid Infrastructure / RAC Provisioning*

Using this Deployment Procedure, you can deploy Oracle Grid Infrastructure and Oracle RAC software using certified gold images, live reference installations, and installation media.

7.3 Supported Releases

Using this Deployment Procedure, you can provision the following release of Oracle Grid Infrastructure and Oracle RAC for clustered environments:

■ Oracle Grid Infrastructure 11g Release 2 (11.2)

■ Oracle Real Application Clusters 11g Release 2 (11.2)

7.4 Cloning an Existing Oracle Grid Infrastructure and Oracle RAC

This section describes how you can clone an existing Oracle Grid Infrastructure and Oracle RAC installation.

In particular, this section covers the following:

■ Prerequisites

■ Cloning Procedure

7.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
    - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

    To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- Meet the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation on the target hosts. For information about the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

- Details of the following groups should be specified in the Deployment Procedure. If these groups do not exist, then the Deployment Procedure automatically creates them. However, if these have to be created on NIS, then you must create them manually before running the Deployment Procedure. For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

    - ASM Database Administrator (OSDBA) group
    - ASM Instance Operator (OSOPER) group
    - ASM Instance Administrator (OSASM) group
    - Database Administrator (OSDBA) group
    - Database Operator (OSOPER) group

- If you want to store the cluster configuration details in a file and use that while running the deployment procedure, then create a cluster configuration file.

- If you want to use a response file to provision the database, then edit the `dbca.rsp` file available on the installation medium, and store it either in the Software Library or in a location accessible from the target hosts.

- If you want to use a custom template to create a structure for the database, then create a template (a `.dbt` file), and store it either in the Software Library or in a location accessible from the target hosts.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

> **Note:**   If you are using a locked account, then ensure that the SSH setup for the locked user is done prior to using the Deployment procedure.

- Ensure that you use an operating system user that has *write* permission on the following locations:
  - Staging areas where software binaries can be stored.
  - Oracle base directory for Grid Infrastructure where diagnostic data files related to Grid Infrastructure can be stored.
  - Oracle base directory for database where diagnostic data files related to database can be stored.
  - Grid Infrastructure software directory where Grid Infrastructure software can be provisioned.
  - Database software location where database software can be provisioned Working directory where cloning-related files can be staged.

- Ensure that the operating system groups you specify for the following groups already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.
  - ASM Database Administrator (OSDBA) group
  - ASM Instance Operator (OSOPER) group
  - ASM Instance Administrator (OSASM) group
  - Database Administrator (OSDBA) group
  - Database Operator (OSOPER) group

> **Note:**   When using Oracle ASM, you need to ensure that the RAC database user is a member of the ASM Database Administrator (OSDBA) group.

7.4.2 Cloning Procedure

To clone an existing Oracle Grid Infrastructure an Oracle RAC from a reference host, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select **Oracle Grid Infrastructure / RAC Provisioning** and click **Schedule Deployment**.

4. On the Select Provisioning Profile page, select **Use an existing Provisioning Profile** if you already have the configuration information captured in a file. Otherwise, select **Do Not Use Provisioning Profile**.

> **Note:** If you select a prosivioning profile, the configuration information captured in the file are automatically displayed on the screens of the deployment procedure. For example, information about the software installation locations as stored in the file are displayed on the Software Locations screen of the deployment procedure. You can choose to retain the default values as they are displayed or modify them to according to your preference. The instructions outlined in this procedure are written assuming a configuration component is not being used.

5. Click **Next**.

6. On the Select Source and Destination page, in the Select Source section, do the following:

   a. Select **Existing Installation**.

   b. For **Select Cluster Database**, click the torch icon and select the Oracle home of an existing cluster database you want to clone. Once the cluster database is selected, Enterprise Manager Grid Control displays the details of the cluster and its associated member nodes.

   c. In the Reference Host Options section, select the reference host you want to use for cloning the Oracle home, and validate the default location of the working directory where cloning-related files can be staged. You can either use the default location or specify another location of your choice.

7. In the Select Destination Hosts section, click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure and Oracle Database.

   If you have a cluster configuration file that lists the destination hosts, then click **Import from File** and select that file. After selecting the destination hosts, if you want to remove one or more hosts from the list, then select the host name and click the remove icon.

8. Click **Next**.

9. On the Specify OS Credentials and Schedule page, in the Reference Host Credentials section, specify the operating system credentials of the reference host you selected for cloning.

   a. Select **Use Preferred Credentials** if you want to use the preferred credentials stored in the Management Repository. Credentials here refer to operating system credentials.

   Otherwise, select **Override Preferred Credentials** if you want to override the preferred credentials and use another set of credentials for this operation.

   b. From the Host Credentials list, select **Same for all Oracle Homes** if you want to use the same path across hosts. Otherwise, select **Different for each Oracle Home** if you want to use different paths for each host.

> **Note:** If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> Also ensure that the deployment procedure has been customized to support locked environments. For information about customizing deployment procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

10. In the Target Hosts Credentials section, specify the operating system credentials of the destination hosts you selected for this deployment procedure. As described in the previous step, select the appropriate options. If you want to provision Oracle Grid Infrastructure using one user and Oracle Database using another user, then specify different credentials.

11. In the OS Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), ASM Instance Administrator (OSASM) group, Database Administrator (OSDBA) group, and Database Operator (OSOPER) group you want to use for operating system authentication. These groups must already exist.

> **Note:** This section appears only if the source you selected is on a UNIX platform.

12. In the Schedule section, if you want to run the procedure immediately, then retain the default selection, that is, **One Time (Immediately)**. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select **End After** and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

13. Click **Next**.

14. On the Specify Software Locations page, specify the locations where the software binaries of Oracle Grid Infrastructure and Oracle RAC can be placed, that is, the $ORACLE_HOME location. In particular, specify the following:

   a. For **Oracle Base for Grid Infrastructure**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with Oracle Automatic Storage Management (Oracle ASM) and Oracle Clusterware can be stored. This location is different from the Oracle home directory where the Grid Infrastructure software will be installed.

If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

The Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to `<user_home>/app/<user>`. For example, `/home/john/app/john`.

b. For **Software Location for Grid Infrastructure**, click the torch icon and select a location on the destination host where the Grid Infrastructure software can be provisioned. This is the Oracle home directory for Grid Infrastructure. Do not select a location that is a subdirectory of the Oracle Base for Grid Infrastructure or database.

Ensure that the directory path you provide meets the requirements described in Section 7.4.2.1, "Requirements for Grid Infrastructure Software Location Path".

c. For **Oracle Base for Database**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with the database can be stored. This location is used for storing only the dump files and is different from the Oracle home directory where the database software will be installed.

If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

d. For **Software Location for Database**, click the torch icon and select a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

Oracle recommends you to select a location that is a subdirectory of the Oracle Base for the database.

e. For **Working Directory**, click the torch icon and select a directory on the destination host where the files related to cloning can be staged temporarily. Ensure that you have approximately 7 GB of space for this directory.

f. For **Additional Installer Parameters**, specify any additional parameters you want to run while provisioning the database. You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. For example, `-debug`.

15. Click **Next**.

16. On the Configure Cluster page, specify a cluster name, a SCAN name, and a SCAN port.

> **Note:** The default SCAN port is port 1521, but you can specify another port of your choice. The deployment procedure verifies that the SCAN port provided is a valid port number, and is not used for any other purpose. After installation, a TNS listener listens to this port to respond to client connections to the SCAN name.

17. Select **Configure GNS and auto-assign with DHCP** and specify the GNS subdomain and VIP address if you want virtual host names outside the cluster to have dynamically assigned names.

18. (Ignore if you selected GNS in the previous step) In the Selected Hosts section, validate the public names and the virtual names of the hosts you selected.

19. In the Network Interface Configuration section, by default, the network interfaces that have the same name and subnet for the selected destination hosts are automatically detected and displayed. Validate these network interface configuration details.

From the **Usage** column, select **Public** to configure the interface as public interface, or **Private** to configure the interface as private interface. If you want to include additional network interfaces, then click **Add** Interface and specify the interface name, subnet, subnet mask, and its usage - public or private.

20. Click **Next**.

21. On the Configure Database page, select **Create Database** if you want to create a simple, starter database as part of this provisioning operation. Specify the global database name for the database and provide its credentials. The credentials are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

22. In the Configuration Type section, select **Admin-Managed** or **Policy-Managed** to indicate the way you want to configure and manage the database. If you select Policy-Managed, then the storage type for this deployment is forced to be ASM.

23. If you want to configure the starter database, then expand the Advanced Database Configuration Assistant (DBCA) Parameters section and provide the following information. Otherwise, go to the next step.

   a. Specify any additional DBCA parameters you want to pass while configuring the database. For example, – `characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

   b. In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you have a custom response file that already has the configuration options enabled.

   If you have stored the file as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select **Destination Hosts** and click the torch icon to select the location on the destination host where the file resides.

   > **Note:** From the Software Library or from the location where an existing database is running, only a .dbt template file can be used. However, from the destination host or a shared location, any template file can be used.

   c. In the Oracle Database Template File subsection, select **Use Custom Database Template File** if you want to use the structure of an existing database and have a custom template to structure the new database.

   If you have stored the template as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select

Destination Hosts and click the torch icon to select the location on the destination host where the file resides.

24. Click **Next**.

25. On the Select Storage Type page, select either **Automatic Storage Management** or **Shared File System** to indicate the storage type for storing voting disk and Oracle Cluster Registry (OCR). Voting disk and OCR are used by Oracle Clusterware to manage its resources. For more information about these storage options, see Overview of Storage Options.

Note: If you select Automatic Storage Management as the storage type for Grid Infrastructure, then the option Shared File System is disabled for Oracle RAC.

---

26. On the Assign Storage Devices page, assign storage devices for storing database files for Grid Infrastructure and Oracle RAC depending on the storage type you selected in the previous step. To understand the different storage type combinations and to know what information you need to provide for each combination, see Section 7.4.2.2, "Assigning Storage Devices"

27. Click **Next**.

28. On the Review page, do the following:

   a. Click **Run Prerequisites Only** if you want to simulate the deployment procedure to check whether the prerequisites are being met and whether the procedure will run successfully.

   b. Click **Save** and specify a provisioning profile name and a descriptoin if you want the configuration information to be stored in a file. The file gets created and uploaded to the Software Library. The file also appears for selection on the first page of the wizard when you run this deployment procedure next time.

   c. Review the details you have provided for the deployment procedure and if you are satisfied with the details, then click **Finish** to run the deployment procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Note: if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

7.4.2.1 Requirements for Grid Infrastructure Software Location Path

Meet the following requirements while specifying a directory path for the Oracle Grid Infrastructure home to deploy the Oracle Grid Infrastructure binaries:

- It should be created in a path outside existing Oracle homes

- It should not be located in a user home directory

- It should be created either as a subdirectory in a path where all files can be owned by root, or in a unique path

- Before installation, it should be owned by the installation owner of Oracle Grid Infrastructure (typically, `oracle`, for a single installation owner for all Oracle software, or grid for role-based Oracle installation owners), and set to 755 permissions

7.4.2.2 Assigning Storage Devices

The Assign Storage page of the Oracle Grid Infrastructure / RAC Provisioning deployment procedure enables you to assign storage devices for storing database files for Grid Infrastructure and Oracle RAC. However, the information about the storage devices you need to provide on this page depends on the storage type. The possible storage type combinations are:

- Automatic Storage Management for Grid Infrastructure and Oracle RAC

- Shared File System for Grid Infrastructure and Oracle RAC

- Shared File System for Grid Infrastructure and Automatic Storage System for Oracle RAC

The following sections describe what information you need to provide for each of these storage type combinations.

7.4.2.2.1 Automatic Storage Management for Grid Infrastructure and Oracle RAC

If you selected Automatic Storage Management as the storage type for Grid Infrastructure as well as Oracle RAC, then do the following:

- In the ASM Storage section, the table displays the disk groups that have been automatically discovered by Enterprise Manager Grid Control. In this table, do the following:

  - Review the details of these disk groups.

  - If you want to include additional disk group, then click **Add Disk Group**.

- Once you have reviewed the disk groups, select the OCR/voting disk and data files, and specify the ASM credentials.

- In the Device Configuration section, select **Clear Shared Storage** if you want to clear the shared storage, which is selected for the cluster and data storage, before creating and configuring the cluster.

7.4.2.2.2 Shared File System for Grid Infrastructure and Oracle RAC

If you selected Shared File System as the storage type for Grid Infrastructure as well as Oracle RAC, then do the following:

- Review the details of the shared storage locations provided in the table. If you want to include additional locations, then click **Add Storage Location** and select another shared storage location. You can search for NFS or OFCS2-based locations.

- In the Oracle Grid Infrastructure section, specify the storage location for Oracle Cluster Registry (OCR) and voting disks. Select **Normal** or **External** to indicate the redundancy level, and specify their locations.

- In the Real Application Cluster (RAC) section, specify the locations for storing the database files and recovering the backed up data.

- In the Device Configuration section, select **Clear Shared Storage** if you want to clear the shared storage, which is selected for the cluster and data storage, before creating and configuring the cluster.

> **Note:** The Deployment Procedure deletes the specified directories or files, that is, oradata directory or OCR file location, and not all the contents of specified storage device.

7.4.2.2.3 Shared File System for Grid Infrastructure and Automatic Storage System for Oracle RAC

If you selected Shared File System as the storage type for Grid Infrastructure and Automatic Storage Management for Oracle RAC, then do the following:

- Rreview the details of the shared storage locations provided in the table. If you want to include additional locations, then click **Add Storage Location** and select another shared storage location. You can search for NFS or OFCS2-based locations.

- In the Oracle Grid Infrastructure section, specify the storage location for Oracle Cluster Registry (OCR) and voting disks. Select **Normal** or **External** to indicate the redundancy level, and specify their locations.

- In the ASM Storage section, the table displays the disk groups that have been automatically discovered by Enterprise Manager Grid Control. In this table, do the following:

  - Review the details of these disk groups.

  - If you want to include additional disk group, then click **Add Disk Group**.

  - Once you have reviewed the disk groups, select the OCR/voting disk and data files, and specify the ASM credentials.

- In the Device Configuration section, select **Clear Shared Storage** if you want to clear the shared storage, which is selected for the cluster and data storage, before creating and configuring the cluster.

7.5 Provisioning Oracle Grid Infrastructure and Oracle RAC Using Certified Gold Images

In particular, this section covers the following:

- Prerequisites
- Deployment Procedure

7.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Meet the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation on the target hosts. For information about the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

- Ensure that the gold image is available either in the Software Library or in a staging location.

  To stage a gold image to the Software Library or to a shared location, see Appendix K, "Staging Gold Image to Software Library or Staging Location".

- Details of the following groups should be specified in the Deployment Procedure. If these groups do not exist, then the Deployment Procedure automatically creates them. However, if these have to be created on NIS, then you must create them manually before running the Deployment Procedure. For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

  - ASM Database Administrator (OSDBA) group

  - ASM Instance Operator (OSOPER) group

  - ASM Instance Administrator (OSASM) group

  - Database Administrator (OSDBA) group

  - Database Operator (OSOPER) group

- If you want to store the cluster configuration details in a file and use that while running the deployment procedure, then create a cluster configuration file.

- If you want to use a response file to provision the database, then edit the `dbca.rsp` file available on the installation medium, and store it either in the Software Library or in a location accessible from the target hosts.

  To understand how you can upload the response file to the Software Library, see Section 5.5.1.1, "Uploading Database Response File to Oracle Software Library".

- If you want to use a custom template to create a structure for the database, then create a template (a `.dbt` file), and store it either in the Software Library or in a location accessible from the target hosts.

To understand how you can upload the template to the Software Library, see Section 5.5.1.2, "Uploading Database Template File to Oracle Software Library".

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

  > **Note:** If you are using a locked account, then ensure that the SSH setup for the locked user is done prior to using the Deployment procedure.

- Ensure that you use an operating system user that has *write* permission on the following locations:
  - Staging areas where software binaries can be stored.
  - Oracle base directory for Grid Infrastructure where diagnostic data files related to Grid Infrastructure can be stored.
  - Oracle base directory for database where diagnostic data files related to database can be stored.
  - Grid Infrastructure software directory where Grid Infrastructure software can be provisioned.
  - Database software location where database software can be provisioned Working directory where cloning-related files can be staged.

- Ensure that the operating system groups you specify for the following groups already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.
  - ASM Database Administrator (OSDBA) group
  - ASM Instance Operator (OSOPER) group
  - ASM Instance Administrator (OSASM) group
  - Database Administrator (OSDBA) group
  - Database Operator (OSOPER) group

  > **Note:** When using Oracle ASM, you need to ensure that the RAC database user is a member of the ASM Database Administrator (OSDBA) group.

7.5.2 Deployment Procedure

To provision Oracle Grid Infrastructure and Oracle RAC using certified gold images, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select **Oracle Grid Infrastructure / RAC Provisioning** and click **Schedule Deployment**.

4. On the Select Provisioning Profile page, select **Use an existing Provisioning Profile** if you already have the configuration information captured in a file. Otherwise, select **Do Not Use Provisioning Profile**.

   > **Note:** If you select a prosivioning profile, the configuration information captured in the file are automatically displayed on the screens of the deployment procedure. For example, information about the software installation locations as stored in the file are displayed on the Software Locations screen of the deployment procedure. You can choose to retain the default values as they are displayed or modify them to according to your preference. The instructions outlined in this procedure are written assuming a configuration component is not being used.

5. Click **Next**.

6. On the Select Source and Destination page, in the Select Source section, do one of the following:

   a. If the Gold Image is stored as a component in Oracle Software Library (Software Library), then do the following:

      Select **Software Library**.

      Click the torch icon for **Grid Infrastructure Component** and select the generic component that has the Gold Image of the Grid Infrastructure.

      Click the torch icon for **RAC Component** and select the generic component that has the Gold Image of the RAC database.

      > **Note:** Ensure that you select only components that are in "Ready" status. When you select a component from the software library, Enterprise Manager Grid Control automatically populates the component location.

   b. If the Gold Image is stored as an image in a staging location, then do the following:

      Select **External Staging Server**.

      Click the torch icon for **Select Host** and select the host where the Gold Image is stored.

      Click the torch icon for **GI Stage Location** and select the location on the host where the Gold Image for Oracle Grid Infrastructure is available.

Click the torch icon for **RAC Stage Location** and select the location on the host where the Gold Image for Oracle RAC is available.

7. In the Select Destination Hosts section, click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure and Oracle Database.

If you have a cluster configuration file that lists the destination hosts, then click **Import from File** and select that file. After selecting the destination hosts, if you want to remove one or more hosts from the list, then select the host name and click the remove icon.

8. Click **Next**.

9. On the Specify OS Credentials and Schedule page, in the Target Hosts Credentials section, specify the operating system credentials of the destination host you selected.

   a. Select **Use Preferred Credentials** if you want to use the preferred credentials stored in the Management Repository. Credentials here refer to operating system credentials.

      Otherwise, select **Override Preferred Credentials** if you want to override the preferred credentials and use another set of credentials for this operation.

   b. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same path across hosts. Otherwise, select **Different for each Oracle Home** if you want to use different paths for each host.

      Note: If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

      Also ensure that the deployment procedure has been customized to support locked environments. For information about customizing deployment procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

      ---

10. In the OS Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), ASM Instance Administrator (OSASM) group, Database Administrator (OSDBA) group, and Database Operator (OSOPER) group you want to use for operating system authentication. These groups must already exist.

    Note: This section appears only if the source you selected is for a UNIX platform.

    ---

11. In the Schedule section, if you want to run the procedure immediately, then retain the default selection, that is, **One Time (Immediately)**. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

    There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the

deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select End After and specify the end time, in hours and minutes, when it should terminate.

> **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

12. Click **Next**.

13. On the Specify Software Locations page, specify the locations where the software binaries of Oracle Grid Infrastructure and Oracle RAC can be placed, that is, the $ORACLE_HOME location. In particular, specify the following:

   a. For **Oracle Base for Grid Infrastructure**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with Oracle Automatic Storage Management (Oracle ASM) and Oracle Clusterware can be stored. This location is different from the Oracle home directory where the Grid Infrastructure software will be installed.

   If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

   The Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to `<user_home>/app/<user>`. For example, `/home/john/app/john`.

   b. For **Software Location for Grid Infrastructure**, click the torch icon and select a location on the destination host where the Grid Infrastructure software can be provisioned. This is the Oracle home directory for Grid Infrastructure.

   Ensure that the directory path you provide meets the requirements described in Chapter 7.4.2.1, "Requirements for Grid Infrastructure Software Location Path".

   c. For **Oracle Base for Database**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with the database can be stored. This location is used for storing only the dump files and is different from the Oracle home directory where the database software will be installed.

   If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

   d. For **Software Location for Database**, click the torch icon and select a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

   e. For **Working Directory**, click the torch icon and select a directory on the destination host where the files related to cloning can be staged temporarily. Ensure that you have approvimately 7 GB of space for this directory.

     f. For **Additional Installer Parameters**, specify any additional parameters you want to run while provisioning the database. You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. For example, `-debug`.

14. Click **Next**.

15. On the Configure Cluster page, specify a cluster name, a SCAN name, and a SCAN port.

16. Select **Configure GNS and auto-assign with DHCP** and specify the GNS subdomain and VIP address if you want virtual host names outside the cluster to have dynamically assigned names.

17. (Ignore if you are using GNS) In the Selected Hosts section, validate the public names and the virtual names of the hosts you selected.

18. In the Network Interface Configuration section, by default, the network interfaces that have the same name and subnet for the selected destination hosts are displayed. Validate these network interface configuration details.

From the **Usage** column, select **Public** to configure the interface as public interface, or **Private** to configure the interface as private interface. If you want to include additional network interfaces, then click **Add** Interface and specify the interface name, subnet, subnet mask, and its usage - public or private.

19. Click **Next**.

20. On the Configure Database page, select **Create Database** if you want to create a simple, starter database as part of this provisioning operation. Specify the global database name for the database and provide its credentials. The credentials are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

21. In the Configuration Type section, select **Admin-Managed** or **Policy-Managed** to indicate the way you want to configure and manage the database. If you select **Policy-Managed**, then the storage type for this deployment is forced to be ASM.

22. If you want to configure the starter database, then expand the **Advanced Database Configuration Assistant (DBCA) Parameters** section and provide the following information. Otherwise, go to the next step.

     a. Specify any additional DBCA parameters you want to pass while configuring the database. For example, `– characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

     b. In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you have a custom response file that already has the configuration options enabled.

     If you have stored the file as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select **Destination Hosts** and click the torch icon to select the location on the destination host where the file resides.

> **Note:** From the Software Library or from the location where an existing database is running, only a .dbt template file can be used. However, from the destination host or a shared location, any template file can be used.

---

c. In the Oracle Database Template File subsection, select U**se Custom Database Template File** if you want to use the structure of an existing database and have a custom template to structure the new database.

If you have stored the template as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select **Destination Hosts** and click the torch icon to select the location on the destination host where the file resides.

23. Click **Next**.

24. On the Select Storage Type page, select either **Automatic Storage Management** or **Shared File System** to indicate the storage type for storing voting disk and Oracle Cluster Registry (OCR). Voting disk and OCR are used by Oracle Clusterware to manage its resources.

> **Note:** If you select Automatic Storage Management as the storage type for Grid Infrastructure, then the option Shared File System is disabled for Oracle RAC.

25. On the Assign Storage Devices page, assign storage devices for storing database files for Grid Infrastructure and Oracle RAC depending on the storage type you selected in the previous step. To understand the different storage type combinations and to know what information you need to provide for each combination, see Chapter 7.4.2.2, "Assigning Storage Devices"

26. Click **Next**.

27. On the Review page, do the following:

a. Click **Run Prerequisites Only** if you want to simulate the deployment procedure to check whether the prerequisites are being met and whether the procedure will run successfully.

b. Click **Save** and specify a provisioning profile name and a descriptoin if you want the configuration information to be stored in a file. The file gets created and uploaded to the Software Library. The file also appears for selection on the first page of the wizard when you run this deployment procedure next time.

c. Review the details you have provided for the deployment procedure and if you are satisfied with the details, then click **Finish** to run the deployment procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

> **Note:** If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.
>
> If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

7.6 Provisioning Oracle Grid Infrastructure and Oracle RAC Using Installation Binaries

In particular, this section covers the following:

- Prerequisites
- Deployment Procedure

7.6.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Meet the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation on the target hosts. For information about the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

- Ensure that the installation binaries are archived and uploaded either as a component in the Software Library or as an archived file in a staging location.

  To understand how you can download the installation binaries, and archive and uload them to the Software Library or to a shared location, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location"

- Details of the following groups should be specified in the Deployment Procedure. If these groups do not exist, then the Deployment Procedure automatically creates them. However, if these have to be created on NIS, then you must create them manually before running the Deployment Procedure. For information about creating these operating system groups, refer to the *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2)*.

  - ASM Database Administrator (OSDBA) group

  - ASM Instance Operator (OSOPER) group

- ASM Instance Administrator (OSASM) group

- Database Administrator (OSDBA) group

- Database Operator (OSOPER) group

- If you want to store the cluster configuration details in a file and use that while running the deployment procedure, then create a cluster configuration file.

- If you want to use a response file to provision the database, then edit the `dbca.rsp` file available on the installation medium, and store it either in the Software Library or in a location accessible from the target hosts.

  To understand how you can upload the response file to the Software Library, see Section 5.5.1.1, "Uploading Database Response File to Oracle Software Library".

- If you want to use a custom template to create a structure for the database, then create a template (a `.dbt` file), and store it either in the Software Library or in a location accessible from the target hosts.

  To understand how you can upload the template to the Software Library, see Section 5.5.1.2, "Uploading Database Template File to Oracle Software Library".

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

  > **Note:** If you are using a locked account, then ensure that the SSH setup for the locked user is done prior to using the Deployment procedure.

- Ensure that you use an operating system user that has *write* permission on the following locations:

  - Staging areas where software binaries can be stored.

  - Oracle base directory for Grid Infrastructure where diagnostic data files related to Grid Infrastructure can be stored.

  - Oracle base directory for database where diagnostic data files related to database can be stored.

  - Grid Infrastructure software directory where Grid Infrastructure software can be provisioned.

  - Database software location where database software can be provisioned Working directory where cloning-related files can be staged.

■ Ensure that the operating system groups you specify for the following groups already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.

– ASM Database Administrator (OSDBA) group

– ASM Instance Operator (OSOPER) group

– ASM Instance Administrator (OSASM) group

– Database Administrator (OSDBA) group

– Database Operator (OSOPER) group

> **Note:** When using Oracle ASM, you need to ensure that the RAC database user is a member of the ASM Database Administrator (OSDBA) group.

7.6.2 Deployment Procedure

To provision Oracle Grid Infrastructure and Oracle Real Application Cluster (RAC) using the archived software binaries, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select **Oracle Grid Infrastructure / RAC Provisioning** and click **Schedule Deployment**.

4. On the Select Provisioning Profile page, select **Use an existing Provisioning Profile** if you already have the configuration information captured in a file. Otherwise, select **Do Not Use Provisioning Profile**.

> **Note:** If you select a prosivioning profile, the configuration information captured in the file are automatically displayed on the screens of the deployment procedure. For example, information about the software installation locations as stored in the file are displayed on the Software Locations screen of the deployment procedure. You can choose to retain the default values as they are displayed or modify them to according to your preference. The instructions outlined in this procedure are written assuming a configuration component is not being used.

5. Click **Next**.

6. On the Select Source and Destination page, in the Select Source section, do the following:

   a. If the software binaries are stored as a component in Oracle Software Library (Software Library), then do the following:

   Select **Software Library**.

   Click the torch icon for **Grid Infrastructure Component** and select the generic component that has the archived software binaries of the Grid Infrastructure.

   Click the torch icon for **RAC Component** and select the generic component that has the archived software binaries of the database.

> **Note:** For information about downloading the installation binaries, and archiving and uploading them to the Software Library, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location". Ensure that you select only components that are in "Ready" status. When you select a component from the software library, Enterprise Manager Grid Control automatically populates the component location.

    b. If the software binaries are stored as an archived file in a staging location, then do the following:

    Select **Stage Location**.

    Click the torch icon for **Select Host** and select the host where the archived file is stored.

    Click the torch icon for **GI Stage Location** and select the location on the host where the archived file of the Grid Infrastructure software is available.

    Click the torch icon for **RAC Stage Location** and select the location on the host where the archived file of the database is available.

> **Note:** For information about downloading the installation binaries, and archiving and uploading them to a shared location, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location".

7. In the Select Destination Hosts section, click **Add** to search and select the destination hosts on which you want to provision Oracle Grid Infrastructure and Oracle RAC.

If you have a cluster configuration file that lists the destination hosts, then click **Import from File** and select that file. After selecting the destination hosts, if you want to remove one or more hosts from the list, then select the host name and click the remove icon.

8. Click **Next**.

9. On the Specify OS Credentials and Schedule page, in the Target Hosts Credentials section, specify the operating system credentials of the destination host you selected.

    a. Select **Use Preferred Credentials** if you want to use the preferred credentials stored in the Management Repository. Credentials here refer to operating system credentials.

    Otherwise, select **Override Preferred Credentials** if you want to override the preferred credentials and use another set of credentials for this operation.

    b. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same path across hosts. Otherwise, select **Different for each Oracle Home** if you want to use different paths for each host.

> **Note:** If this environment is secure and has locked accounts, then make sure that the credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> Also ensure that the deployment procedure has been customized to support locked environments. For information about customizing deployment procedures for locked accounts, see Chapter 31, "Customizing Deployment Procedures".

10. In the OS Group Details section, specify the ASM Database Administrator (OSDBA) group, ASM Instance Operator (OSOPER), ASM Instance Administrator (OSASM) group, Database Administrator (OSDBA) group, and Database Operator (OSOPER) group you want to use for operating system authentication. These groups must already exist.

    > **Note:** This section appears only if the source you selected is for a UNIX platform.

11. In the Schedule section, if you want to run the procedure immediately, then retain the default selection, that is, **One Time (Immediately)**. If you want to run the procedure later, then select **One Time (Later)** and provide time zone, start date, and start time details.

    There may be occasions when the Management Agent on that target host is down and as a result, Oracle Management Service is unable to start the deployment job. To handle such eventualities, you can specify a grace period. You can specify until when Oracle Management Service should continue to attempt to restart the deployment job if at all it stops for some reason. The default grace period is indefinite. If you want to have a defninte period, then select End After and specify the end time, in hours and minutes, when it should terminate.

    > **Note:** If the deployment job does not start within a specified time frame, then Oracle Management Service marks the job "skipped".

12. Click **Next**.

13. On the Specify Software Locations page, specify the locations where the software binaries of Oracle Grid Infrastructure and Oracle RAC can be placed, that is, the $ORACLE_HOME location. In particular, specify the following:

    a. For **Oracle Base for Grid Infrastructure**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with Oracle Automatic Storage Management (Oracle ASM) and Oracle Clusterware can be stored. This location is different from the Oracle home directory where the Grid Infrastructure software will be installed.

       If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

The Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to `<user_home>/app/<user>`. For example, `/home/john/app/john`.

b. For **Software Location for Grid Infrastructure**, click the torch icon and select a location on the destination host where the Grid Infrastructure software can be provisioned. This is the Oracle home directory for Grid Infrastructure.

Ensure that the directory path you provide meets the requirements described in Section 7.4.2.1, "Requirements for Grid Infrastructure Software Location Path".

c. For **Oracle Base for Database**, click the torch icon and select a location on the destination host where the diagnostic and administrative logs, and other logs associated with the database can be stored. This location is used for storing only the dump files and is different from the Oracle home directory where the database software will be installed.

If the user (operating system credentials) used for provisioning Oracle Grid Infrastructure is the same as the one used for provisioning the database, then you can specify a common Oracle base location for them. Otherwise, ensure that this Oracle base location is different from the Oracle base location for the database.

d. For **Software Location for Database**, click the torch icon and select a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

e. For **Working Directory**, click the torch icon and select a directory on the destination host where the files related to cloning can be staged temporarily. Ensure that you have approximately 7 GB of space for this directory.

f. For **Additional Installer Parameters**, specify any additional parameters you want to run while provisioning the database. You can specify any Oracle Universal Installer (OUI) parameter that can be used in this provisioning operation. For example, `-debug`.

14. Click **Next**.

15. On the Configure Cluster page, specify a cluster name, a SCAN name, and a SCAN port.

16. Select **Configure GNS and auto-assign with DHCP** and specify the GNS subdomain and VIP address if you want virtual host names outside the cluster to have dynamically assigned names.

17. (Ignore if you are using GNS) In the Selected Hosts section, validate the public names and the virtual names of the hosts you selected.

18. In the Network Interface Configuration section, by default, the network interfaces that have the same name and subnet for the selected destination hosts are displayed. Validate these network interface configuration details.

From the **Usage** column, select **Public** to configure the interface as public interface, or **Private** to configure the interface as private interface. If you want to include additional network interfaces, then click **Add Interface** and specify the interface name, subnet, subnet mask, and its usage - public or private.

19. Click **Next**.

20. On the Configure Database page, select **Create Database** if you want to create a simple, starter database as part of this provisioning operation. Specify the global

database name for the database and provide its credentials. The credentials are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

21. In the Configuration Type section, select **Admin-Managed** or **Policy-Managed** to indicate the way you want to configure and manage the database. If you select **Policy-Managed**, then the storage type for this deployment is forced to be ASM.

22. If you want to configure the starter database, then expand the **Advanced Database Configuration Assistant (DBCA) Parameters** section and provide the following information. Otherwise, go to the next step.

    a. Specify any additional DBCA parameters you want to pass while configuring the database. For example, `- characterSet AL32UTF8`. To view a list of DBCA parameters that can be specified here, see Appendix D, "Creating Templates for Oracle Database".

    b. In the Oracle Database Response File subsection, select **Use Custom Database Response File** if you have a custom response file that already has the configuration options enabled.

    If you have stored the file as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select **Destination Hosts** and click the torch icon to select the location on the destination host where the file resides.

    > **Note:** From the Software Library or from the location where an existing database is running, only a .dbt template file can be used. However, from the destination host or a shared location, any template file can be used.

    c. In the Oracle Database Template File subsection, select **Use Custom Database Template File** if you want to use the structure of an existing database and have a custom template to structure the new database.

    If you have stored the template as a component in the software library, then select **Software Library** and click the torch icon to select the component. If you have stored in it in a accessible location on the destination host, then select **Destination Hosts** and click the torch icon to select the location on the destination host where the file resides.

23. Click **Next**.

24. On the Select Storage Type page, select either **Automatic Storage Management** or **Shared File System** to indicate the storage type for storing voting disk and Oracle Cluster Registry (OCR). Voting disk and OCR are used by Oracle Clusterware to manage its resources.

    > **Note:** If you select Automatic Storage Management as the storage type for Grid Infrastructure, then the option Shared File System is disabled for Oracle RAC.

25. On the Assign Storage Devices page, assign storage devices for storing database files for Grid Infrastructure and Oracle RAC depending on the storage type you selected in the previous step. To understand the different storage type combinations and to know what information you need to provide for each combination, see Section 7.4.2.2, "Assigning Storage Devices".

26. Click **Next**.

27. On the Review page, do the following:

   a. Click **Run Prerequisites Only** if you want to simulate the deployment procedure to check whether the prerequisites are being met and whether the procedure will run successfully.

   b. Click **Save** and specify a provisioning profile name and a descriptoin if you want the configuration information to be stored in a file. The file gets created and uploaded to the Software Library. The file also appears for selection on the first page of the wizard when you run this deployment procedure next time.

   c. Review the details you have provided for the deployment procedure and if you are satisfied with the details, then click **Finish** to run the deployment procedure according to the schedule set. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Note: if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

Note: If Deployment Procedure failed at the prerequisites stage, then ensure that the latest prerequisites are downloaded from *My Oracle Support* and uploaded to the Software Library.

If your Enterprise Manager Grid Control has connectivity to *My Oracle Support*, then the latest prerequisites are automatically downloaded every night by a predefined job. However, if your Enterprise Manager Grid Control does not have any connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*, then manually download the prerequisite-specific ZIP files from *My Oracle Support* and upload them to the Software Library. For more information, see Appendix J, "Manually Downloading Prerequisites from My Oracle Support".

8

Extending Oracle Real Application Clusters

This chapter explains how you can extend and scale up an existing Oracle RAC stack (Oracle Clusterware, Oracle ASM, Oracle RAC database), in a single click. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedure
- Supported Releases
- Extending Oracle RAC

> **Note:** Before you run the Deployment Procedure as outlined in this chapter to extend an Oracle RAC, Oracle strongly recommends you to check the prerequisites to ensure that you meet all the necessary requirements for extending an Oracle RAC. To do so, simulate the Deployment Procedure as described in Chapter 4, "Simulating Database Provisioning Deployment Procedures to Check Prerequisites", and resolve all operational conflicts and issues beforehand.

8.1 Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in extending an existing Oracle RAC stack. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully extend an existing Oracle RAC. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 8–1   Getting Started with Extending Oracle RAC

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for extending Oracle RAC. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 8.2, "Deployment Procedure".

Table 8–1 (Cont.) Getting Started with Extending Oracle RAC

Step	Description	Reference Links
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle RAC can be extended by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 8.3, "Supported Releases".
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for extending Oracle RAC, see Section 8.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully extend an existing Oracle RAC.	To extend Oracle RAC, follow the steps explained in Section 8.4.2, "Procedure".

8.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for extending Oracle RAC:

- *One Click Extend Cluster Database*

8.3 Supported Releases

Using this Deployment Procedure, you can extend the following releases of Oracle RAC:

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

- Oracle Database 11g Release 2 (11.2.x.x)

8.4 Extending Oracle RAC

This section describes how you can extend an existing Oracle RAC to include as many additional nodes as you need, in just one click.

This section covers the following:

- Prerequisites

- Procedure

8.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers
- Ensure that you meet the prerequisites described in the following sections:

  – Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

- – Section 3.2, "Optional Infrastructure Requirements"

- Ensure that *oinstall* and *dba* groups are available.

- Ensure that operating system users such as *oracle* and *crsuser* are available on all nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.

- Ensure that the credentials being used to run this operation along with the group ID are the same on all nodes of the selected cluster.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that the shared storage used for existing cluster nodes are accessible to the nodes you want to add.

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

  Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

8.4.2 Procedure

To extend an existing Oracle RAC, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **One Click Extend Cluster Database**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Extend Real Application Clusters page.

4. On the Extend Real Application Clusters page, do the following:

a. In the Select Real Application Clusters (RAC) section, select the Oracle RAC you want to extend. The associated clusterware and Automatic Storage Management (ASM) also get extended if they are not already existing.

Expand All	Collapse All				
Select	Name	Member Nodes	Oracle Home	Platform	Product
	▼ Available Cluster Databases				
⦿	▼ pqtest (3)	strrac13.us.oracle.com, strrac15.us.oracle.com, strrac14.us.oracle.com	/opt/oracle /oracle11g/db	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle Database 11.1.0.7.0
	str13_cluster (3)	strrac13.us.oracle.com, strrac14.us.oracle.com, strrac15.us.oracle.com	/opt/crs11g	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle Clusterware 11.1.0.7.0
	+ASM1_strrac13.us.oracle.com (3)	strrac13.us.oracle.com, strrac15.us.oracle.com, strrac14.us.oracle.com	/opt/oracle /oracle11g/asm	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle ASM 11.1.0.7.0

You can use the **Search** section to search for a particular Oracle RAC. From the **Search** list, select the target type based on which you want to search, and click **Go**. You can use wildcards such as % and *.

Note: If the cluster database you want to extend does not appear on this page, then:

- Specify the clusterware home location for the cluster target in Enterprise Manager Grid Control. In Grid Control, click the **Targets** tab, and then click the **All Targets** subtab. On the All Targets page, from the Search list, select **Cluster** and click **Go**. From the results table, select the cluster for which you want to specify the clusterware home location. On the Cluster Home page, in the Related Links section, click **Monitoring Configuration**. On the Configure Target page, specify the clusterware home location and click **Update**.

- Configure Enterprise Manager Grid Control settings to display them. In Grid Control, click the **Deployments** tab, and on the Deployments page, in the Configuration section, click **Refresh Host Configuration**. On the following page, from the Available Hosts pane, select the hosts and add them to the **Selected Hosts** pane. Then, click **Refresh Hosts** and wait for the job to succeed. Return to this Deployment Procedure page and run the search query again to view the hosts in the search results.

b. In the Reference Host Options section, from the **Reference Host** list, select a host that you want to use as the primary host for performing this operation. Reference Host is the host that is selected for creation of clone archives and then transferred to the new target nodes being added.

▼ Reference host options - (strrac13.us.oracle.com)	
Reference host	strrac13.us.oracle.com ▼
	Select the reference host for extend.
Working directory	/tmp
	A working directory on the source host, is required to stage files for cloning. Supply the name of an existing directory on the host.
Home Location	**Files To Exclude**
/opt/crs11g	log/*,crs/init,racg/dump,srvm/log,cdata,EMStagedP
/opt/oracle/oracle11g/db	log/*,EMStagedPatches,oratab,*.trc,*.dbf,cdump,list
/opt/oracle/oracle11g/asm	log/*,EMStagedPatches,oratab,*.trc,*.dbf,cdump,list

For **Working directory**, specify the full path to an existing directory on the selected host that can be used for staging files for cloning. If the directories

you specify do not exist on the target hosts, then they will be created by the Deployment Procedure. Ensure that the working directory is NOT shared on the nodes.

For **Files To Exclude**, for each Oracle home, specify the files you want to exclude while performing this operation. Note that any file or folder corresponding to the regular expressions provided here will be excluded.

c. In the Oracle Home Shared Storage Options section, select the Oracle home locations that are on shared storage.

▼ Oracle Home Shared Storage Options	
Home Location	**Is Shared**
/opt/crs11g	☑
/opt/oracle/oracle11g/db	☐
/opt/oracle/oracle11g/asm	☐

d. In the Select New Nodes section, click **Add** to add new nodes that you want to include to the selected Oracle RAC. After adding the node, specify the virtual node name for each new node and verify the values displayed by default.

Select New Nodes
Select the destination hosts and enter the respective Virtual Node Names.

(Add...) (Import From File...)   ▶ Show Options
Host   **Virtual Node Name**   **Remove**
strrac16.us.oracle.com | strrac16-vip.us.oracle.com   🗑
                                              (Remove all)
✔ **TIP** Select hosts that are managed by agents of version 10.2.0.5.0

Note: Ensure that you select nodes that are monitored by Oracle Management Agents 10g Release 5 (10.2.0.5) or higher.

Optionally, you can click **Show Options** to specify Private Node Name, Private IP, Virtual IP, and Working Directory. **Private Node Name** and **Private IP** are required only if you want to set up a private network as part of the procedure.**Virtual Node Name** and **Virtual IP** are required only if they are fixed and not DHCP-based. If the node is already part of the Oracle Real Application Clusters (RAC) system, it will be ignored. If the node is part of the Oracle Clusterware, the private network and virtual host information will be ignored. For **Working Directory**, ensure that the location you specify is NOT shared on the nodes.

If you already have these details stored in cluster configuration file, then click **Import From File** to select that cluster configuration file. This file typically contains information about the new nodes to be added. It may also include information about the private node name, private IP address, virtual host name, and virtual IP address to which the Oracle RAC should be extended.

e. In the User Credentials (Override Preferred Credentials) section, retain the default selection, that is, **Use Preferred Credentials**.

> **Note:** You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.
>
> The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

From the **Host Credentials** list, select **Different for each Oracle Home** if you want to use different operating system credentials for each Oracle home, or **Same for all Oracle Homes** if you want to use the same set of credentials for all Oracle homes. Depending on the selection you make, specify the credentials. Ensure that the users belong to the same group (*dba/oinstall*).

> **Note:** If you are using vendor clusterware, then ensure that `root` and the operating system users, such as `oracle` and `crsuser`, owning the clusterware and various Oracle homes are a part of the operating system groups required by the vendor clusterware.
>
> For example, if your system uses High Availability Cluster Multiprocessing (HACMP) clusterware, then create or check for the existence of the group `hagsuser`. Ensure that the relevant operating system users and root user are members of this group.
>
> For more information, refer to the *Oracle Clusterware and Oracle Real Application Clusters Installation and Configuration Guide*.
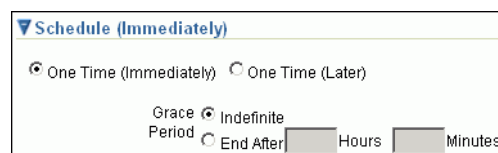
    f. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

    g. In the Prerequisites (Run Prerequisites and Fix-Ups) section, by default, **Skip prerequisites and fix-ups** is not selected and therefore, the deployment procedure runs the prerequisite checks and fix-ups on the selected nodes.

The prerequisite checks are required to ensure that the nodes meet all the requirements of this operation and are ready to be added to the cluster. The option is not selected assuming that you have not already run the prerequisite checks on the selected nodes beforehand.

If you have already run the prerequisite checks on the selected nodes and want the deployment procedure to skip running them all over again, then select **Skip prerequisites and fix-ups**.

If you have never run prerequisite checks on the selected nodes and if you want the deployment procedure to run them, then deselect **Skip prerequisites and fix-ups**. The deployment procedure runs the prerequisite checks, fixes issues if there are any, and then proceeds with the extend cluster operation.

If you want to check the prerequisites only but not proceed with the operation at this point, then click **Run Prerequisites Only**.

h. Click **Review**.

5. On the Review page, review the details you have provided for extending Oracle RAC, and click **Submit**.

Note: if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

Note: When you run the Deployment Procedure on Linux Itanium x64, if the *CVU Run to verify shared locations* step fails, then manually fix it before proceeding to the next step. No automated fix-ups are available for this platform.

9

Deleting or Scaling Down Oracle Real Application Clusters

This chapter describes how you can delete or scale down Oracle RAC. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedures
- Core Components Deleted
- Supported Releases
- Deleting the Entire Oracle RAC
- Scaling Down Oracle RAC by Deleting Some of Its Nodes

9.1 Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in deleting or scaling down an existing Oracle RAC stack. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully delete or scale down an existing Oracle RAC stack. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 9–1   Getting Started with Deleting or Scaling Down an Existing Oracle RAC

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for deleting or scaling down an existing Oracle RAC. Know how the Deployment Procedure functions, what use cases it covers, what core components it deletes or scales down, and so on.	To learn about the Deployment Procedure, see Section 9.2, "Deployment Procedures". To learn about the core components that are deleted or scaled down, see Section 9.3, "Core Components Deleted".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle RAC can be deleted or scaled down by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 9.4, "Supported Releases".

Table 9–1   (Cont.) Getting Started with Deleting or Scaling Down an Existing Oracle RAC

Step	Description	Reference Links
Step 3	**Selecting the Use Case** This chapter covers a few use cases for deleting and scaling down an existing Oracle RAC. Select the use case that best matches your requirement.	■ To learn about deleting the entire Oracle RAC stack, see Section 9.5, "Deleting the Entire Oracle RAC". ■ To learn about scaling down an existing Oracle RAC stack by deleting one or more of its own nodes, see Section 9.6, "Scaling Down Oracle RAC by Deleting Some of Its Nodes".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for deleting an entire Oracle RAC stack, see Section 9.5.1, "Prerequisites". ■ To learn about the prerequisites for scaling down an existing Oracle RAC stack by deleting one or more of its own nodes, see Section 9.6.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully delete or scale down an existing Oracle RAC.	■ To delete an entire Oracle RAC stack, follow the steps explained in Section 9.5.2, "Procedure". ■ To scale down an existing Oracle RAC stack by deleting one or more of its own nodes, follow the steps explained in Section 9.6.2, "Procedure".

9.2 Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for provisioning Oracle RAC:

■   *Delete/Scale down Oracle Real Application Clusters*

This Deployment Procedure enables you to desacle an entire cluster database stack (Oracle Clusterware, Oracle ASM, and Oracle RAC database) or one or more nodes of an Oracle RAC cluster, in a single click. This includes cluster databases of various releases as described in Section 9.4, "Supported Releases", and across different platforms.

The procedure can descale or delete clusters that have:

■   Oracle CRS, Oracle ASM, and Oracle Database homes owned by the same or different users.

■   Separate Oracle CRS, Oracle ASM, and Oracle Database homes present on a shared storage, which is shared by all member nodes.

■   Partially provisioned or failed installations of Oracle RAC clusters (may include one or more tiers of the stack installed. For example, cleanup after the clusterware installation failed or cleanup when the clusterware was only partially provisioned).

■   Nodes that were reimaged or shut down, and the existing configuration has to be resolved to remove all references to this node in Enterprise Manager Grid Control.

9.3 Core Components Deleted

Using this Deployment Procedure, you can delete either a node of an existing Oracle RAC or the entire Oracle RAC. As a result, the Deployment Procedure deinstalls the Oracle Clusterware, listeners, and Oracle RAC and ASM homes associated with the nodes selected for deletion.

9.4 Supported Releases

Using this Deployment Procedure, you can delete or scale down the following releases of Oracle RAC:

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

- Oracle Database 11g Release 2 (11.2.x.x)

9.5 Deleting the Entire Oracle RAC

This section describes how you can delete the entire Oracle RAC. In particular, this section covers the following:

- Prerequisites

- Procedure

9.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers
- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that *oinstall* and *dba* groups are available.

- Ensure that operating system users such as *oracle* and *crsuser* are available on all nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.

- Ensure that the credentials being used to run this operation along with the group ID are the same on all nodes of the selected cluster.

Prerequisites for Operators
- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment

Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

9.5.2 Procedure

To delete the entire Oracle RAC, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Delete/Scale down Oracle Real Application Clusters**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Delete/Scale down Oracle Real Application Clusters page.

4. On the Delete/Scale down Oracle Real Application Clusters page, do the following:

   a. In the Select Cluster section, click the torch icon for **Select Cluster** and select an Oracle Clusterware instance that you want to delete. Along with the selected Oracle Clusterware, the associated Oracle RAC and ASM instances will also be deleted. The table displays details about the member nodes that are part of the selected Oracle Clusterware.

Note:

When you use the torch icon to search for Oracle Clusterware, if you do not find the Oracle Clusterware that you are looking for, then from the tip mentioned below the table, click **here** to manually provide details about that clusterware and search for it.

This is particularly useful when you want to delete partially-provisioned or configured Oracle Clusterware instances because, by default, when you click the torch icon, only the fully-provisioned clusterware instances appear for selection. In this case, to search, select, and delete the partially-provisioned instances, click **here**, and in the Enter Cluster Details window, manually provide details about the cluster node that contains the partially-provisioned instance and click **OK**. You can then select the host that appears in the Select Nodes to Delete section and mark for deletion.

---

b. In the Reference Host Options section, from the **Cluster Node** list, select a node that you want to use as the primary node for all cleanup operations.

For **Working directory**, specify the full path to an existing directory on the selected node that can be used for staging files temporarily.

c. In the Select Nodes to Delete section, click **Mark all** to select all the nodes for deletion. On clicking **Mark all**, you should see a cross icon against all the nodes in the **Deletion** column. These cross icons indicate that the nodes have been selected for deletion.

d. In the User Credentials (Override Preferred Credentials) section, retain the default selection, that is, **Use Preferred Credentials**.

> **Note:** You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.
>
> For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.
>
> The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

From the **Host Credentials** list, select **Different for each Oracle Home** if you want to use different operating system credentials for each Oracle home, or **Same for all Oracle Homes** if you want to use the same set of credentials for all Oracle homes. Depending on the selection you make, specify the credentials. Ensure that the users belong to the same group (*dba/oinstall*).

e. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

f. Click **Review**.

5. On the Review page, review the details you have provided for deleting Oracle RAC, and click **Submit**.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

9.6 Scaling Down Oracle RAC by Deleting Some of Its Nodes

This section describes how you can scale down Oracle RAC by deleting one or more nodes that are part of it, and also other nodes that are part of it but do not appear as

targets in Enterprise Manager Grid Control. In particular, this section covers the following:

- Prerequisites

- Procedure

9.6.1 Prerequisites

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that *oinstall* and *dba* groups are available.

- Ensure that operating system users such as *oracle* and *crsuser* are available on all nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.

- Ensure that the credentials being used to run this operation along with the group ID are the same on all nodes of the selected cluster.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

9.6.2 Procedure

To scale down Oracle RAC by deleting one or more nodes that are part of it, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Delete/Scale down Oracle Real Application Clusters**. Then click **Schedule Deployment**.

Enterprise Manager Grid Control displays the Delete/Scale down Oracle Real Application Clusters page.

4. On the Delete/Scale down Oracle Real Application Clusters page, do the following:

   a. In the Select Cluster section, click the torch icon for **Select Cluster** and select an Oracle Clusterware instance that you want to scale down. Along with the selected Oracle Clusterware, the associated Oracle RAC and ASM instances will also be deleted. The table displays details about the member nodes that are part of the selected Oracle Clusterware.

   > **Note:** When you use the torch icon to search for Oracle Clusterware, if you do not find the Oracle Clusterware that you are looking for, then you can manually provide details about that clusterware and search for it. To do so, from the tip mentioned below the table, click **here**.

   b. In the Reference Host Options section, from the **Cluster Node** list, select a node that you want to use as the primary node for all cleanup operations.

   For **Working directory**, specify the full path to an existing directory on the selected node that can be used for staging files.

   c. In the Select Nodes to Delete section, select the nodes you want to delete, and click **Mark for delete**. On clicking **Mark for delete**, you should see a cross icon against the selected nodes in the **Deletion** column. These cross icons indicate that the nodes have been selected for deletion.

   If you do not see the nodes that are part of the cluster, then click **Add more nodes** to add those nodes so that nodes that do not appear as targets in Enterprise Manager Grid Control also are selected for deletion.

If you want to deselect a node, click **Unmark**. If you want to select all nodes at a time, click **Mark all**, and if you want to deselect all nodes, click **Unmark all**.

d. In the User Credentials (Override Preferred Credentials) section, retain the default selection, that is, **Use Preferred Credentials**.

Note: You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

From the **Host Credentials** list, select **Different for each Oracle Home** if you want to use different operating system credentials for each Oracle home, or **Same for all Oracle Homes** if you want to use the same set of credentials for all Oracle homes. Depending on the selection you make, specify the credentials. Ensure that the users belong to the same group (*dba/oinstall*).

e. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

f. Click **Review**.

5. On the Review page, review the details you have provided for deleting or scaling down Oracle RAC, and click **Submit**.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

10

Provisioning Oracle Database Replay Client

This chapter explains how you can provision Oracle Database Replay Client. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedures

- Supported Releases

- Cloning a Running Oracle Database Replay Client

- Provisioning Oracle Database Replay Client Using Gold Image

- Provisioning Oracle Database Replay Client Using Installation Binaries

10.1 Getting Started

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

Table 10–1   Getting Started with Provisioning Oracle Database Replay Client

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Database Replay Client. Know how the Deployment Procedure functions, what use cases it covers, what core components it deletes or scales down, and so on.	To learn about the Deployment Procedure, see Section 10.2, "Deployment Procedures".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle Database Replay Client can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 10.3, "Supported Releases".

Table 10–1 (Cont.) Getting Started with Provisioning Oracle Database Replay Client

Step	Description	Reference Links
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle Database Replay Client. Select the use case that best matches your requirement.	■ To learn about cloning an existing Oracle Database Replay Client, see Section 10.4, "Cloning a Running Oracle Database Replay Client". ■ To learn about provisioning Oracle Database Replay Client using a gold image, see Section 10.5, "Provisioning Oracle Database Replay Client Using Gold Image". ■ To learn about provisioning a standalone Oracle Database Replay Client, see Section 10.6, "Provisioning Oracle Database Replay Client Using Installation Binaries".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for cloning an existing Oracle Database Replay Client, see Section 10.4.1, "Prerequisites". ■ To learn about the prerequisites for provisioning Oracle Database Replay Client using a gold image, see Section 10.5.1, "Prerequisites". ■ To learn about the prerequisites for provisioning a standalone Oracle Database Replay Client, see Section 10.6.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Database Replay Client.	■ To clone an existing Oracle Database Replay Client, follow the steps explained in Section 10.4.2, "Cloning Procedure". ■ To provision Oracle Database Replay Client using a gold image, follow the steps explained in Section 10.5.2, "Provisioning Procedure". ■ To provision a standalone Oracle Database Replay Client, follow the steps explained in Section 10.6.2, "Provisioning Procedure".

10.2 Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for provisioning Oracle RAC:

■ *Oracle Database Replay Client Provisioning*

This Deployment Procedure is much like the Oracle Database Provisioning Deployment Procedure that helps you deploy a standalone Oracle Database across your enterprise configuration.

The Oracle Database Replay Client you want to provision may be available in different formats such as an installer in the installation medium, a running instance on a host monitored by Enterprise Manager Grid Control, or even a gold image in Oracle

Software Library (Software Library). Using this Deployment Procedure, you can provisioning any of these formats.

For example, if you want to have a fresh installation of Oracle Database Replay Client, the recommended option is to provision the database replay client from an installation medium. However, if you want to have a copy of a running instance that is stable and has all the latest patches applied, the recommended option is to clone that existing instance so that you retain the same configuration settings.

10.3 Supported Releases

Using these Deployment Procedures, you can provision the following releases of Oracle Database Replay Client:

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

- Oracle Database 11g Release 2 (11.2.x.x)

10.4 Cloning a Running Oracle Database Replay Client

This section describes how you can clone an existing Oracle Database Replay Client that is running on a host monitored by Enterprise Manager Grid Control.

This option is best suited when you have a running instance of Oracle Database Replay Client that is stable and has all the latest patches applied, and you want to make identical copies of it on multiple hosts. However, the risk involved in using an existing instance is that the instance may be deleted or deinstalled anytime without prior notice, and as a result, the Deployment Procedure may fail. Therefore, use this option when you know that the running instance is available for cloning.

In particular, this section covers the following:

- Prerequisites

- Cloning Procedure

10.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

    If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

    For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

    For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Database Replay Client.

    Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

    ```
    $ umask
    ```

    Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

10.4.2 Cloning Procedure

To clone an existing instance of Oracle Database Replay Client, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Replay Client Provisioning**. Then click **Schedule Deployment**.

    Enterprise Manager Grid Control displays the Select Source and Destination page of the Deployment Procedure.

4. On the Select Source and Destination page, do the following:

    a. In the Select Source section, select **Existing Database Replay Client Installation**. Then click the torch icon for **Source Host** and select the host on which the existing Oracle Database Replay Client is running.

    In the Source Host Details section, by default, **Oracle Home**, **Working Directory**, and **Files to exclude** are prefilled. **Oracle Home** shows where the existing instance is installed, but this is a non-editable field. For **Working**

Directory, specify the full path to a directory on source host where the files related to cloning can be staged temporarily. For **Files to exclude**, specify file names that must not be cloned to the source host. Use a comma to separate the file name, and use the wildcard (*) to indicate all files with the same extension. For example, *.trc. Note that any file or folder corresponding to the regular expressions provided here will be excluded.

In the Source Host Credentials section, select **Use Preferred Credentials** to use the credentials stored in the Management Repository.

b. In the Specify Destination Host Settings section, click **Add** and select the target hosts on which you want to clone the existing instance of Oracle Database Replay Client.

> **Note:** On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

> **Note:** You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.
>
> The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

> **Note:** If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced Installation Parameters section, specify any additional parameters you want to run while installing the Oracle Database Replay Client. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

While installing software binaries from an existing Oracle Database Replay Client location, if you want to also stage them to a shared location, then select **Stage to Shared Location** and specify a location that is shared across all destination hosts. This staged location can also act as a source for future deployments.

        d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

        e. Click **Continue**.

5. On the Review page, review the details you have provided for provisioning an Oracle Database Replay Client, and click **Submit**.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

10.5 Provisioning Oracle Database Replay Client Using Gold Image

This section describes how you can provision a gold image of Oracle Database Replay Client from the Software Library.

This option is best suited when you have a copy of a stable, well-tested, and patched Oracle Database Replay Client stored in the Software Library. This option scores over a fresh installation because you save time in patching and testing a fresh instance.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

10.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that the gold image is available either in the Software Library or in a shared, staging location.

  To stage a gold image to the Software Library or to a shared location, follow the instructions outlined in Appendix K, "Staging Gold Image to Software Library or Staging Location".

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Database Replay Client.

  Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

  Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

10.5.2 Provisioning Procedure

To provision a gold image of Oracle Database Replay Client from the software library, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Replay Client Provisioning**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source and Destination page of the Deployment Procedure.

4. On the Select Source and Destination page, do the following:

   a. In the Select Source section, do one of the following:

      If the gold image is stored as a component in the Software Library, then select **Software Library**. Then, click the torch icon for **Component** and select the component that has the gold image. Ensure that you select only components that are in "Ready" status.

      > **Note:** If you do not see the required component in the Software Library, then follow the workaround described in Appendix A, "Troubleshooting Issues".

If the gold image was stored as an image in a staging location while provisioning a database in the past, then select **External Staging Server** and then **Gold Image**. Click the torch icon for **Select Host** and select the host where the gold image is stored. Then click the torch icon for **Stage Location** and select the location on the host where the gold image is available. For Product version, specify the version of the product you are provisionoing.

b. In the Specify Destination Host Settings section, click **Add** and select the target hosts on which you want to install the gold image of Oracle Database Replay Client.

Note: On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

> **Note:** You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.
>
> The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

> **Note:** If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced Installation Parameters section, specify any additional parameters you want to run while installing the Oracle Database Replay Client. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

While installing the Oracle Database Replay Client location, if you want to also stage them to a shared location, then select **Stage to Shared Location** and specify a location that is shared across all destination hosts. This staged location can also act as a source for future deployments.

d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

> **e.** Click **Continue**.

5. On the Review page, review the details you have provided for provisioning an Oracle Database Replay Client, and click **Submit**.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

10.6 Provisioning Oracle Database Replay Client Using Installation Binaries

This section describes how you can provision Oracle Database Replay Client that is identical to the one available on the installation medium.

This option is best suited when you want a completely new installation to be provisioned across multiple hosts. Of course, understandably, this is a fresh installation and you will have to update it with all the latest patches that have been released so far.

> **Note:** The Oracle Database Replay Client version to be used for replaying workload must be the same version as the version of the test database on which the workload has to be replayed. Oracle Database Replay Client is supported in Oracle Database 10g Release 4 (10.2.0.4) and higher. While you can use archived software binaries for installing Oracle Database Replay Client of Oracle Database 11g Release 1 (11.1.0.6), for test database versions 10.2.0.4, 10.2.0.5 (future) and 11.1.0.7, you must create a gold image of the respective versions of Oracle Database Replay Client homes and use the same.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

10.6.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers
- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"

> **Note:** If you want to create a component for the software binaries of Oracle Database Replay Client, then before you access the Software Library, see My Oracle Support note 815567.1. This note explains the different requirements for each OS platform prior to using the media with Grid Control Deployment Procedure.

- Ensure that the installation binaries are downloaded, and archived and uploaded either as a component in the Software Library or as an archived file in a staging location.

  For information about downloading the installation binaries, and archiving and uploading them to the Software library or to a staging location, see Appendix L, "Downloading Installation Binaries and Uploading to Software Library or Staging Location".

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you use an operating system user that has *write* permission on the staging areas used for placing software binaries of Oracle Database Replay Client.

  Deployment Procedures allow you to use staging locations for quick file-transfer of binaries and prevent high traffic over the network. While providing a staging location, ensure that the operating system user you use has *write* permission on those staging locations. For information about staging area, see Section 2.11.3, "Staging Area".

- Ensure that the umask value on the target host is 022. To verify this, run the following command:

  ```
  $ umask
  ```

Depending on the shell you are using, you can also verify this value in `/etc/profile`, `/etc/bashrc`, or `/etc/csh.cshrc`.

10.6.2 Provisioning Procedure

To provision a fresh Oracle Database Replay Client, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Database Replay Client Provisioning**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source and Destination page of the Deployment Procedure.

4. On the Select Source and Destination page, do the following:

   a. In the Select Source section, do one of the following:

      If the software binaries are stored as a component in the Software Library, then select **Software Library**. Then, click the torch icon for **Component** and select the component that has the archived software binaries. Ensure that you select only components that are in "Ready" status. When you select a component from the Software Library, Enterprise Manager Grid Control automatically populates the component location.

      > **Note:** If you do not see the required component in the Software Library, then follow the workaround described in Appendix A, "Troubleshooting Issues".

      If the software binaries are stored as an archived file in a staging location, then select **External Staging Server** and then **Shiphome**. Click the torch icon for **Select Host** and select the host where the archived file is stored. Then click the torch icon for **Stage Location** and select the location on the host where the archived file is available. For Product version, specify the version of the product you are provisionoing.

   b. In the Specify Destination Host Settings section, click **Add** and select the target hosts on which you want to install the Oracle Database Replay Client.

      > **Note:** On clicking Add, a window appears with a list of suitable host. If you do not see your desired host, then select **Show All Hosts** and click **Go** to view all other hosts.

By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are prefilled with sample values. Edit them and specify values that match with your environment and standards. If the directories you specify do not exist on the target hosts, then they will be created by the Deployment Procedure.

From the **Credentials** list, retain the default selection, that is, **Preferred**, so that the preferred credentials stored in the Management Repository can be used. Credentials here refer to operating system credentials.

Note: You can optionally override these preferred credentials. For example, if you have added two destination hosts where the users are A and B, then you can choose to override the preferred credentials with different credentials for each of the hosts. Similarly, if the destinations hosts have same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with same credentials for all hosts.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

If you have selected multiple hosts, then from the **Path** list, select **Same for all hosts** if you want to use the same path across hosts, or select **Different for each host** if you want to use different paths for each host.

Note: If you select **Same for all hosts**, then ensure that the Oracle home and the user are present on all the hosts.

If you want to customize the host settings, then click **Customize Host Settings**. For example, you can specify the Management Agent home credentials, a name for your installation, or an alternate host name instead of the first hostname found on the system.

c. (Optional) In the Advanced Installation Parameters section, specify any additional parameters you want to run while installing the Oracle Database Replay Client. For example, `-force` (to override any warnings), `-debug` (to view more debug information), and `-invPtrLoc <Location>` (for UNIX only). Ensure that the parameters are separated by white space.

While installing software binaries from the Software Library, if you want to also stage them to a shared location, then select **Stage to Shared Location** and specify a location that is shared across all destination hosts. This staged location can also act as a source for future deployments.

d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

e. Click **Continue**.

5. On the Review page, review the details you have provided for provisioning an Oracle Database Replay Client, and click **Submit**.

> **Note:** if the Deployment Procedure fails, then review log files described in Appendix M, "Reviewing Log Files".

Part III

Provisioning Linux Operating System

This part describes how you can provision Linux operating system on bare metal servers and virtual servers.

This part contains the following chapters:

- Chapter 11, "Provisioning Linux Operating System"

1

Provisioning Linux Operating System

This chapter explains how you can provision Linux on bare metal servers and virtual servers using the bare metal application. In particular, this chapter covers the following. In particular, this chapter contains the following sections:

- Getting Started
- Supported Releases
- Provisioning Bare Metal Boxes
- Provisioning Oracle VM Servers
- Provisioning Single Server Images

> **Tip:** Before you begin provisioning of Linux on bare metal boxes, make sure that preferred credentials are set for the Stage server and Reference Host. If not, follow instructions in Section 3.7.8, "Setting Up Preferred Credentials for Bare Metal Provisioning in Enterprise Manager" to set up preferred credentials.

11.1 Getting Started

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 11–1   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.	To learn about the releases supported for Linux Provisioning, see Section 11.2, "Supported Releases".
Step 1	**Selecting the Use Case** This chapter covers a few use cases for provisioning Linux. Select the use case that best matches your requirement.	- To learn about provisioning bare metal boxes, see Section 11.3, "Provisioning Bare Metal Boxes". - To learn about provisioning Oracle VM Servers, see Section 11.4, "Provisioning Oracle VM Servers". - To learn about provisioning Single Server Images, see Section 11.5, "Provisioning Single Server Images".

Table 11–1 (Cont.) Getting Started with Provisioning Linux Operating System

Step	Description	Reference Links
Step 1	**Meeting the Prerequisites** Before you perform Linux provisioning, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites to be met for provisioning bare metal boxes, see Section 11.3.1, "Prerequisites". ■ To learn about the prerequisites to be met for provisioning Oracle VM Servers, see Section 11.4.1, "Prerequisites". ■ To learn about the prerequisites to be met for provisioning Single Server Images, see Section 11.5.1, "Prerequisites".
Step 1	**Provisioning** Provisioning Linux on bare metal boxes, Oracle VM Servers, or performing single image provisioning.	■ To provision Linux on bare metal boxes, follow the steps explained in Section 11.3.2, "Creating a Default Image", Section 11.3.3, "Setting Up a Default Image", and Section 11.3.4, "Staging the Default Image" in the same order. ■ To provision Oracle VM Servers, follow the steps explained in Section 11.4.2, "Provisioning Procedure". ■ To provision Single Server Images, follow the steps explained in Section 11.5.2, "Creating a Network Component", Section 11.5.3, "Creating Disk Layout Component", Section 11.5.4, "Creating OS Component", Section 11.5.5, "Creating Single Server Image", Section 11.5.6, "Creating Assignments", and Section 11.5.7, "Provisioning Assignments" in the same order.

11.2 Supported Releases

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

- Oracle Enterprise Linux 5.0 or higher

- Oracle Enterprise Linux 4.0 or higher

- RedHat Enterprise Linux (RHEL) 5.0 or higher

- RedHat Enterprise Linux (RHEL) 4.0 update 2 or higher

- RedHat Enterprise Linux (RHEL) 3.0 update 6 or higher

- SuSE Linux (SLES) 10

11.3 Provisioning Bare Metal Boxes

Provisioning only Linux Operating system on target servers involves creation of a **Default Image** and then deploying the image. The following sections explain how to provision Linux on bare metal boxes:

- Prerequisites

- Creating a Default Image

- Setting Up a Default Image

- Staging the Default Image

- Ready to Go

> **Note:** For information about downloading the Management Agent RPM kits, access the following URL:
>
> http://www.oracle.com/technology/software/products/oem/htdocs/provisioning_agent.html
>
> For instructions to install a Management Agent RPM kit, read the README file associated with the Management Agent RPM kit you are downloading.

11.3.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

11.3.2 Creating a Default Image

A **Default Image** is a special type of image that installs operating system based on a reference Linux installation. It then deploys an Oracle Management Agent on the bare metal box so that it can subsequently be managed as an Enterprise Manager target. A default image can be created for a particular IP prefix or MAC address and provisioned on bare metal boxes as they are added to a given subnet.

The Default Image setup consists of following steps:

1. Go to the **Images** tab and click **Create Image**. The following page is displayed:

In the **Describe** step, choose appropriate type of Default Image depending on the Operating System that needs to be provisioned on the bare metal servers. Enter a name for the image and other information that you may want to associate with it. Click **Next**.

2. The Configuration step comes up as shown below.

The **Configure** step depends on the type of the type of Image selected in the earlier step and is different for different types of Images. On this page, specify the following details:

Reference Installation is the reference host or gold host used for creating OS components for this image. For information about reference host, see Section 2.9.1, "Overview of the Bare Metal Provisioning Environment".

RPM repository allows you to choose an RPM repository from a list of RPM repositories configured for use with the provisioning application.

Additional RPM allows you to add RPMs for installation on the servers to be provisioned.

Exclude RPM allows you to specify the RPMs to be excluded from the image.

Time Zone and **UTC Clock** are used for setting the time zone of the bare-metal boxes after provisioning the operating system on them. Select the time zone corresponding to the OMS time zone.

Root Password is the root password that needs to be set for the newly provisioned Linux box.

User List is the list of operating system users you want to configure for the newly provisioned Linux box. You can create users with sudo permission.

Clicking **Next** brings up the **Customize** and **Set Directives** steps. For Default Image case, the out-of-the-box directives are automatically associated during creation.

Final step is the **Review** step where all the entries made in the earlier steps can be reviewed. Clicking **Finish** on the review page kicks off a job to create the necessary components and directives required for the Default Image.

Clicking on the job link will show the various steps executed by the job.

Once the job is finished a new Image can be seen in the Images table on the Images tab. This Image is created with version 0.1 is in **Ready** state for use. Any subsequent changes to the Image will lead to an increment of 0.1 in revision.

As a check, go to the Components tab to check that the components have been created automatically for the Default Image under the node "Oracle Components". It is recommended that you do not edit the components created for the Default Image.

> **Note:** If you want to remove an existing image from the Software Library, do the following:
>
> 1. Make sure the Image is not set as a default image in the Deployments->Administration->Default Images section. If it is, remove the entry from the table.
>
> 2. Remove the assignments (if any) for the image from the Deployment->Provisioning->Assignments table.
>
> 3. Remove the image from the Software Library by removing it from Deployment->Provisioning->Images table.
>
> 4. Remove the components associated with this image from the Software Library by removing it from the Deployments->Provisioning->Components->Oracle Components section.

11.3.3 Setting Up a Default Image

1. Go to the Administration tab and click **Add** in the Default Images section on the Administration tab as shown below.

2. The Add Default Image page appears.

Enter the following values on this page:

- **IP Address Prefix** or **MAC Address**:

  - **IP Address Prefix** is used to specify the IP prefix of subnet where the bare-metal machines would be added. For sample usecases, see Appendix G, "PXE Booting and Kickstart Technology".

  - **MAC Address** is used to specify the MAC addresses of the bare metal machines to be added.

- Choose a desired **Boot Server** and **Stage Server** from the list of servers configured with the provisioning application.

- Specify the **Image** name to be used as the Default Image, select **Use Preferred Credentials** for the Stage server.

  Specify the number of network cards for the target machine in the **Number of NICs** field.

- Click **Edit Advanced Properties**.

Administration > Add Default Image >
Edit Advanced Properties

Cancel OK

Agent Settings

Agent Install User: spudukol
Agent will be installed with this user

Agent Install Group: g900
Agent will be installed with this group permissions

Agent Base Installation Directory: /ade/spudukol_emcor
A directory called agent10g will be created under the base installation directory

Agent RPM location URL:
<file/http>://<hostname>/<complete path including rpm name>. Agent RPM to be installed will be picked from this URL location. If not specified, then Agent RPM will be picked from stage location

Oracle Central Inventory:
Enter a location like /u01/oraInventory under which the inventory will be created

Security Settings

Require tty: ☐
If unchecked, /etc/sudoers file is modified so that sudo does not require logging into a real terminal.

Security Parameter Settings:
```
#
# inittab       This file describes how the INIT process should set up
#               the system in a certain run-level.
#
# Author:       Miguel van Smoorenburg, <miquels@drinkel.nl.mugnet.org>
#               Modified for RHS Linux by Marc Ewing and Donnie Barnes
#
# Default runlevel. The runlevels used by RHS are:
#   0 - halt (Do NOT set initdefault to this)
#   1 - Single user mode
```
These entries will be added to /etc/inittab file

SELinux: ⦿ Disabled ○ Enforcing ○ Permissive
This will enable selinux

Firewall Settings

Firewall Settings:
```
firewall --disabled
```
These firewall setting will be applied on the target.

Kernel Settings

Kernel Parameter Settings:
```
#####
#
# File: /etc/sysctl.conf
# Created on: 08/05/2003
#
# Entries compiled from OTN recommedations and
# FST environment standards
#
################
# COMMON SECTION to OTN recommendations and FST env
################
```
These entries will be added to the /etc/sysctl.conf file

Disk Partition Settings

Disk partition commands:
```
zerombr yes
clearpart --all --initlabel
part / --fstype ext3 --size=3000 --grow
part swap --size=2048
```
These entries are picked up from the reference anaconda installation file available at /root/anaconda-ks.cfg on the reference host. You can modify these entries to suit the disk partition requirements of the target machines being provisioned.i Refer to the official Linux distribution documentation for details.

Boot Configuration Settings

Bootable Network Device: eth0
Enter the network device (name of the network interface card like eth0 or its MAC address or options like bootif and link). Refer to the official Linux distribution documentation for details.

Mount Point Settings

Mount Points Settings:

These entries will be added to the /etc/fstab file

NTP Settings

NTP Settings:
```
# Server Technologies
# Created 12/02/02 Jose Salas / Chris Trevino
# updated: <Date> <Name>
#
# Servers - routers are considered servers within Oracle.
#
# Systems Engineering routers that should be around for a long time
#
#
# Undisciplined Local Clock. This is a fake driver intended for backup
```
These entries will be added to the /etc/ntp.conf file

NIS Settings

NIS Settings:
```
domain us.oracle.com server 10.229.172.18    # dadvil0097.us.oracle.com
domain us.oracle.com server 10.229.79.146    # dadvig0137.us.oracle.com
domain us.oracle.com server 10.229.79.11     # dadvig0002.us.oracle.com
domain us.oracle.com server 10.229.172.59    # dadvil0138.us.oracle.com
```
These entries will be added to the /etc/yp.conf file

Post Section

Post Section Scripts:

These entries will be added to the post section of the Kick Start file

Cancel OK

You can provide the following properties:

■ In the **Agent Settings** section, specify the agent user, group, and Oracle home for the install, and the location of the agent RPM. If the agent RPM location is not specified, the RPMs will be picked up from the stage location. You can also specify the location where you want the Oracle Central Inventory to be created.

- In the **Security Settings** section, you can add a script in the Security Parameters field. You can enable or disable SELinux.

- In the **Firewall Settings** section, specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Enterprise Manager agent is open for its communication with the Management Service. By default, Enterprise Manager agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

- In the **Kernel Settings** section, you can add a script in the Kernel Parameters field.

- In the **Disk Partition Settings** section, you can specify the disk partition for new Linux machines. These entries are picked up from the reference anaconda installation file available at /root/anaconda-ks.cfg on the reference host. You can modify these entries to suit the disk partition requirements of the target machines being provisioned.

- In the **Boot Configuration Settings** section, specify the name of the network interface card such as eth0 or its MAC address or options like bootif and link.

- In the **Mount Point Settings** section, you can 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.

- In the **NTP Settings** section, you can 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.

- In the **NIS Settings** section, 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.

- In the **Post** section, 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.

Note: If setting up of the default image fails, delete the image, delete the component, and retry.

Running this job may take a few minutes.

Note: If the user has migrated default image components from an earlier version of Grid Control, say 10.2.0.5, to 11.1, the new properties for the default image will not be available.

11.3.4 Staging the Default Image

Once the Image is created, click **Stage** in the **Default Images** section and select the Image created earlier.

Default Images
Specify the Images to be provisioned to newly discovered Hardware Servers.

Select Default Image	Version	Description	IP Address Prefix	MAC Address	Status
○ Images/oel	0.1		1.1.1.1	N/A	Staged
○ Images/oel1	0.1		1.1.1.8	N/A	Staged
○ Images/oel2	0.1		1.1.1.9	N/A	Staged
● Images/DefaultImage1	0.1		10.177.226.5	N/A	Not Staged

This prepares the default image ready for installation on the staging server. You will be asked to enter the OMS registration password before staging starts.

An Enterprise Manager job is submitted to do the staging of default image. The on-going status can be seen using the Enterprise Manager job console.

> **Note:** If staging of the image fails, correct the pre-requisite failures mentioned in the failed job and retry the staging operation. Click **Edit** in the Default Images section and re-submit the staging job by clicking **Finish**.
>
> You can also clean up the image from the Default Images table and try again.
>
> Running this job may take a few minutes.

11.3.5 Ready to Go

Once the image is in the `Staged` status, you are ready to go. Plug in the bare metal boxes in the network and the provisioning application will automatically provision them to have the default image and management agent on it.

> **Note:** Once Linux is provisioned on the bare metal box, out-of-box Deployment Procedures can be used to provision Database and other Oracle products on the server.

11.4 Provisioning Oracle VM Servers

Enterprise Manager supports bare metal provisioning of Oracle VM servers.

11.4.1 Prerequisites

■ Ensure that you meet the prerequisites described in the following sections:

– Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

– Section 3.2, "Optional Infrastructure Requirements"

11.4.2 Provisioning Procedure

To provision a virtual server, do the following:

1. In the Images tab, click **Create Image**.

2. In the Create Image: Describe page, select Type as **OVS Default Image**, and specify a name and other details for the image. Click **Next**.

3. In the Create Image: Configure page, select the RPM repository.

Select the Time Zone and specify the root password for the operating system and the OVM agent password. The OVM agent password is used when registering the Virtual Server with the Server Pool.

In the User List section, you can add users for the operating system during the provisioning process.

> **Note:** Unlike bare metal provisioning of Linux boxes, when you provision virtual servers, you do not require a reference installation.

4. In the Create Image: Set Directives page, the directives required for OVS provisioning are selected by default. You can customize the provisioning operation by adding directives.

5. In the Create Image: Review page, review all the information you have provided and submit the job for creating the image.

6. After the image is created, in the Administration tab, go to the Default Images section.

7. In the Default Images section, click **Add**.

8. In the Add Default Image page, specify the target type as a MAC address or IP Address Prefix.

Specify the Boot Server, Staging Storage, and Number of NICs.

9. Select the OVS Image you created earlier.

10. Click **Edit Advanced Properties**.

11. In the Edit Advanced Properties page, specify advanced properties for the operating system.

In the **Security Settings** section, specify the following values:

- Request tty - Select this option to set the "Defaults require tty" property in the /etc/sudoers file.

- SELinux - Choose to disable or enable SELinux.

- Firewall Settings - Specify firewall settings for the virtual server target.

- Dom0 Memory (MB) - Specify memory size for the virtual server.

In the **Disk Partition Settings** section, specify the disk partition configuration to be applied to the target machines.

In the **Boot Configuration Settings** section, specify the network device name.

In the **Package** section, you can specify the packages to be installed during the provisioning operation.

In the **Xend Configuration** section, you can specify entries for the /etc/xen/xend-config.sxp file.

In the **Post** section, 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. You can specify Kernel, NFS, NTP, and other configurations in the Post section.

> **Note:** Unlike bare metal provisioning of Linux boxes, when you provision virtual servers, you will not configure the agent, since the virtual server will not explicitly contain an agent to monitor it.

12. Click **OK**.

13. Once the Image is created, click **Stage** in the **Default Images** section and select the Image created earlier.

14. Plugin the baremetal box in the subnet where the OVS image has been staged.

After the virtual server is provisioned, you can register it with a server pool. See Chapter 12, "Provisioning Virtualization Systems" for more information.

11.5 Provisioning Single Server Images

A Single server image contains an operating system and other applications. During provisioning, the operating system and other applications are installed on the machines being provisioned.

For provisioning single server images, follow these steps:

Table 11–2   Getting Started with Provisioning Single Server Images

Step	Description	Reference Links
Step 1	**Creating A Network Component** Create a network component to provide network details for the machines you are provisioning.	To learn about creating a network component, see Section 11.5.2, "Creating a Network Component".
Step 2	**Creating A Disk Layout Component** Create a disk layout component to provide profiles for hard disk, physical disk partition, RAID configuration, logical volumes, and logical volume groups for the machines you are provisioning.	To learn about creating a disk layout component, see Section 11.5.3, "Creating Disk Layout Component".
Step 3	**Creating An OS Component** Create an OS component to provide operating system details for the machines you are provisioning.	To learn about creating a disk layout component, see Section 11.5.4, "Creating OS Component".

Table 11–2   (Cont.) Getting Started with Provisioning Single Server Images

Step	Description	Reference Links
Step 4	**Creating A Single Server Image** Create a Single Server Image and add the disk layout and OS components created earlier.	To learn about creating a single server image, see Section 11.5.5, "Creating Single Server Image".
Step 5	**Creating An Assignment** Create an Assignment by associating the single-server image with the network component and any hardware component.	To learn about creating an assignment, see Section 11.5.6, "Creating Assignments".
Step 6	**Provisioning The Assignment** Provision the assignment containing single server images.	To learn about provisioning the assignment, see Section 11.5.7, "Provisioning Assignments".

11.5.1 Prerequisites

■   Ensure that you meet the prerequisites described in the following sections:

   –   Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

   –   Section 3.2, "Optional Infrastructure Requirements"

11.5.2 Creating a Network Component

The network component created in this section will provide the network details for the machines you provision using this component.

To create a network profile, do the following:

1.   In the Networks tab, click **Create Network**.

2.   In the Create Network: Describe page, specify the name, description, and product details for the network component. Click **Next**.

3. In the Create Network: Configure page, add the network configuration. Select:

   ■ Static if you want to specify the IP addresses

   ■ Dynamic if you want the DHCP server to assign a network address

   ■ Network Configuration if you want to assign network addresses from a network pool. Ensure that the network pool is configured.

4. Specify network information for bonding or non bonding interface. Click **Next**.

5. In the Create Network: Set Directives page, choose the out-of-box network directives for the network component for single server component creation. The directive will be available at `Oracle Directives/Baremetal Provisioning/Single Server/Linux/Network Layout`. Click **Next**.

6. Review the information you have provided and click **Finish**.

The network component is created.

11.5.3 Creating Disk Layout Component

Follow these steps to create a disk layout component:

1. In the Components tab, click **Create Component**.

2. In the Create Component: Describe page, select Type as **Linux Disk Layout**, and specify the name and other product details. Click **Next**.

3. In the Create Component: Configure page, add the profiles for hard disk, physical disk partition, RAID configuration, logical volumes, and logical volume groups. Click **Next**.

4. In the Create Component: Set Directives page, choose the out-of-box disk layout directives for the disk layout component for single server component creation. The directive will be available at `Oracle Directives/Baremetal Provisioning/Single Server/Linux/Disk Layout`. Click **Next**.

5. Review the information you have provided and click**Finish**.

The disk layout component is created.

11.5.4 Creating OS Component

Follow these steps to create an OS component:

1. In the Components tab, click **Create Component**.

2. In the Create Component: Describe page, select Type as **Linux Operating System**, (RedHat Linux, Oracle Enterprise Linux, or SuSE Linux), and specify the name and other product details. Click **Next**.

3. In the Create Component: Configure page, select the target host for the reference installation for the operating system, select the RPM repository, add or delete RPMs, specify the root password for the operating system and the list of operating system users you want to configure when the machine is provisioned. Click Next.

4. In the Create Component: Set Directives page, choose the out-of-box operating system directives for the operating system component for single server component creation. The directive will be available at `Oracle Directives/Baremetal Provisioning/Single Server/Linux/OS Staging`. Click **Next**.

5. Review the information you have provided and click**Finish**.

The operating system component is created.

11.5.5 Creating Single Server Image

Follow these steps to create a single server image:

1. In the Images tab, click **Create Images**.

2. In the Create Image: Describe page, select Type as **Single-Server** and specify the name and other product details. Click **Next**.

3. In the Create Image: Customize page, add properties if you want to customize the image. Click **Next**.

4. In the Create Image: Set Directives page, choose the out-of-box single server directives. The directives are available at `Oracle Directives/Baremetal Provisioning/Single Server/Linux/Image Staging`. Click **Next**.

5. In the Create Image: Add Components page, add the disk layout component and OS component created earlier. Click **Next**.

6. Review the information you have provided and click **Finish**.

The single server image is created.

11.5.6 Creating Assignments

Follow these steps to create an assignment:

1. In the Assignments tab, click **Create New Assignment**.

2. In the Create New Assignment: Select Type page, select **Create Image Assignment**. Click **Next**.

3. In the Create New Assignment: Select Image/Suite page, select the single server image you created earlier. Click **Next**.

4. In the Create New Assignment: Configure page, specify the name, boot server, staging storage, and staging server credentials. Select the network profile created earlier to assign network properties during provisioning.

    Specify Advanced Properties for the assignment. Select the hardware servers (Enterprise Manager host targets) you want to provision. Click **Next**.

5. Review the information you have provided and click **Finish**.

11.5.7 Provisioning Assignments

After creating the assignment, follow these steps to provision assignments:

1. In the Assignments tab, select the assignment and click **Provision**.

2. In the Provision Assignments page, specify a schedule, provide the agent password, and select the assignment. Click **Provision**.

The hardware specified will be provisioned after the reboot with the specified image assignment.

Part IV

Provisioning Virtualization Systems

This part describes how you can provision virtualization systems.

This part contains the following chapters:

- Chapter 12, "Provisioning Virtualization Systems"

12

Provisioning Virtualization Systems

Virtualization Management (VM) involves the monitoring, administration, and maintenance of virtual servers and guest virtual machines in your enterprise. Oracle provides the customer increased value by extending Enterprise Manager Grid Control capabilities to monitor virtualized systems alongside the physical infrastructure and perform complete lifecycle management of virtual servers and software running on them.

This chapter explains how you can provision Oracle VM. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedures

- Supported Releases

- Server Consolidation

- Provisioning Guest Virtual Machines

- Virtualization Central Dashboard

- Live Migrating Guest Virtual Machines

- High Availability and Fault Tolerance

- Advanced Operations

- Migrating from Oracle VM Manager to Enterprise Manager

12.1 Getting Started

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

Table 12–1   Getting Started with Oracle VM Provisioning

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedures** Understand the Deployment Procedures offered by Enterprise Manager Grid Control for provisioning Oracle VM.	To learn about the Deployment Procedures, see Section 12.2, "Deployment Procedures".

Table 12–1   (Cont.) Getting Started with Oracle VM Provisioning

Step	Description	Reference Links
Step 2	**Knowing About the Supported** Releases Know what releases are supported by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 12.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle VM. Select the use case that best matches your requirement	■ To learn about provisioning guest virtual machines, see Section 12.5, "Provisioning Guest Virtual Machines". ■ To learn about provisioning Oracle VM Servers, see Section 12.9.3, "Bare Metal Provisioning of Oracle VM Servers". ■ To learn about cloning guest virtual machines, see Section 12.9.1, "Cloning Guest Virtual Machines". ■ To learn about live migrating guest virtual machines, see Section 12.7, "Live Migrating Guest Virtual Machines".
Step 4	**Meeting the Prerequisites** Before you perform Oracle VM provisioning, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about prerequisites for provisioning guest virtual machines, see Section 12.5.2.4.1, "Prerequisites". ■ To learn about prerequisites for cloning guest virtual machines, see Section 12.9.1.1, "Prerequisites". ■ To learn about prerequisites for live migrating guest virtual machines, see Section 12.7.1, "Prerequisites".
Step 5	**Provisioning** Provisioning guest virtual machines, cloning guest virtual machines, or live migrating guest virtual machines.	■ To learn about provisioning guest virtual machines, see Section 12.5.3, "Creating Guest Virtual Machines". ■ To learn about cloning guest virtual machines, see Section 12.9.1.2, "Cloning Procedure". ■ To learn about live migrating guest virtual machines, see Section 12.7.2, "Live Migrating Selected Guest Virtual Machine" and Section 12.7.3, "Live Migrating All Guest Virtual Machines".
Step 6	**Configuring for High Availability** Configure Oracle VM for failover support.	To learn about high availability and fault tolerance, see Section 12.8, "High Availability and Fault Tolerance".

12.2 Deployment Procedures

Enterprise Manager Grid Control provides the following Deployment Procedures for provisioning Oracle VM:

■ *Virtual Machine Cloning*

This deployment procedure clones a guest virtual machine to one or more guest virtual machines

- *Virtual Machine Provisioning Using ISO*

  This deployment procedure creates a new guest virtual machines using an ISO image.

- *Virtual Machine Provisioning Using PXE*

  This deployment procedure creates a new guest virtual machines using Preboot Execution Environment. The new guest virtual machines will be PXE-booted and they will get provisioned via network.

- *Virtual Machine Provisioning Using Template*

  This deployment procedure creates new guest virtual machines using Oracle VM templates. The new virtual machines will be based on the Oracle VM template.

- *Edit Virtual Machine*

  This deployment procedure edits guest virtual machine configuration.

- *Live Migration*

  This deployment procedure performs live migration of a particular guest virtual machine or live migration of all guest virtual machines in a particular virtual server.

- *Save Virtual Machine as Template*

  This deployment procedure saves a guest virtual machine as template and exports the template to an external location, if provided.

- Import Template

  This deployment procedure imports an Oracle VM template into a server pool.

- *Import ISO Image*

  This deployment procedure imports an ISO Image into a server pool.

- *Create Shared Disk*

  This deployment procedure creates a shared disk image and registers it with the server pool.

- *Import Guest Virtual Machine (P2V)*

  This deployment procedure converts a physical machine to an Oracle virtual machine on a server pool.

- *Import Guest Virtual Machine (V2V)*

  This deployment procedure converts a virtual machine to an Oracle virtual machine on a server pool.

- *Import Template From P2V*

  This deployment procedure imports an Oracle VM template into a server pool by converting a Physical Machine(P2V).

- *Import Template From V2V*

  This deployment procedure imports an Oracle VM template into a server pool by converting a VMWare or VI Machine (V2V).

- *Discover Shared Disks*

  This deployment procedure discovers shared disks.

Deployment Procedures cannot be run directly. Deployment procedures for Oracle VM provisioning run when provisioning is performed through Virtualization Central.

12.3 Supported Releases

Table 12–2 shows the supported operating system versions for virtualization.

Table 12–2   Supported Releases

Operating System	Version
Linux	Oracle Enterprise Linux 4 - 32-bit and 64-bit
	Oracle Enterprise Linux 5 - 32-bit and 64-bit
	Red Hat Enterprise Linux 3 - 32-bit and 64-bit
	Red Hat Enterprise Linux 4 - 32-bit and 64-bit
	Red Hat Enterprise Linux 5 - 32-bit and 64-bit
Microsoft Windows	Microsoft Windows Server 2000
	Microsoft Windows Server 2003 Edition - 32-bit and 64-bit
	Microsoft Windows Vista
	Microsoft Windows XP
UNIX	Solaris 10

Table 12–3 shows the minimum agent version required for different Oracle VM Server versions.

Table 12–3   Minimum Agent Version for Different Oracle VM Server Versions

Oracle VM Server Version	Minimum Agent Version
2.1.2	2.2.70
2.1.5	2.2.86
2.2	2.3.37

12.4 Server Consolidation

One of the important goals in an enterprise network/cloud is to ensure faster dynamic resource provisioning for changing demands that invariably leads to frequent migrations of machines. Since migration is a basic operation in virtualization, Enterprise Manager provides the following functionalities:

- Conversion of physical machines to Oracle virtual machines
- Conversion of existing VMware/Virtual Iron virtual machines to Oracle virtual machines

> **Note:** This feature is available for the following conversion:
>
> - P2V as virtual machine (for Oracle VM version 2.1.2 and above)
> - VMWare V2V as virtual machine (for Oracle VM version 2.1.2 and above)
> - Virtual Iron V2V as virtual machine (for Oracle VM version 2.2 and above)

This section consists of the following subsections:

- P2V Conversion
- V2V Conversion

12.4.1 P2V Conversion

This section explains the steps to convert a Linux, Unix, or Windows computer to an Oracle VM guest virtual machine. This section contains the following subsections:

- Prerequisites
- Performing P2V Conversion

12.4.1.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

    - Ensure that the physical computer is booted in P2V mode using Oracle VM Server Installation Media, before initiating the P2V conversion process.

12.4.1.2 Performing P2V Conversion

Follow these steps to convert a Linux, Unix, or Windows computer to an Oracle VM guest virtual machine.

1. In Virtualization Central, select Convert P2V, select Action as Virtual Server Pool, and click **Go**.

2. In the Convert P2V page, specify the values in Table 12–4.

3. Click **OK**.

    The physical machine will be converted to an Oracle VM machine in the selected server pool.

    Note: If the physical machine had an agent installed before conversion, then after the conversion to an Oracle VM machine, the agent will lose its network identity, and therefore, will stop communicating with the OMS. To resolve this issue, you must reconfigure the agent. To do so, run the following commands from the Oracle home of the agent:

    1. Unsecure the agent:

        ```
        <ORACLE_HOME>/bin/emctl unsecure agent
        ```

    2. Invoke the agentca script:

        ```
        <ORACLE_HOME>/bin/agentca -f
        ```

    3. Secure the agent:

        ```
        <ORACLE_HOME>/bin/emctl secure agent
        ```

Table 12–4   Convert P2V

Element	Description
Server Pool	Specify the server pool where the guest virtual machine will be created. This field is prepopulated if you selected a server pool in Virtualization Central when invoking the P2V conversion action.
Web Server URL	Specify the Web Server URL in HTTPS format of the host booted in P2V mode, to access the disk images and configuration.
VM Name	Specify the name of the guest virtual machine to be created on the server pool.
OS Type	Specify the type of operating system.
OS Name	Specify the name of operating system.
VNC Console Password	Enterprise Manager Grid Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
Enable High Availability Mode	Select this option if you want the guest virtual machine to run automatically on a different virtual server if the earlier virtual server shuts down. This field can be set only if the server pool is enabled for high availability.
Start VM After Creation	Select this checkbox to start the guest virtual machine after conversion.
Preferred Server List	Select the virtual server on which you want the guest virtual machine to run. Choose **Automatic** to select the virtual server automatically or **Manual** to add a virtual server.

12.4.2 V2V Conversion

This section explains the procedure to convert a Virtual Iron or VMWare guest virtual machine to an Oracle VM guest virtual machine. This section contains the following subsections:

- Prerequisites
- Performing V2V Conversion

12.4.2.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

    - Ensure that the source machine is shut down and the binaries are available in an HTTP/FTP accessible location to start the conversion.

        For Virtual Iron, the location should be `http:///virtualMachineLocation/` or `ftp:///virtualMachineLocation/` where `virtualMachineLocation` contains the vm.cfg and the disk images for the virtual machine.

        For VMware, the location should be `http:///virtualMachineLocation/.vmx` or `ftp:///virtualMachineLocation/.vmx` where `virtualMachineLocation` contains the vmx configuration files and the disk images for the virtual machine.

12.4.2.2 Performing V2V Conversion

Follow these steps for V2V conversion:

1. In Virtualization Central, select Convert V2V, select Action as Virtual Server Pool, and click **Go**.

2. In the Convert V2V page, specify the values in Table 12–5.

3. Click **OK**.

   The source VMWare machine should be completely shut down when the conversion is in progress.

   After the process is completed, the specified virtual machine will be converted to an Oracle VM guest virtual machine.

Table 12–5   V2V Conversion

Element	Description
Server Pool	Specify the server pool where the guest virtual machine will be created. This field is prepopulated if you selected a server pool in Virtualization Central when invoking the P2V conversion action.
Virtual Machine Location	Specify the URL for accessing virtual machine.
	For VMWare virtual machines, specify the full http/ftp path to the configuration file (.vmx file). For example, ftp://host/vm/vm.vmx. For Virtual Iron virtual machines, specify the http/ftp path to the folder location, which contains the configuration and image files, for example, ftp://host/vm/
VM Name	Specify the name of the guest virtual machine being imported to the server pool.
OS Type	Specify the type of operating system.
OS Name	Specify the name of operating system.
VNC Console Password	Enterprise Manager Grid Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
Enable High Availability Mode	Select this option if you want the guest virtual machine to restart automatically on a different virtual server if the earlier virtual server shuts down. This field can be set only if the server pool is enabled for high availability.
Start VM After Creation	Select this checkbox to start the guest virtual machine after conversion.
Preferred Server List	Select the virtual server on which you want the guest virtual machine to run. Choose **Automatic** to select the virtual server automatically or **Manual** to add a virtual server.

12.5 Provisioning Guest Virtual Machines

Guest Virtual machines can be provisioned using Oracle VM templates, ISO images, or by PXE boot. Oracle VM templates and ISO images can be downloaded from OTN or other sources.

You can create guest virtual machines from the following three sources:

- Using Oracle VM Templates
- Using ISO images of Linux and Windows operating systems

- PXE bootable guest virtual machines that can boot over a network

Oracle VM template or ISO binaries can be available in any of the following locations:

- In the Software Library as a software component

- In the virtual server pool as a resource

- In a HTTP or NFS accessible external location

This section describes the following:

- Creating Software Library Components

- Managing Server Pool Resources

- Creating Guest Virtual Machines

- Configuring the Guest Virtual Machine for Agent and Other Software

12.5.1 Creating Software Library Components

You can create software components of type Oracle VM template and ISO images from the Software Components tab. These components can be later used to create Oracle guest virtual machines.

Enterprise Manager Grid Control allows you to create guest virtual machines from Oracle VM Templates and Linux and Windows ISO images. Oracle VM templates may contain operating system and pre-installed and pre-configured Oracle software. This section explains the following:

- Prerequisites

- Creating a Software Component from an Oracle VM Template

- Creating a Software Component from an ISO Image

12.5.1.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

12.5.1.2 Creating a Software Component from an Oracle VM Template

To create a software component from an Oracle VM template, do the following:

1. In the Software Components tab, select Components as Template and click **Create**.

2. In the Create Component: Describe page, select the Creation Method as Oracle VM Template Binary. Specify the name and description for the component, and operating system details.

3. In the Create Component: Configure page, select Virtualization Type as Para Virtualized or Hardware Virtualized based on the template you are trying to upload. Ensure that you select the virtualization type correctly.

   Select the Binary Upload Mechanism option to choose the option to store the template binary in the Software Library or specify an external location where the binary is stored.

   If you choose Upload from Agent Machine, select the host and the location of the binary.

If the binary is available at an external location, specify the HTTP or NFS location of the binary. For example, `http://myvm_server/template` or `file://myvm_server/template`. The binary will not be stored in the Software Library. It will be directly downloaded to the server pool when this resource is imported to a pool.

4. In the Create Component: Review page, review the information you have provided and submit the component creation job.

   After the job is submitted and has run successfully, the software component you have created will be displayed on the Software Components page.

12.5.1.3 Creating a Software Component from an ISO Image

To create a software component from an ISO image, do the following:

1. In the Software Components tab, select **Components** as ISO and click **Create**.

2. In the Create ISO page, specify the name and description of the ISO component and the operating system details.

   Select the Binary Upload Mechanism option to choose the option to store the template binary in the Software Library or specify an external location where the template binary is stored.

   If you choose Upload from Agent Machine, select the host and the location of the binary. Specify the host credentials to select the file from the agent machine.

   You can set up Preferred Credentials for hosts and virtual servers and use them or choose to Override Preferred Credentials when creating the ISO image.

If the binary is available at an external location, specify the HTTP or NFS location of the binary. For example, `http://myvm_server/iso` or `file://myvm_server/iso`.

The binary will not be stored in the Software Library. It will be directly downloaded to the server pool when this resource is imported to a server pool.

Click **OK**.

After the job is submitted and has run successfully, the software component you have created will be displayed on the Software Components page.

12.5.2 Managing Server Pool Resources

To create guest virtual machines, you can discover or import Oracle VM Templates and ISO binaries to the server pools and use them for provisioning. This section explains the following:

- Discovering Oracle VM Templates and ISO Images
- Creating Shared Disk
- Discovering Shared Disk

- Importing Oracle VM Templates and ISO Binaries as Server Pool Resources

12.5.2.1 Discovering Oracle VM Templates and ISO Images

You can store your Oracle VM templates in the following directories and add them as resources to the server pool using the Discover functionality:

- `/OVS/seed_pool`

- `/OVS/<repos_guid>/seed_pool`

`<repos_guid>` is the GUID of storage repository that has been added to the server pool.

You can store ISO images in the `/OVS/iso_pool` or /OVS/<repos_guid>/iso_pool directories and add them as resources to the server pool using the Discover functionality. Ensure that you create a directory with the same name as the ISO image in the `iso_pool` directory and then store the image. For example, `/OVS/iso_ pool/Win2008R2/Win2008R2.iso`.

To discover Oracle VM templates, do the following:

1. In Virtualization Central, select the Virtual Server Pool, select Action Virtual Server Pool - Manage Virtual Server Pool and click **Go**.

2. In the Oracle VM Server Pool: Resources page, select Resources as **Oracle VM Templates** and click **Discover**.

3. In the Discover GVM Template page, discover an existing Oracle VM template or VMWare/Virtual Iron guest virtual machine binaries. This template must be present in the shared location of the server pool. All discovered resources are available as Oracle VM templates in the server pool.

   - If the undiscovered template is of type Oracle VM, select the undiscovered Oracle VM template, and specify the description, operating system name and type, and the virtualization type of the template.

- If the undiscovered template is of type VMWare or Virtual Iron, select the undiscovered template, and specify the description, operating system name and type. The virtualization type for VMWare/Virtual Iron format templates will be hardware virtualized. Select **Delete source after conversion** to delete VMWare or Virtual Iron binaries after the discovery to free up space on the virtual server.

Note: The VMWare/VI Format import also involves a background conversion from VMware to Oracle VM template. Hence the time taken for VMware image import will be much more than the time taken by an Oracle VM template import.

The VMWare/Virtual Iron template discovery feature is available for Oracle VM version 2.2 or higher.

Click **OK**.

After clicking **OK**, the deployment procedure to discover the templates of type VMware or Virtual Iron will be submitted. The template you are discovering will appear in the Resources tab with the status `Creating`. Once the deployment

procedure completes successfully, the status will be `Ready`. In case the deployment procedure fails, the status will be `Error`. In case you have specified templates of type Oracle VM for discovery, they will be discovered without initiating any deployment procedure.

To discover ISO images, do the following:

1. In Virtualization Central, select the Virtual Server Pool, select Action Virtual Server Pool - Manage Virtual Server Pool and click **Go**.

2. In the Oracle VM Server Pool: Resources page, select Resources as ISO Images and click **Discover**.

3. In the Discover ISO Image page, select the undiscovered ISO image and provide a description. Select the OS Type and Operating System details. Click **OK**.

After clicking **OK**, the ISO image you are discovering will appear in the Resources tab with the status `Ready`.

12.5.2.2 Creating Shared Disk

You can also create shared disks that can be used by guest virtual machines and assign them at the time of guest virtual machine creation.

Oracle VM allows sharing of disks between multiple guest virtual machines. To create a shared disk that the guest virtual machine can access, follow these steps:

1. In Virtualization Central, select the Virtual Server Pool, and select Action Virtual Server Pool - Manage Virtual Server Pool and click **Go**.

2. In the Oracle VM Server Pool: Resources page, select Resources as Shared Disks and click **Create**.

3. In the Create Shared Disk page, specify the name, size of the disk, and provide a description. Click **OK**.

   After clicking **OK**, the deployment procedure to create shared disk will be submitted. The initial status will be `Creating`. Once the deployment procedure completes successfully, the status will be `Ready`. In case the deployment procedure fails, the status will be `Error`.

12.5.2.3 Discovering Shared Disk

Discover shared disk in the server pool shared location or using multipath device.

> **Note:** This feature is available for OVS version 2.2 and above.

12.5.2.3.1 Prerequisites

You must manually configure the multipath device on each virtual server in a server pool, then you can add the multipath device to the server pool as a shared virtual disk. The prerequisite steps to use a multipath device as a shared virtual disk are:

- Manually set up the SAN connection to the multipath device on each virtual server in the server pool.

- Manually configure the multipath device on each virtual server in the server pool.

- Make sure the multipath device has the same device identifier and device path on each virtual server in the server pool.

  For example, if you want to use a multipath device with the device identifier mpath-36090a028301ffd14fca5a48246976baa a sharable disk for the virtual machines in the server pool, it must have same device path name on each Virtual Machine Server in the server pool, for example, /dev/mpath/mpath1. Do not use the path /dev/mapper/x for multipath devices, only use the path /dev/mpath/x.

- Make sure the multipath device is not used for any other purpose (for example, mounted as an FTP server) to avoid any possible data corruption.

12.5.2.3.2 Procedure

Follow these steps to discover shared disk:

1. In Virtualization Central, select the Virtual Server Pool, and select Action Virtual Server Pool - Manage Virtual Server Pool and click **Go**.

2. In the Oracle VM Server Pool: Resources page, select Resources as Shared Disks and click **Discover**.

3. In the Discover Shared Disk page, select the shared disk type:

   - Select **File Backed** to discover virtual disk image file placed in the shareddisk directory of storage repositories. Select one or more image files and click **OK**.

   - Select **Multipath Backed** to discover multipath devices to be used as shared disks. Select one or more devices and click **OK**.

   After clicking **OK**, the deployment procedure to discover shared disk will be submitted. The initial status will be Creating. Once the deployment procedure completes successfully, the status will be Ready. In case the deployment procedure fails, the status will be Error.

12.5.2.4 Importing Oracle VM Templates and ISO Binaries as Server Pool Resources

- Prerequisites

- Importing Software Components from Software Components Tab

- Importing Oracle VM Templates from Server Pool Resources Tab

- Importing ISO Images from Server Pool Resources Tab

12.5.2.4.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:

  – Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

- Section 3.2, "Optional Infrastructure Requirements"
- Only a server pool administrator can import Oracle VM templates.

12.5.2.4.2 Importing Software Components from Software Components Tab

To import a software component, do the following:

1. In the Software Components tab, select Components as Template or ISO depending on the component you are importing, select the component, and click **Import**.

2. In the Import Template or Import ISO Image Page, specify the name of the imported resource, add the server pools to which you want to import the component, and specify monitoring server and virtual server credentials.

Import Template

* Name	Test_import_image
Component	AGENT_TEMPLATE_RC
Description	123
Virtualization Type	Paravirtualized
OS Type	Linux

Add Server Pools

(Delete) | (Add)

Select	Name
⊙	ca-emhostsPool

Monitoring Server Credentials

Specify the credentials to be used for provisioning. You can either choose to use the preferred credentials

Monitoring Server **stbdh19.us.oracle.com**

⊙ Use Preferred Credentials   ○ Override Preferred Credentials

Virtual Server Host Credentials

Specify the credentials to be used for provisioning. You can either choose to use the preferred credentials that the user you specify here must be allowed to use sudo on that virtual server host.

Virtual Server Host **ca-emhost3.us.oracle.com**

⊙ Use Preferred Credentials   ○ Override Preferred Credentials

3. Click **OK**.

After importing the software component, it will appear in the Server Pool Resources page as an imported resource.

12.5.2.4.3 Importing Oracle VM Templates from Server Pool Resources Tab

You can import an Oracle VM template to the server pool by using an existing Oracle VM template, or operating system and software on a physical computer, or guest virtual machine image file of VMWare/VI format.

To import an Oracle VM Template, do the following:

1. In Virtualization Central, select the Server Pool, select Action as Virtual Server Pool Manage Virtual Server Pool and click **Go**.

2. In the Resources tab, select Resources as Oracle VM Templates and click **Import**.

3. In the Import Template Page, select Mechanism as either Existing Binary, P2V Conversion, or V2V Conversion:

    ■ Select **Existing Binary** to import guest virtual machine template from existing binary. Specify a name and description for the Oracle VM template to be imported.

        – If you select the Binary Upload Mechanism as Associate Global Component, select the software component from the Software Library.

        If you specify Binary Upload Mechanism as External Location, provide the NFS or HTTP location of the Oracle VM template for the import, operating system name and type, and the virtualization type.

        Specify monitoring server and virtual server host credentials.

    ■ Select **P2V Conversion** to convert a Linux or Unix or Windows computer to an Oracle VM guest virtual machine. Specify the Oracle VM template name, description, and operating system name and type. Specify the Web Server URL in HTTPS format to access the disk images and configuration.

> **Note:** P2V conversion is supported for Oracle VM version 2.1.2 or higher.

- Select **V2V Conversion** to convert a Virtual Iron or VMWare guest virtual machine to an Oracle VM guest virtual machine. Specify the Oracle VM template name, description, and operating system name and type. Specify the URL for accessing the virtual machine. It can be either an FTP or HTTP location. For example,
`http://my-http-storage.domain.com/savedVMWareMachine/config.vmx` for VMWare guest virtual machine image files or
`http://my-http-storage.domain.com/savedVIIronMachine/config.cff` for Virtual Iron guest virtual machine image files.

> **Note:** V2V conversion is supported for Oracle VM version 2.2 or higher.
>
> The source VMWare machine should be completely shut down when the conversion is in progress.

4. Click **OK**. The deployment procedure to import the Oracle VM template is submitted and the template will appear with the status `Importing` in the Server Pool Resources page.

   After importing the Oracle VM template, it will appear in the Server Pool Resources page as an imported resource with the status `Ready`.

12.5.2.4.4 Importing ISO Images from Server Pool Resources Tab

To import an ISO image, do the following:

1. In Virtualization Central, select the Server Pool, select Action as Virtual Server Pool Manage Virtual Server Pool and click **Go**.

2. In the Resources tab, select Resources as ISO Images and click **Import**.

3. In the Import ISO Image Page, specify the name of the imported resource, select the component to be imported or specify the external location of the component, and specify monitoring server and virtual server credentials.

4. Click **OK**. The deployment procedure to import the ISO image is submitted and the template will appear with the status `Importing` in the Server Pool Resources page.

   After importing the ISO image, it will appear in the Server Pool Resources page as an imported resource with the status `Ready`.

12.5.3 Creating Guest Virtual Machines

You can create guest virtual machines from software components, imported resources on the server pool, or by providing an external location. You can also create PXE bootable guest virtual machines with no operating system on them. This section explains the following:

- Prerequisites
- Provisioning Guest Virtual Machines Using Oracle VM Templates
- Provisioning Guest Virtual Machines Using ISO Images
- Provisioning Guest Virtual Machines Using PXE Boot

12.5.3.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- The template for guest virtual machine creation must be present as a Software Library component, or at an external HTTP or NFS location, or imported to the server pool shared location.

12.5.3.2 Provisioning Guest Virtual Machines Using Oracle VM Templates

To provision a guest virtual machine using Oracle VM template, follow these steps:

1. Create a guest virtual machine template in any of the following ways:
   - Download the Oracle VM template to an external location (either HTTP or NFS location).

- Import the Oracle VM template to the server pool. For steps to import templates to server pool, see Section 12.5.2, "Managing Server Pool Resources".

- Use the Software Components tab in Virtualization Central to create a software component in the Software Library. For steps to create software components, see Creating Software Library Components.

2. In Virtualization Central, select Action as Guest VM, Create Guest VM and click **Go**.

3. In the Create Guest Virtual Machines: Server Pool page, specify the number of guest virtual machines to be created and the server pool with which you want to associate the guest virtual machines.

   'Specify the preferred virtual servers where the guest virtual machines will run. Choose **Automatic** to select the virtual server automatically. Click **Next**.

4. In the Create Guest Virtual Machines: Source page, select **Oracle VM Template**.

   – If the Oracle VM template is of type imported resources, select the template from the table.

   – If the Oracle VM template binary is stored in an external location, specify the NFS or HTTP location where the Oracle VM template is stored. For example, `http://myvm_server/template` or `file://myvm_server/template`. The template will be directly imported to the server pool without creating a

software library component. Specify the template name for the binary to be imported. Specify the operating system details and virtualization type.

– If the Oracle VM template is a software component stored in the Software Library, specify the resource name for the software component to be imported into the selected server pool and description. If the software component is not imported, it will be imported and the guest virtual machine will be created.

Click **Next**.

5. In the Create Guest Virtual Machines: Configure page, specify the configuration details shown in . Click **Configure Network** to provide network details. By default, the network configuration will be DHCP.

6. In the Create Guest Virtual Machines: Schedule/Credentials page, specify the schedule for creating the guest virtual machine and the monitoring server and virtual server host credentials.

7. In the Create Guest Virtual Machines: Review page, review the information you have provided and submit the job for guest virtual machine creation.

   After clicking **Finish**, the deployment procedure to create the guest virtual machine will be submitted. The initial status of the guest virtual machine(s) being created will be `Creating`. Once the guest virtual machine gets created, the status will be `Halted`. If the guest virtual machine is started successfully, then the status will be `Running` and it will appear under the virtual server on which it started in the Virtualization Central.

Table 12–6   Guest Virtual Machine Configuration Details for Oracle VM Templates

Element	Description
Virtual Machine Name Prefix	This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager Grid Control Console, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.
Number of Cores	When creating guest virtual machine from imported template resource, specify the number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (vCPU) configuration parameter for the guest virtual machine.

Table 12–6   (Cont.) Guest Virtual Machine Configuration Details for Oracle VM Templates

Element	Description
Memory Size (MB)	When creating guest virtual machine from imported template resource, specify the amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
Root Password/Confirm Root Password	When creating guest virtual machine from a Linux-based Oracle VM template, specify the operating system "root" password to log in to the guest virtual machine once it is up and running.
VNC Console Password	Enterprise Manager Grid Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
Number of Network Interfaces	Number of virtual network interfaces this guest virtual machine should have. The minimum value is 1 and the maximum value is 8. Specify the network details by clicking **Configure Network**. To configure the network, specify the following values: **Virtual Machine** - Select the guest virtual machine you are configuring. **Configuration Type** - Specify the type of configuration whether dynamic or static. **MAC Address** - Specify the MAC address of the network interface. This is not a mandatory field. If not specified, MAC address will be automatically generated. The format of MAC address is XX:XX:XX:XX:XX:XX, where the first three octets should be 00:16:3E. **Bridge** - Select a network bridge from the available bridges. **IP Address** - Specify the IP address to be assigned to the network interface. Specify the value in /ddd.ddd.ddd.ddd/ format. **Host Name** - Specify the Hostname associated with the specified IP address. For example, myhost1. The hostname specified for "VIF0" will be set as host name of the machine. For remaining interfaces, the hostname entries are made to /etc/hosts file. **Domain Name** - Specify the domain name associated with the specified IP address. **Subnet Mask** - Specify the mask used to determine what subnet the specified IP address belongs to. For example, 255.255.252.0 **Default Gateway** - Specify the IP address of the node (or router) on the network that serves as an access point to another network. The default gateway specified for "VIF0" is set as the default gateway of the machine. **DNS Server** - Specify the DNS server configured for the network.
CPU Scheduling Priority	Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
CPU Cap (%)	Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.

Table 12–6   (Cont.) Guest Virtual Machine Configuration Details for Oracle VM Templates

Element	Description
Enable High Availability Mode	When creating the guest virtual machine, if the "Enable High Availability Mode" flag is selected, this guest virtual machine will participate in the High Availability functionality. If its parent virtual server fails, this guest virtual machine will be auto-restarted on another Oracle VM server. Select this option if you want the guest virtual machine to restart automatically on a different virtual server if the earlier virtual server shuts down.
Local Devices from Source Template	Displays the local devices from source template. You can enable disk priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
Configure Agent	Select this option if you want to configure the agent to monitor the guest virtual machine.
Load Balancer Host	In case of multiple OMS setup, specify the load balancer host name to be used for agent configuration. By default, the agent will be configured with one of OMS hosts in a multiple OMS setup. If load balancer host is provided, then agent will be configured with the load balancer host.
Load Balancer Port	In case of multiple OMS setup, specify the load balancer port.
Management Server Registration Password	Password to register the agent.
Email	Specify the Email address for contact.
My Oracle Support Password	Password credentials for My Oracle Support.
Proxy Server	HTTP Proxy server used to connect to My Oracle Support.
Proxy Port	Proxy server port number
Proxy User name/Proxy Password	User name and password for the proxy server.
Start VM after creation	When creating guest virtual machine from template, specify whether the guest virtual machine should be started automatically after it is created.
Additional Disks	Every guest virtual machine is created with at least one disk. Specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log in to the guest virtual machine and reformat and mount the additional disks. Specify the name of the hard disk, its size, and QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

Table 12–6   (Cont.) Guest Virtual Machine Configuration Details for Oracle VM Templates

Element	Description
Shared Disks	Oracle VM allows sharing of disks between multiple guest virtual machines. You can specify multiple shared disks that the guest virtual machines can access.
	If the server pool has no shared disk available, the **Add** button will be disabled. Only virtual machines within the same server pool can share a shared disk.
	Specify the shared device name, the shared disk on it, and the QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

12.5.3.3 Provisioning Guest Virtual Machines Using ISO Images

To provision a guest virtual machine using ISO image, follow these steps:

1. Create an ISO image in any of the following ways:

    ■   Import the ISO image to the server pool. For steps to import ISO images to server pool, see Section 12.5.2, "Managing Server Pool Resources".

    ■   Use the Software Components tab in Virtualization Central to create a software component in the Software Library pointing to the ISO image in the shared storage.

    ■   Download the ISO image to an external HTTP or NFS location.

2. In Virtualization Central, select Action as Guest VM, Create Guest VM and click **Go**.

3. In the Create Guest Virtual Machines: Server Pool page, specify the number of guest virtual machines to be created and the server pool with which you want to associate the guest virtual machines.

    You can select the virtual servers for the guest virtual machines manually or allow the server pool to automatically allocate the virtual servers.

    Click **Next**.

4. In the Create Guest Virtual Machines: Source page, specify the ISO image source from which you want to create the guest virtual machine. Also specify the NFS Mount Point, that is, the NFS location where the ISO image will be mounted.

   – If the ISO image is of type imported resource, select Type as Imported Resource and then select the ISO image from the table.

   – If the ISO image is stored in an external location, specify the HTTP/NFS location where the ISO image is stored. For example, `http://myvm_server/iso` or `file://myvm_server/iso`. Specify the resource name using which the binary will be imported, OS name, and OS type. Also, choose the virtualization mode, either hardware or para, in which you want the virtual machine to run.

- If the ISO image is a software component stored in the Software Library, specify Name and Description. For para-virtualized type, specify the NFS Mount Point also.

5. In the Create Guest Virtual Machines: Configure page, specify the configuration details shown in Table 12–7.

6. In the Create Guest Virtual Machines: Schedule/Credentials page, specify the schedule for creating the guest virtual machine and the monitoring server and virtual server host credentials.

7. In the Create Guest Virtual Machines: Review page, review the information you have provided and submit the job for guest virtual machine creation.

Clicking **Finish** will submit the deployment procedure to create guest virtual machine(s). The initial status of the guest virtual machine will be 'Creating'. If the creation is successful, the status will be 'Running'.

> **Note:** For guest virtual machines created with Windows ISO images, the guest virtual machine will be shut down after creation. You must explicitly restart the guest virtual machine from Virtualization Central.

Table 12–7   Guest Virtual Machine Configuration Details for ISO Images and PXE Boot

Element	Description
Virtual Machine Name Prefix	This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager Grid Control Console, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.
Number of Cores	The number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (vCPU) configuration parameter for the guest virtual machine.
Memory Size (MB)	The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
Primary Disk Name	For guest virtual machine provisioning using ISO image or PXE boot, specify the name of the primary disk that will be used to create the guest virtual machine.
Primary Disk Size (MB)	For ISO and PXE boot, specify the size of the hard disk in megabytes. The disk size should be atleast 1024MB.
Primary Disk QoS	You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
VNC Console Password	Enterprise Manager Grid Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.

Table 12–7   (Cont.) Guest Virtual Machine Configuration Details for ISO Images and PXE

Element	Description
Number of Network Interfaces	Number of virtual network interfaces this guest virtual machine should have. The minimum value is 1 and the maximum value is 8.
	Specify the network details by clicking **Configure Network**. To configure the network, specify the following values:
	Virtual Machine - Select the guest virtual machine you are configuring.
	Configuration Type - Specify the type of configuration whether dynamic or static.
	MAC Address - Specify the MAC address of the network interface. This is not a mandatory field. If not specified, MAC address will be automatically generated. The format of MAC address is XX:XX:XX:XX:XX:XX, where the first three octets should be 00:16:3E.
	Bridge - Select a network bridge from the available bridges.
	IP Address - Specify the IP address to be assigned to the network interface. Specify the value in /ddd.ddd.ddd.ddd/ format.
	Host Name - Specify the Hostname associated with the specified IP address. For example, myhost1.
	The hostname specified for "VIF0" will be set as host name of the machine. For remaining interfaces, the hostname entries are made to /etc/hosts file.
	Domain Name - Specify the domain name associated with the specified IP address.
	Subnet Mask - Specify the mask used to determine what subnet the specified IP address belongs to. For example, 255.255.252.0
	Default Gateway - Specify the IP address of the node (or router) on the network that serves as an access point to another network. The default gateway specified for "VIF0" is set as the default gateway of the machine.
	DNS Server - Specify the DNS server configured for this network.
CPU Scheduling Priority	Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
CPU Cap (%)	Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
Enable High Availability Mode	When creating the guest virtual machine, if the "Enable High Availability Mode" flag is selected, this guest virtual machine will participate in the High Availability functionality. If its parent virtual server fails, this guest virtual machine will be auto-restarted on another Oracle VM server.
	Select this option if you want the guest virtual machine to restart automatically on a different virtual server if the earlier virtual server shuts down.

Table 12–7   (Cont.) Guest Virtual Machine Configuration Details for ISO Images and PXE

Element	Description
Additional Disks	Every guest virtual machine is created with at least one disk. Specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log in to the guest virtual machine and reformat and mount the additional disks.
	Specify the name of the hard disk, its size, and QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
Shared Disks	Oracle VM allows sharing of disks between multiple guest virtual machines. You can specify multiple shared disks that the guest virtual machines can access.
	If the server pool has no shared disk available, the **Add** button will be disabled. Only virtual machines within the same server pool can share a shared disk.
	Specify the shared device name, the shared disk on it, and the QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

12.5.3.4 Provisioning Guest Virtual Machines Using PXE Boot

To provision a guest virtual machine using PXE boot, follow these steps:

1. In Virtualization Central, select Action as Guest VM, Create Guest VM and click **Go**.

2. In the Create Guest Virtual Machines: Server Pool page, specify the number of guest virtual machines to be created and the server pool with which you want to associate the guest virtual machines.

   You can select the virtual server for the guest virtual machine manually or allow the server pool to automatically allocate the virtual server.

   Click **Next**.

3. In the Create Guest Virtual Machines: Source page, specify the PXE Boot source from which you want to create the guest virtual machine. Specify if you want hardware virtualized or para-virtualized type.

   Specify the MAC addresses for the guest virtual machines:

   – If you want Enterprise Manager Grid Control to assign a MAC addresses for the new guest virtual machines, select Automatic.

   – If you want to specify MAC addresses manually, select Manual and provide the MAC addresses in the text fields. The format of MAC address is XX:XX:XX:XX:XX:XX where the first three octets are 00:13:6E.

4. In the Create Guest Virtual Machines: Configure page, specify the configuration details shown in Table 12–7.

5. In the Create Guest Virtual Machines: Schedule/Credentials page, specify the schedule for creating the guest virtual machine and the monitoring server and virtual server host credentials.

6. In the Create Guest Virtual Machines: Review page, review the information you have provided and submit the job for guest virtual machine creation.

   Clicking Finish will submit the deployment procedure to create guest virtual machine(s). The initial status of guest virtual machine will be 'Creating'. If the creation is successful, the status will be 'Running'.

12.5.4 Configuring the Guest Virtual Machine for Agent and Other Software

To enable Enterprise Manager Grid Control to monitor your guest virtual machine, you will need to configure the agent.

Table 12–8 lists all the options available for configuring the agent when creating a guest virtual machine from Oracle VM template and the action you can take.

Table 12–8   Agent Configuration Options

Option	Result
Oracle VM template used for guest virtual machine creation has a Oracle Management Agent (Management Agent)	The fields to specify agent configuration details are displayed in the Create Guest Virtual Machine: Configure page.
- **Configure Agent** option is selected when creating the guest virtual machine	The agent is configured when the guest virtual machine is started for the first time. For details about configuring the agent, see Section 12.5.3, "Creating Guest Virtual Machines".

Table 12–8   (Cont.) Agent Configuration Options

Option	Result
- Oracle VM template used for guest virtual machine creation has a Management Agent - **Configure Agent** option is **not** selected when creating the guest virtual machine	The fields to specify agent configuration details are not displayed in the Create Guest Virtual Machine: Configure page. The agent is not configured during guest virtual machine creation. You will need to manually log in to the guest virtual machine using the VNC Console and configure the agent. To configure the agent, see *Oracle Enterprise Manager Grid Control Installation Guide*.
- Oracle VM template used for guest virtual machine creation does not have a Management Agent - **Configure Agent** option is selected when creating the guest virtual machine	Even though the **Configure Agent** option is selected, when creating the guest virtual machine, the agent will not be configured. You will need to manually log in to the guest virtual machine using the VNC Console and install and configure the agent. To install and configure the agent, see *Oracle Enterprise Manager Grid Control Installation Guide*.
- Oracle VM template used for guest virtual machine creation does not have a Management Agent - **Configure Agent** option is **not** selected when creating the guest virtual machine	When creating the guest virtual machine, the agent will not be configured. You will need to manually log in to the guest virtual machine using the VNC Console and configure the agent. To install and configure the agent, see *Oracle Enterprise Manager Grid Control Installation Guide*.

12.6 Virtualization Central Dashboard

The Enterprise Manager Virtualization Central page allows you to monitor Oracle VM, guest virtual machines, and manage resources. This section covers the following:

- Overview
- Editing Guest Virtual Machines
- Updating Preferred Virtual Server List
- Opening a VNC Console to the Guest Virtual Machine
- Deleting Server Pools, Oracle VM Servers, and Guest Virtual Machines

12.6.1 Overview

The Enterprise Manager Virtualization Central page allows you to monitor Oracle VM servers and guest virtual machines, and manage resources. Resource management includes packaging and management of various resources like ISO images, Oracle VM templates, and shared disk.

Managed resources can be imported or associated with a server pool. Resources associated with server pools can be chosen for virtualization life cycle management operations such as creating, cloning, and live migrating guest virtual machines.

Enterprise Manager Grid Control provides support for packaging and associating Oracle VM template and ISO component with server pools.

In addition, the Virtualization Central page also provides the following information:

- A quick view of the status of the virtual server and guest virtual machines and basic information about them

- CPU and Memory utilization of virtual servers

- Information about the server pool to which the virtual servers are registered

- Drilldowns to server pools, virtual servers, and guest virtual machines home pages from where you can view more information about these targets

- Administrative tasks such as starting and stopping guest virtual machines

- Provisioning activities such as creating virtual server pools, registering virtual servers, and creating guest virtual machines

- VNC Support to access guest virtual machines

Starting from the Virtualization Central page, you can access additional details and management functions by drilling down through the user interface. In addition, the Virtualization Central page provides a list of Related Links. From these links you can perform activities such as bare metal provisioning of Oracle VM server software, customization of table columns, and deployment of Management Agent on a guest virtual machine. The following figure shows the Virtualization Central Page.

For more information about administration and monitoring of Oracle VM, see *Enterprise Manager Concepts Guide* available at:

http://www.oracle.com/technology/documentation/oem.html

> **Note:** To support the monitoring and management of Oracle VM 2.2 Server in Grid Control 11g Release 1, when using a monitoring server of version 10.2.0.5, perform following manual configuration:
>
> 1. As a root user, create the following symlink:
>
>    ```
>    ln -s /opt/ovs-agent-2.2 /opt/ovs-agent-2.3
>    ```
>
> 2. If the `/opt/ovs-agent-2.2` directory already exists, back up the directory to another location first.
>
> This step is required before registering Oracle VM Server version 2.2 with Enterprise Manager during new virtual server pool creation or when registering a virtual server with virtual server pool, and also when upgrading a registered virtual server to version 2.2.

For information about the privileges required for administrative operations, see Section 2.8.2, "Privileges Required for Administrative and Provisioning Operations".

12.6.2 Editing Guest Virtual Machines

Enterprise Manager Grid Control allows you to modify the virtual CPU count, allocated memory, and other configuration of the guest virtual machine.

The Edit Guest Virtual Machine page has four tabs - General, Disk Storage, Network Interface, and Advanced:

- The General tab displays details such as guest virtual machine name, virtual server name, memory details, CPU scheduling and priority and other details.

- The Disk Storage tab displays disk storage details for the guest virtual machine.

- The Network Interface tab displays network configuration details for the guest virtual machine.

- The Advanced Configuration tab displays the boot order details.

Table 12–9 shows the editable fields for the guest virtual machine and the time when the changes will take effect.

Table 12–9   Guest Virtual Machines: Editable Fields

Editable Field	Action for Changes to Take Effect
Virtual Machine Name	This field can be edited only for guest virtual machines in the Halted status. Changes will take effect when the guest virtual machine is started.
Memory	This field can be edited only for guest virtual machines in the Halted status. Changes will take effect immediately when the guest virtual machine is started.
Dynamic Memory	Dynamic memory specified should be lesser than or equal to the Memory specified. For para-virtualized guest virtual machines, changes to the memory take effect without stopping and starting the guest virtual machine. For hardware-virtualized guest virtual machines, changes to memory take effect only after stopping and starting the guest virtual machine.

Table 12–9   (Cont.) Guest Virtual Machines: Editable Fields

Editable Field	Action for Changes to Take Effect
Virtual CPUs	For para-virtualized guest virtual machines, number of virtual CPUs can be edited and changes take effect without stopping and starting the guest virtual machine. For hardware-virtualized guest virtual machines, changes to virtual CPUs take affect only after stopping and starting the guest virtual machine.
	For Oracle VM of version 2.2 onwards, the number of virtual CPUs cannot be increased beyond the value when the virtual machine was last started.
VNC Console Password	New password takes effect after the guest virtual machine is stopped and restarted.
CPU Scheduling Priority	Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
CPU Cap (%)	Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
Enable High Availability Mode	If High Availability was enabled during guest virtual machine creation, you can edit the guest virtual machine to enable or disable High Availability.
Add/Remove Disk Storage	Changes take place immediately. CDs cannot be attached with para-virtualized guest machines.
QoS Priority Class	You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
Add/Remove/Edit Network Interface	If a guest VM is running, all the interfaces that are present on it since the last restart can be removed, but not edited. But the newly added interfaces can be edited or removed. Changes take place immediately. The primary network interface VIF0 cannot be removed.
Network Rate	Allows you to control the bandwidth granted to each virtual network interface in the running guest virtual machine.
Boot Order	Requires the guest virtual machine to be stopped and restarted. Boot Order cannot be set to CD for para-virtualized guest virtual machines.

Editing a Guest Virtual Machine

1. In Virtualization Central, select the Guest Virtual Machine you want to edit, and select Action Guest VM, Edit Guest VM and click **Go**.

2. In the Edit Guest Virtual Machine page, edit the values. Click **Continue**.

The Edit deployment procedure will be submitted. The guest virtual machine being edited will appear with the status 'Editing' in Virtualization Central. Once the deployment procedure completes successfully, the guest virtual machine status will be changed to the current state.

12.6.3 Updating Preferred Virtual Server List

If you want to restrict certain guest virtual machines to run on only certain Oracle VM servers, use the Preferred server list feature. Preferred server list contains all the Oracle VM servers on which this guest virtual machine is allowed to run. If the "Automatic" option is selected when creating a guest virtual machine, Enterprise Manager Grid Control picks the Oracle VM server with maximum resources to run the guest virtual machine. When auto-restarting guest virtual machines during fail-over scenario, Enterprise Manager Grid Control picks a preferred Oracle VM server to auto-restart the guest virtual machine.

To update preferred virtual server list, follow these steps:

1. In Virtualization Central, select the Guest Virtual Machine for which you want to update the preferred server list, and select Action Guest VM, Update Preferred Server List and click **Go**.

2. In the Update Preferred Server List page, choose the preferred server list as manual or automatic. If you choose Manual, add the preferred virtual servers for the guest virtual machine. Click **Apply**.
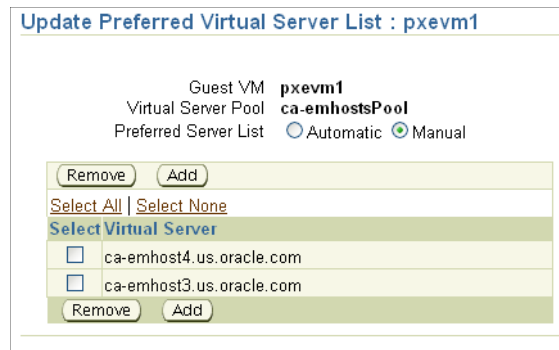
Enterprise Manager picks the preferred Oracle VM server to start the guest virtual machine.

12.6.4 Opening a VNC Console to the Guest Virtual Machine

You can log in to guest virtual machines using VNC Console.

1. In Virtualization Central, select the guest virtual machine you want to access and select Action Guest VM VNC Console.

   The VNC Console to the guest virtual machine opens in a separate window.
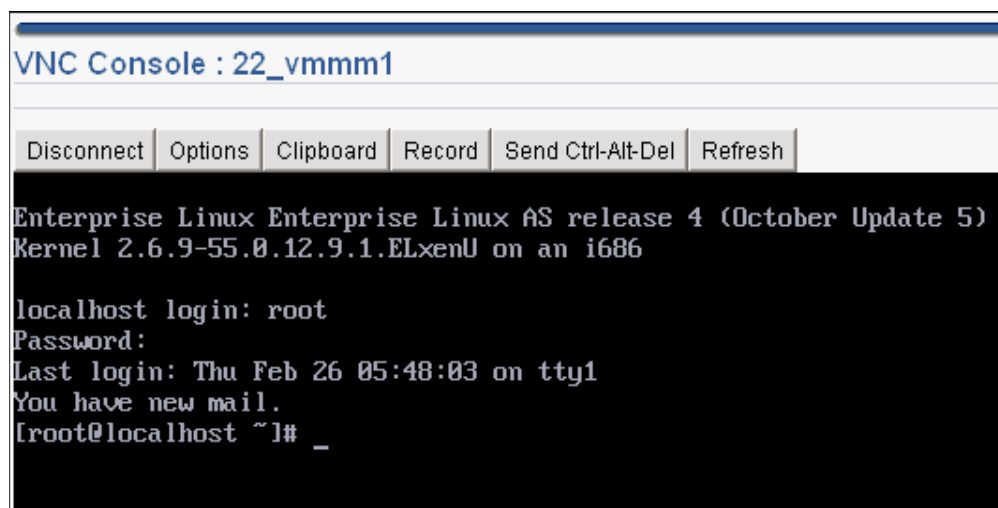
2. In the VNC Console, specify the VNC password.

12.6.5 Deleting Server Pools, Oracle VM Servers, and Guest Virtual Machines

On removing a server pool, the virtual servers associated with the server pool will be removed. This section explains the following:

- Prerequisites for Deleting a Server Pool

- Deleting a Server Pool

- Prerequisites for Deleting an Oracle VM Server

- Deleting an Oracle VM Server

- Deleting a Guest Virtual Machine

12.6.5.1 Prerequisites for Deleting a Server Pool

Following are the prerequisites for deleting a server pool:

- The virtual server pool must not have any running guest virtual machines.

- No Import Resource operation must be running in the virtual server pool.

- Agents running on all the virtual servers in the virtual server pool should be accessible.

12.6.5.2 Deleting a Server Pool

Follow these steps to delete a server pool:

1. In Virtualization Central, select the Server Pool, select Action Virtual Server Pool, Remove Server Pool, and click **Go**.

2. Confirm that you want to remove the server pool. You can choose **Remove all working directories** and/or **Force Remove**. Select Remove all working directories from the server pool if you want to delete all the related directories (under /OVS). If you leave this option unselected, only the server pool data is removed from the database, while the relevant directories and files of the server pool remain on the server.

   If you select the **Force Remove** option, the virtual servers in the server pool will not be updated with the changes.

   > **Note:** If you select the **Force Remove** option, the validations mentioned earlier will not be performed.

3. Click **Finish**.

On clicking **Finish**, the deployment procedure to delete the server pool is submitted and the status of the Server Pool changes to Unregistering. Once the job is completed, the deleted Server Pool and the associated virtual servers will not appear in Virtualization Central.

12.6.5.3 Prerequisites for Deleting an Oracle VM Server

Following are the prerequisites for deleting an Oracle VM server:

- The master agent of the virtual server should be connected and accessible.

- There should not be any running guest virtual machines on the virtual server.

- The virtual server should not be a preferred server of any guest virtual machine.

- The virtual server should not be the master server. The master virtual server cannot be directly removed. To remove the master virtual server, the virtual server pool has to be deleted.

12.6.5.4 Deleting an Oracle VM Server

Follow these steps to delete an Oracle VM server:

1. In Virtualization Central, select the Oracle VM Server, select Action Virtual Server, Deregister Virtual Server, and click **Go**.

2. Confirm that you want to deregister the virtual server. You cannot deregister the master server from a server pool.

> **Note:** If you select the **Force Deregister** option, the validations mentioned above will not be performed.

3. Click **Finish**.

On clicking **Finish**, the deployment procedure to delete the virtual server is submitted and the status of the Virtual Server changes to `Unregistering`. Once the job is completed, the deleted Virtual Server and the associated guest virtual machines will not appear in Virtualization Central.

12.6.5.5 Prerequisite for Deleting a Guest Virtual Machine

Following is the prerequisite for deleting a guest virtual machine:

- The guest virtual machine to be deleted should be in the Halted status.

12.6.5.6 Deleting a Guest Virtual Machine

Follow these steps to delete a guest virtual machine:

1. In Virtualization Central, select the guest virtual machine, select Action Guest VM, Delete Guest VM, and click **Go**.

2. Confirm that you want to delete the guest virtual machine.

3. Click **Finish**.

On clicking **Finish**, the deployment procedure to delete the guest virtual machine is submitted and the status of the guest virtual machine changes to `Deleting`. Once the job is completed, the deleted guest virtual machine will not appear in Virtualization Central.

When you delete a guest virtual machine, all the files and data associated with this guest virtual machine are removed from the virtual server and server pool.

12.7 Live Migrating Guest Virtual Machines

Administrators may want to move guest virtual machines on an Oracle VM server to another Oracle VM server for balancing the application workload, or during planned downtime of the Oracle VM server. Enterprise Manager Grid Control allows users to live migrate one or all guest virtual machines. Applications running in the guest virtual machines continue to be accessible during live migration. Live migration ensures high availability of guest virtual machines.

Enterprise Manager Grid Control supports the following live migration cases:

- Live migrate one guest virtual machine to another virtual server
- Live migrate all guest virtual machines from one virtual server to another

This section explains the following:

- Prerequisites
- Live Migrating Selected Guest Virtual Machine
- Live Migrating All Guest Virtual Machines

12.7.1 Prerequisites

Live migration requires the following preconditions to be met:

- The virtual servers participating in the live migration must belong to the same server pool and reside in the same subnet.

- The computer make and model number of both the source computer and the destination computer must be identical.

- The destination computer must have enough memory and free space.

12.7.2 Live Migrating Selected Guest Virtual Machine

To live migrate a guest virtual machine, do the following:

1. In Virtualization Central, select the guest virtual machine you want to live migrate and select Action Guest VM Live Migrate Guest VM. Ensure that the guest virtual machine is in the Running state. Click **Go**.

2. In the Live Migrate: Guest Virtual Machine page, specify the virtual server to which you want to live migrate the running guest virtual machine and also schedule the guest virtual machine live migration job to run either immediately or later. Select the virtual server allocation type as either Automatic or Manual.

   - If you choose **Automatic**, the system will pick up a suitable virtual server from the server pool for the guest virtual machine live migration. This virtual server will be selected from the virtual servers configured in the preferred server list for the guest virtual machine.

   - If you choose **Manual**, click **Add** and select the virtual server from the list of virtual servers available in the server pool.

3. Click **OK**.

After performing the live migration activity, navigate to Virtualization Central. The status of the migrated guest virtual machines will change from `Running` to `Migrating` to `Running`. The guest virtual machines will be displayed under the new virtual server after a few minutes.

12.7.3 Live Migrating All Guest Virtual Machines

To live migrate all guest virtual machines, do the following:

1. In Virtualization Central, select the virtual server from which you want to live migrate all guest virtual machines, and select Action Virtual Server Live Migrate All Guest VM. Ensure that it is in the Up state.

2. In the Live Migrate: Guest Virtual Machine page, schedule the guest virtual machine live migration job to run either immediately or later. The virtual server allocation type is Automatic.

3. Click **OK**. If there are no running guest virtual machines, a warning message is displayed. You can ignore the warning and continue with the live migration operation.

After performing the live migration activity, navigate to Virtualization Central. The migrated guest virtual machines must now be displayed under the new virtual server.

12.8 High Availability and Fault Tolerance

Oracle VM server uses Oracle Clusterware to maintain network and storage heartbeat in a cluster of Oracle VM servers. When an Oracle VM server failure is detected, Enterprise Manager Grid Control automatically restarts the guest virtual machines that were running on the failed server on a new Oracle VM server.

This section explains how you can create a cluster and enable high availability (HA). You can create either an NFS-based cluster, or an OCFS2 cluster. In particular, this section covers the following:

- Creating an NFS Cluster for High Availability

- Creating an OCFS2 Cluster for High Availability

- Preparing a Storage Repository (For Oracle VM 2.2)

- Configuring Oracle VM Servers and Guest Virtual Machines for High Availability

- Master Server Failover

12.8.1 Creating an NFS Cluster for High Availability

The following sections list the procedure to create an NFS cluster for high availability for Oracle VM version 2.2, 2.1.2., and 2.1.5.

Creating an NFS-based cluster for High Availability for Oracle VM 2.2

To configure HA, you must first create external storage, and then create a storage repository to be used in the server pool.

To prepare external storage to be used in a storage repository using NFS:

1. To prepare external storage to be used in a storage repository using NFS, find an NFS mount point to use. For example:

```
mynfsserver:/vol/vol1/data/ovs
```

When you have identified an NFS mount point to use for the external storage, prepare the storage repository.

Creating an NFS-based cluster for High Availability for Oracle VM 2.1.2 and 2.1.5

To create an NFS-based cluster for HA for Oracle VM 2.1.2 and 2.1.5:

1. On the master virtual server, start the cluster service:

   ```
   # /usr/lib/ovs/ovs-cluster-configure
   ```

   This script generates the /etc/ocfs2/cluster.conf file if it does not already exist, and configures and starts the o2cb service.

2. On the master virtual server, configure the shared storage with /usr/lib/ovs/ovs-makerepo source C description, for example.

   ```
   # /usr/lib/ovs/ovs-makerepo example.com:/OVS C "cluster root"
   ```

3. On the master virtual server, check the cluster has been configured correctly and is running:

   ```
   # /usr/lib/ovs/ovs-cluster-check --master --alter-fstab
   ```

   The --master parameter should only be used on the master virtual server. The --alter-fstab parameter modifies the /etc/fstab file. A backup of the original /etc/fstab file is created in /tmp.

4. Log in to Oracle Enterprise Manager and create an HA-enabled server pool. See Configuring Oracle VM Servers and Guest Virtual Machines for High Availability for steps to enable high availability on server pools.

12.8.2 Creating an OCFS2 Cluster for High Availability

The following sections list the procedure to create an OCFS2 cluster for high availability for Oracle VM version 2.2, 2.1.2., and 2.1.5.

Creating an OCFS2-based cluster for High Availability for Oracle VM 2.2

To configure HA, you must first create external storage, and then create a storage repository to be used in the server pool.

To prepare external storage to be used in a storage repository using OCFS2 on SAN:

1. Review the partitions by checking /proc/partitions:

   ```
   # cat /proc/partitions
   major minor  #blocks  name
      8     0   71687372 sda
      8     1     104391 sda1
      8     2   71577607 sda2
    253     0   70516736 dm-0
    253     1    1048576 dm-1
      8    16    1048576 sdb
   ```

2. Determine the share disk volume you want to use. You may need to use the fdisk utility to create at least one partition on the shared disk volume:

   ```
   # fdisk /dev/sdb
   ```

Alternatively, if the external storage is more than 2 TB, use the gparted utility to create a GPT partition:

```
# gparted /dev/sdb
```

3. Format the partition from any of the Oracle VM Servers in the (intended) server pool with a command similar to the following:

```
# mkfs.ocfs2 -Tdatafiles -N8 /dev/sdb1
```

The -Tdatafiles parameter makes mkfs.ocfs2 use a large cluster size. The size chosen depends on the device size.

Add an additional parameter, --fs-features=local, if you are not mounting the volume in the server pool.

The -N8 parameter allocates 8 slots to allow that many nodes to mount the file system concurrently. Increase the number if the server pool has (or will have) greater than 8 nodes. It is recommended that you have at least 8 slots, even for a local file system. Creating slots for a local file system allows users to easily cluster-enable the shared volume later.

For information about preparing external storage using OCFS2 on iSCSI, see "Chapter 8, Preparing External Storage and Storage Repositories" in the *Oracle VM Server User's Guide* available at:

http://download.oracle.com/docs/cd/E15458_01/doc.22/e15444/storage.htm#sthref68

When you have prepared the external storage, prepare the storage repository.

Creating an OCFS2-based cluster for High Availability for Oracle VM 2.1.2 and 2.1.5

To create an OCFS2-based cluster for HA for Oracle VM 2.1.2 and 2.1.5:

1. On the master virtual server, start the cluster service:

```
# /usr/lib/ovs/ovs-cluster-configure
```

This script generates the /etc/ocfs2/cluster.conf file if it does not already exist, and configures and starts the o2cb service.

2. On the master virtual server, configure the shared storage with /usr/lib/ovs/ovs-makerepo source C description, for example.

```
# /usr/lib/ovs/ovs-makerepo example.com:/OVS C "cluster root"
```

3. On the master virtual server, check the cluster has been configured correctly and is running:

```
# /usr/lib/ovs/ovs-cluster-check --master --alter-fstab
```

The --master parameter should only be used on the master virtual server. The --alter-fstab parameter modifies the /etc/fstab file. A backup of the original /etc/fstab file is created in /tmp.

4. Log in to Oracle Enterprise Manager and create an HA-enabled server pool. See Configuring Oracle VM Servers and Guest Virtual Machines for High Availability for steps to enable high availability on server pools.

Creating an OCFS2-based cluster for High Availability for Oracle VM 2.1.2 and 2.1.5

To create an OCFS2-based cluster for HA:

1. On the master virtual server, format a volume to be an OCFS2 partition with the `mkfs.ocfs2` utility. For example, to create an OCFS2 partition with a 4k block size, 4k cluster size, and 16 node slots:

   ```
   # mkfs.ocfs2 -L mylabel -b 4K -C 4K -J size=64M -N 16 /dev/sda1
   ```

2. On the master virtual server, start the cluster service:

   ```
   # /usr/lib/ovs/ovs-cluster-configure
   ```

   This script generates the /etc/ocfs2/cluster.conf file if it does not already exist, and configures and starts the o2cb service.

3. On the master virtual server, configure the OCFS2 partition as shared storage with /usr/lib/ovs/ovs-makerepo source C description, for example.

   ```
   # /usr/lib/ovs/ovs-makerepo /dev/sdb1 C "cluster root"
   ```

4. Check the cluster has been configured correctly and is running:

   ```
   # /usr/lib/ovs/ovs-cluster-check --master --alter-fstab
   ```

   The --master parameter should only be used on the master virtual server. The --alter-fstab parameter modifies the /etc/fstab file. A backup of the original /etc/fstab file is created in /tmp.

5. Log in to Oracle Enterprise Manager and create an HA-enabled server pool. See Configuring Oracle VM Servers and Guest Virtual Machines for High Availability for steps to enable high availability on server pools.

12.8.3 Preparing a Storage Repository (For Oracle VM 2.2)

For Oracle VM version 2.2 or higher, in addition to preparing external storage for a server pool, you must also prepare a storage repository. The storage repository is used to host the resources used in the server pool, such as guest virtual machines and shared virtual disks.

To prepare a storage repository, do the following:

1. On the Server Pool Master, create a storage repository with the script:

   ```
   /opt/ovs-agent-2.3/utils/repos.py -n storage_location
   ```

2. The repos.py -n command from the previous step displays the UUID of the storage repository you created. You need this UUID to set the cluster root. Copy the UUID for the storage repository you want to use as the cluster root.

   For example, for an NFS set up, you might use something similar to:

   ```
   # /opt/ovs-agent-2.3/utils/repos.py -n mynfsserver:/vol/vol1/data/ovs
   ```

   Or for an OCFS2 using iSCSI or SAN set up, you might use something similar to:

   ```
   # /opt/ovs-agent-2.3/utils/repos.py -n /dev/sdb1
   ```

3. Paste the UUID for the storage repository and use it to set the cluster root with the command:

   ```
   # /opt/ovs-agent-2.3/utils/repos.py -r UUID
   ```

The external storage and storage repository are configured and ready to use.

12.8.4 Configuring Oracle VM Servers and Guest Virtual Machines for High Availability

You can set up High Availability for Oracle VM to guarantee the availability of guest virtual machines if the Oracle VM server they are running on fails or restarts. When a virtual server is restarted or shut down, the guest virtual machines running on it are either restarted on, or migrated to, another virtual server.

Enabling High Availability mode for the virtual server pool allows all the Oracle VM servers in the pool to be clustered.

From Oracle VM version 2.2 onwards, Enterprise Manager provides the master server failover feature. If the master virtual server fails, another virtual server is selected from the server pool to act as the master virtual server. You can configure a virtual IP for the entire server pool. For more information about master server failover, see Master Server Failover.

To enable High Availability, you will need the following configuration:

- High Availability mode should be enabled for the virtual server pool.
- Clustering configuration should be done for master server and other virtual servers in the server pool.
- High Availability flag must be set for the guest virtual machines.

Steps to enable High Availability

1. When creating a server pool, in the Create Server Pool page, select **Enable** for the High Availability Mode field.

2. When creating a guest virtual machine, in the Create Guest Virtual Machine: Configure page, select **Enable High Availability Mode** check box.

3. When editing a guest virtual machine, in the Edit Guest Virtual Machine page, select **Enable High Availability Mode** check box.

If High Availability is configured for both server pool and the guest virtual machine, if the virtual server is shutdown, the guest virtual machine will be stopped and restarted on another virtual server.

12.8.5 Master Server Failover

To support the automatic failover and assignment of master virtual server in a pool, you can configure a virtual IP for the entire server pool. This will be a floating IP, and upon failover, the IP will be assigned to the new master virtual server.

> **Note:** This feature is available only for Oracle VM 2.2 or higher.

To enable master server failover, follow these steps:

■ During virtual server pool creation, in the Create Virtual Server Pool page, enable the High Availability Mode setting and specify the Master Server Virtual IP as follows:

■ You can also modify the master server virtual IP address from the Server Pool Home page as follows:

12.9 Advanced Operations

This section explains the advanced operations that can be performed in your virtualization environment. In particular, this section explains the following:

- Cloning Guest Virtual Machines
- Saving Guest Virtual Machine as Template
- Bare Metal Provisioning of Oracle VM Servers
- Server Maintenance Mode

12.9.1 Cloning Guest Virtual Machines

Cloning creates a copy of a guest virtual machine on the same server pool. When a guest virtual machine is cloned, the software on the source guest virtual machine is also copied onto the cloned guest virtual machine.

12.9.1.1 Prerequisites

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Guest virtual machine to be cloned must be in the `Halted` status.

12.9.1.2 Cloning Procedure

During cloning, new guest virtual machines are created from a source guest virtual machine.

To clone a guest virtual machine, do the following:

1. In Virtualization Central, select the guest virtual machine you want to clone and select Action Guest VM Clone Guest VM. Ensure that it is in the Halted state. Click **Go**.

2. In the Clone Guest Virtual Machines: Server Pool page, specify the number of guest virtual machines to be created. The new guest virtual machines will be associated with the same server pool as the guest virtual machine you are cloning. You can also select the virtual server for the guest virtual machine manually or allow the server pool to automatically allocate the virtual server. Click **Next**.

3. In the Clone Guest Virtual Machines: Configure page, configure the guest virtual machines by specifying the values in Table 12–10.

Clone Guest Virtual Machines: Configure

Virtual Machine Name Prefix `clo`
Guest Virtual Machine names will be prefix1, prefix2 depending on the number of guest virtual machines to be created.

Root Password ••••

Confirm Root Password ••••

VNC Console Password •••••

Confirm VNC Console Password •••••

Number of Network Interfaces `1` [Configure Network]
Number of Network Cards

CPU Scheduling Priority [Intermediate ▼] `50`
Oracle VM Server will schedule the physical CPU cycles, based on the CPU Scheduling Priority of the Guest VM

CPU Cap (%) [High ▼] `100`
Oracle VM Server will limit the allocation of physical CPU cycles, based on the CPU Cap of the Guest VM

Start VM after creation ☑

Enable High Availability Mode ☑

Configure Agent ☐

Local Devices from Source VM

Disk Name	Capacity(MB)	QoS Priority Class
linux	3000	Not Enabled ▼

Additional Disks

[Add]

Select Disk Name	Capacity(MB)	QoS Priority Class
No Additional Disks		

☑ TIP Specify a disk name for the additional disk to be added. A valid disk name can contain only alpha-numeric and underscore characters.

Shared Disks

[Add]

Select Shared Disk Name	QoS Priority Class
No Shared Disk	

4. In the Clone Guest Virtual Machines: Schedule/Credentials page, provide a schedule for the provisioning and user credentials for the monitoring server and Virtual Server host. The credentials you specify for the Virtual Server host must have sudo privileges.

5. In the Clone Guest Virtual Machines: Review page, review the information you have provided and submit the job.

The cloned guest virtual machine will appear with the status `Creating..` in Virtualization Central under the same server pool as the source guest virtual machine. The source guest virtual machine will appear with the status `Cloning..` in Virtualization Central. If you have selected the **Start VM after creation** option during cloning, the guest virtual machines will be started after the clone process.

Table 12–10   Clone Guest Virtual Machines

Element	Description
Virtual Machine Name Prefix	This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager Grid Control Console, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.

Table 12–10   (Cont.) Clone Guest Virtual Machines

Element	Description
VNC Console Password	Enterprise Manager Grid Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
Number of Network Interfaces	Number of virtual network interfaces this guest virtual machine should have. The minimum value is 1 and the maximum value is 8.
	Specify the network details by clicking **Configure Network**.
CPU Scheduling Priority	Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
CPU Cap (%)	Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
Start VM after creation	Specify whether the guest virtual machine should be started automatically after it is created.
Enable High Availability Mode	Oracle VM allows sharing of disks between multiple guest virtual machines. You can specify multiple shared disks that the guest virtual machines can access.
	If the server pool has no shared disk available, the Add button will be disabled. Only virtual machines within the same server pool can share a shared disk.
	Specify the shared device name, the shared disk on it, and the QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
Additional Disks	Every guest virtual machine is created with at least one disk. Specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log in to the guest virtual machine and reformat and mount the additional disks.
	Specify the name of the hard disk, its size, and QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.
Shared Disks	Oracle VM allows sharing of disks between multiple guest virtual machines. You can specify multiple shared disks that the guest virtual machines can access.
	If the server pool has no shared disk available, the **Add** button will be disabled. Only virtual machines within the same server pool can share a shared disk.
	Specify the shared device name, the shared disk on it, and the QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

12.9.2 Saving Guest Virtual Machine as Template

A guest virtual machine can be saved as an Oracle VM template in the same server pool where the guest virtual machine exists. Additionally, if an external NFS location

is provided, a Software Library component is created for this template that can later be used for creating guest virtual machines on other server pools.

Saving a Guest Virtual Machine as a Template

1. In Virtualization Central, choose the guest virtual machine you want to save, select Action Guest VM, Save Guest VM as Template and click **Go**. Ensure that the guest virtual machine is in the Halted status.

2. In the Create Component: Describe page, specify the creation method as:

   ■ **From Existing Guest VM** to store an existing guest VM configuration details.

Specify the template name and OS details. Click **Next**.

3. In the Create Component: Configure page, optionally specify the external location where the template is stored and used for creating software library component. The software library component can then be imported to server pools to create guest virtual machine templates. Specify the monitoring server and virtual server credentials.

   Click **Next**.

4. In the Create Component: Review page, review the information you have provided and click **Submit**.

The template will be saved in the location you specified. The template saved will be imported to the server pool to which the Halted guest virtual machine belongs, you can view the imported template in the Server Pool Resources page.

12.9.3 Bare Metal Provisioning of Oracle VM Servers

Enterprise Manager Grid Control allows you to provision Oracle VM server software on bare metal machines. For information about bare metal provisioning, see Provisioning Linux.

During bare metal provisioning of Oracle VM server, the Management Agent will not be installed. After completing bare metal provisioning, register the Oracle VM server with the server pool using the steps explained in Section 3.6.3, "Registering Virtual Servers".

12.9.4 Server Maintenance Mode

The server maintenance mode will enable administrators to perform maintenance activity on virtual servers.

Virtual servers which are put in maintenance mode will lose their utility server role when the maintenance mode is enabled. Once the maintenance mode is removed, their roles will be restored. However, if there is only one utility server in the pool, it will not be allowed to be put into maintenance mode until another server is assigned with that role. A master virtual server cannot be put in maintenance mode.

When the virtual server is in maintenance mode, Enterprise Manager allows you to enable blackout on that virtual server to avoid any unnecessary alerts coming in to the system.

The following actions cannot be performed when a virtual server is in the maintenance mode:

- Live migrate to the virtual server
- Start and Resume operations on guest virtual machines

12.9.5 Enabling Server Maintenance Mode

To enable server maintenance mode, follow these steps:

1. In Virtualization Central, select Action as Enable Server Maintenance Mode, and click **Go**.

2. In the Enable Server Maintenance Mode page, select the server pool and add the virtual servers on the server pool for which you want to enable server maintenance. Select the blackout options:

   - **Blackout Oracle VM Server and all Host targets running on it** will blackout virtual servers and all hosts running on it. You can also choose to run jobs on hosts during the blackout and blackout targets associated with the selected hosts and monitored by the same agent.

   - **Blackout Only Oracle VM Server** will blackout only the virtual servers put in maintenance mode.

   - **No Blackout** will not blackout Oracle VM servers or any associated hosts.

3. Click **OK**.

The selected virtual servers will be put in maintenance mode.

12.9.6 Removing Server Maintenance Mode

To remove server maintenance mode, follow these steps:

1. In Virtualization Central, choose the virtual server pool, select Action as Remove Server Maintenance Mode, and click **Go**.

2. In the Confirmation Page, click **Yes** to remove server maintenance mode.

   The virtual server will be removed from maintenance mode.

12.10 Migrating from Oracle VM Manager to Enterprise Manager

If you are using Oracle VM Manager to manage Oracle VM servers, guest virtual machines, and resources, and want to move to Enterprise Manager to manage your Oracle VM systems, follow these steps:

1. Delete the virtual server from Oracle VM Manager. For information about deleting a virtual server from Oracle VM Manager, see Oracle VM Manager User's Guide available at `http://download.oracle.com/docs/cd/E11081_01/doc/doc.21/e10901/toc.htm`

2. Shutdown the virtual server.

3. Create virtual server pool in Enterprise Manage. For information about creating virtual server pools, see Section 3.6.2, "Creating Server Pools".

4. Register virtual servers with the server pool. For information about registering virtual servers, see Section 3.6.3, "Registering Virtual Servers".

5. Create guest virtual machines. For information about creating guest virtual machines, see Section 12.5.3, "Creating Guest Virtual Machines".

Part V

Middleware Provisioning

This part describes the Deployment Procedures you can run for provisioning Middleware targets. In particular, this part the Deployment Procedures for:

13

Provisioning Oracle Fusion Middleware

This chapter assumes that you have applied the *Cumulative Fixes for Oracle Enterprise Manager 11g Grid Control Fusion - May 2010* (*interim patch number 9534192*). For more details, refer to *My Oracle Support* note 427577.1. You must apply all the recommended patches listed in the *Readme.txt* shipped with this patch.

This chapter explains how you can automate common provisioning operations for Oracle Fusion Middleware 11g and Oracle WebLogic Domains.It covers the following:

- Getting Started

- Deployment Procedure

- Supported Releases

- Best Practices while Running the Fusion Middleware Provisioning and Scale Up Deployment Procedures

- Fusion Middleware Domain Provisioning

- Fusion Middleware Domain Scale Up

13.1 Getting Started

This chapter helps you get started by providing an overview of the steps involved in provisioning WebLogic domains, Fusion Middleware homes or scaling up a WebLogic domain. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision or extend a WebLogic domain. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 13–1   Getting Started with Provisioning and Scaling Up Fusion Middleware Domains

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedures offered by Enterprise Manager Grid Control for provisioning WebLogic domains, Fusion Middleware homes, and scaling up / out WebLogic domains. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedures, see Section 13.2, "Deployment Procedure".

Table 13–1 (Cont.) Getting Started with Provisioning and Scaling Up Fusion Middleware Domains

Step	Description	Reference Links
Step 2	**Knowing About the Supported Releases** Know what releases can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 13.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers the use cases for provisioning a WebLogic Domain, Fusion Middleware Home, and scaling a WebLogic Domain. Select the use case that best matches your requirement.	■ To learn about provisioning a WebLogic Domain or a Fusion Middleware Home, see Section 13.5, "Fusion Middleware Domain Provisioning". ■ To learn about scaling a WebLogic Domain, see Section 13.6, "Fusion Middleware Domain Scale Up".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as space requirements and user privileges. When you run the Deployment Procedure, these prerequisites are verified.	■ To learn about the prerequisites for provisioning a WebLogic domain, see Section 13.5.1, "Prerequisites". ■ To learn about the prerequisites for scaling up a Fusion Middleware domain, see Section 13.6.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully clone a WebLogic domain, Fusion Middleware Home, or scale up a WebLogic domain.	■ To provision a WebLogic domain, follow the steps explained in Section 13.5.2, "Running the Fusion Middleware Provisioning Deployment Procedure". ■ To scale up a WebLogic Domain, follow the steps explained in Section 13.6.2, "Running the Fusion Middleware Domain Scale Up Deployment Procedure".

13.2 Deployment Procedure

Enterprise Manager Grid Control offers the following deployment procedures for provisioning WebLogic domains, Fusion Middleware homes, and scaling up WebLogic domains:

■ *Fusion Middleware Domain Provisioning*: Using this Deployment Procedure, you can automate the provisioning of WebLogic domains and Fusion Middleware homes using a live reference installation. Specifically, you can automate the provisioning of:

■ WebLogic Domains with a corresponding Oracle Fusion Middleware Home which includes WebLogic Servers and one or more Oracle Homes.

Figure 13–1   WebLogic Domain Provisioning
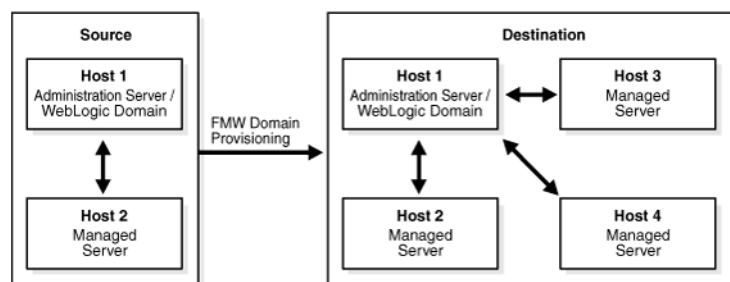
- Oracle Fusion Middleware Homes.

Figure 13–2   Fusion Middleware Home Provisioning

> **Note:** The Fusion Middleware Domain Provisioning deployment procedure does not support provisioning from a gold image.

- *Fusion Middleware Domain Scale Up*: Using this Deployment Procedure, you can expand or scale up a domain by increasing the capacity of the WebLogic server cluster. You can scale up a domain by:
  - Adding a new managed server to an existing cluster.
  - Provisioning a managed server.

Figure 13–3   Fusion Middleware Domain Scale Up

> **Note:** These deployment procedures start or create (using the FMW Provisioning Deployment Procedure) an Administration Server, Managed Servers, and Node Managers on destination hosts. With these deployment procedures, users can also perform generic domain configuration tasks on the destination domain.
>
> But, application specific or product specific artifacts present in the source domain will not be cloned on the destination domain. For example, in a SOA domain, these artifacts may include SOA Composites, Web Service Policies, and so on. You need to run the SOA Artifacts Provisioning Deployment Procedure to configure them. For more information, see Chapter 14, "Provisioning SOA Artifacts and Composites".

13.2.1 Steps Involved in Provisioning WebLogic Domains

This section lists the sequence of operations that are performed when you run the Fusion Middleware Provisioning Deployment Procedure in Enterprise Manager Grid Control.

- Verify Source
- Verify Destination
- Prepare Source Binary Archive
- Create Binary Clones on Destination
- Create Source Domain Archive
- Create Domain on Administration Server Host (Target)
- Configure Domain
- Create Domain Archive
- Create Domain on Managed Server Hosts
- Start Up Node Managers
- Start Up Admin Server
- Start Up Managed Servers
- Validate Domain
- Register Fusion Middleware homes and domains with OMS

13.2.2 Steps Involved in Scaling Up / Out WebLogic Domains

You can scale up a WebLogic Domain by adding a new managed server to existing cluster or by provisioning a managed server to an existing host machine that is already running a managed server in the domain or cluster. You can scale out the domain by adding a managed server or cloning a managed server to a remote, new host machine that does not yet exist in the domain. This section lists the sequence of steps that are involved in scaling up / out a WebLogic Domain.

- Verify Source
- Verify Destination
- Prepare Source Binary Archive
- Create Binary Clones on Destinations

- Clone / Add Managed Servers
- Create Domain Archive
- Create Domain on Managed Server Host
- Start Up Node Managers
- Start Up Managed Servers
- Register Fusion Middleware homes and domains with OMS

13.3 Supported Releases

Using these Deployment Procedures, you can provision and scale up Oracle Fusion Middleware Homes and WebLogic Domains using Oracle Fusion Middleware 11g Release 1 Patch Set 1 (11.1.1.2.0) which includes Oracle WebLogic Server Release 10.3.2.

> **Note:** To view a list of supported releases, refer to *Oracle Enterprise Manager Grid Control Certification Checker (ID 412431.1)* note on My Oracle Support.

13.4 Best Practices while Running the Fusion Middleware Provisioning and Scale Up Deployment Procedures

This section lists some of the best practices to be followed while using Fusion Middleware Domain Provisioning and Domain Scale Up.

- **Configuration of the source domain should not be changed**: While executing these deployment procedures, ensure that no administrative activities (such as configuration changes on the source domain and software patching) are actively performed on the source domain. If you change the configuration, the managed server may not respond to requests and the Administration Server will have an Unknown status.

- **Provisioning on the same machine**: If you are using the Deployment Procedure to provision or scale up to the same machine as the source, ensure that the source working directory and the destination working directory are different. For example, if the working directory is /tmp/source for the source machine, it should be /tmp/dest on the destination directory.

- **Unique Farm Prefix**: While using the Fusion Middleware Domain Provisioning Deployment Procedure, ensure that the farm prefix is unique. The farm prefix gets appended to the domain name to uniquely identify a given domain in Enterprise Manager Grid Control.

- **JDBC Configuration**: While configuring the JDBC data sources, you need to provide appropriate passwords. If the passwords are not specified, the source passwords will be used.

- **Custom Java Applications and their Deployment Plan**: These deployment procedures support custom java applications in staged mode. Externally staged applications need to be manually deployed. For instructions on manually deployment, see the *WebLogic Administration Guide*.

- **Multi NIC Machines**: If the destination machine is a multi NIC system, enter a listen address that is accessible to both the Administration Server and Managed Server.

13.5 Fusion Middleware Domain Provisioning

This section describes how to provision a WebLogic Domain or a Fusion Middleware Home from an existing installation. It contains the following:

- Prerequisites

- Running the Fusion Middleware Provisioning Deployment Procedure

13.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- The target machine should contain sufficient space. If the size of the Fusion Middleware Home on the source machine is 3 GB, you need approximately 3 GB in the working directory on the source and destination machines. Additionally, the destination machine should also have 3 GB of space for the Fusion Middleware Home. The working directory is cleaned up after deployment procedure has been successfully completed.

- The Fusion Middleware Home directory you specify on the destination machine must be a new directory or an empty one.

- The operating system user who runs the deployment procedure must belong to the same group and have the same permissions as the user who owns the Oracle Universal Installer Inventory. This will ensure that the group ownership on the Middleware Home directory on the destination machine will match the group ownership on the OUI inventory.

- Ensure that the 11.1 Management Agent has been installed on the source (on which the Administration Server is running) and destination machines.

13.5.2 Running the Fusion Middleware Provisioning Deployment Procedure

To provision a Fusion Middleware Domain, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. Click **Deployment Procedures** in the Deployment Procedure Manager section.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Fusion Middleware Provisioning.** Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source page of the Oracle Fusion Middleware Deployment Procedure.

   > **Note:** You cannot customize the deployment procedure and include additional steps that require user input.

4. On the Select Source page, do the following:

   a. In the Select Existing installation section, you can choose a WebLogic Domain or a Fusion Middleware Home. Select WebLogic Domain if you want to provision the domain and the software. Select Fusion Middleware Home if you want to provision the software only. The Fusion Middleware Home is provisioned first in a new environment. Domains are provisioned later and associated with the same software. Click the torch icon in the **Select WebLogic Domain** field and select a domain from the list. In the Source Information section, the WebLogic Domain Components including the **Host**, **Fusion**

Middleware Home, **WebLogic Server Home**, and **Fusion Middleware Domain Location** are displayed.

Note: The deployment procedure may fail when the following situations occur:

- You select a WebLogic Domain that is not supported. To successfully complete the deployment, you must select a supported version from the list.

- The corresponding Oracle homes are not listed under the selected domain. Before proceeding with the deployment, click **Refresh Hosts** in Enterprise Manager Grid Control console to ensure that the Oracle homes are properly discovered or select a different domain.

---

b. If you select a Fusion Middleware Home, click the torch icon in the **Select Fusion Middleware Home** field and select a Fusion Middleware Home from the list. In the Source Information section, the **Host**, **Fusion Middleware Home**, and **WebLogic Server Home** are displayed.

c. In the Working Directory field, specify the directory on the Administration Server machine on which the provisioning related files are temporarily stored. If this directory is not present, it will be created. When the provisioning operation has been completed, the contents of this directory will be deleted.

d. In the Host Credentials section, select **Preferred Credentials** so that the preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Enterprise Manager Grid Control. For more information on setting up Preferred Credentials, see the *Enterprise Manager Security* chapter in the *Enterprise Manager Administration Guide*.

Note: You can optionally override these preferred credentials. For example, if you have added two source hosts where the users are A and B, then you can choose **Override Credentials - Different for Each Host** and specify different credentials for each of the hosts. Similarly, if the destination hosts have the same credentials, which may be different from the preferred credentials, then you can choose **Override Credentials - Same for All Hosts**.

The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation. These include privileges to run `sudo` and `powerbroker` commands.

- The Deployment Procedure has been customized to support locked environments.

For more information, see Section 31.7, "Customizing Steps to Run as Another User"

Figure 13–4 Select Source Page

e. Click **Next**.

Enterprise Manager Grid Control displays the Select Targets page.

5. On the Select Targets page you can specify the target hosts on which the domain is to be created.

Figure 13–5 Select Targets Page

You can do the following:

a. Click on the torch icon in the Host column and select a host from the list.

b. If you choose to override the preferred credentials, specify the user name and password for the selected host.

c. In the Fusion Middleware Home field, enter the full path to the directory in which the Fusion Middleware Home is to be created.

d. If the JDK on the destination machine is in an external location, enter the full path to the location in the JDK Home Location field.

> **Note:** If you are using the WebLogic Server generic installer, the external JDK location must be specified. If this location is not specified, the default JDK of the Management Agent will be used.

    e. Check the **Agent Credentials** checkbox to specify the agent credentials. This field is applicable if the user who installed the Management Agent is different from the user installing Oracle Fusion Middleware.

    f. Click **Next** to proceed to the Configure Domain page.

6. On this page, you can enter the basic properties that are most likely to be reconfigured like domain name, administration server, clusters, machines, server, and JDBC Data Sources. For more information on configuring the domain, see the *Creating WebLogic Domains Using the Configuration Wizard* guide.

    a. In the Domain Properties page, enter a unique name for the domain, the username for the Administration Server and the password. You also need to enter the name of the Farm and the Domain Location (path to the directory in which the domain is to be created). The Domain Location directory must be empty.

Figure 13–6   Domain Properties Page
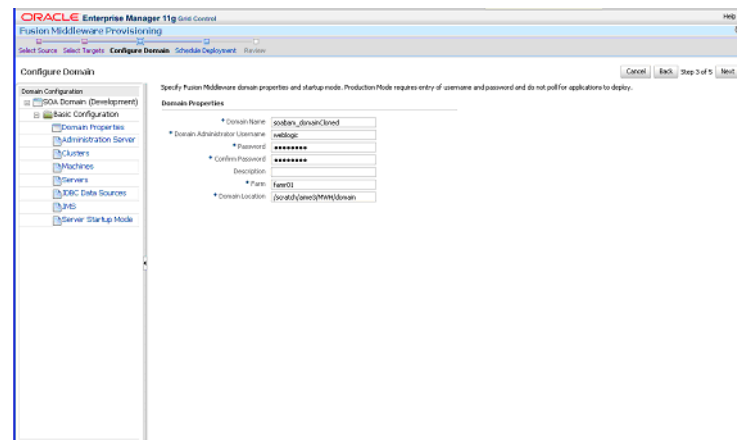
> **Tip:** A farm is a collection of components managed by Enterprise Manager Grid Control. It can contain Oracle WebLogic Server domains, one Administration Server, one or more Managed Servers, and the Oracle Fusion Middleware components that are installed, configured, and running in the domain. For more information, see the *Oracle Fusion Middleware Concepts Guide*.

    b. In the Administration Server page, enter the name of the Administration Server that will be used to manage the domain, listen address and port, and the name of the host on which the Administration Server is to be installed. If you checked the Enable SSL checkbox, enter the number of the SSL port for secure requests. You must ensure that the port numbers you specify for the Listen Port and SSL Listen Port are available.

> **Note:** If you are using the SSL configuration, you must ensure that the security/identity stores are present in the file system under the same path as on the source and are configured with certificates generated for the destination hosts.

Figure 13–7   Administration Server Page
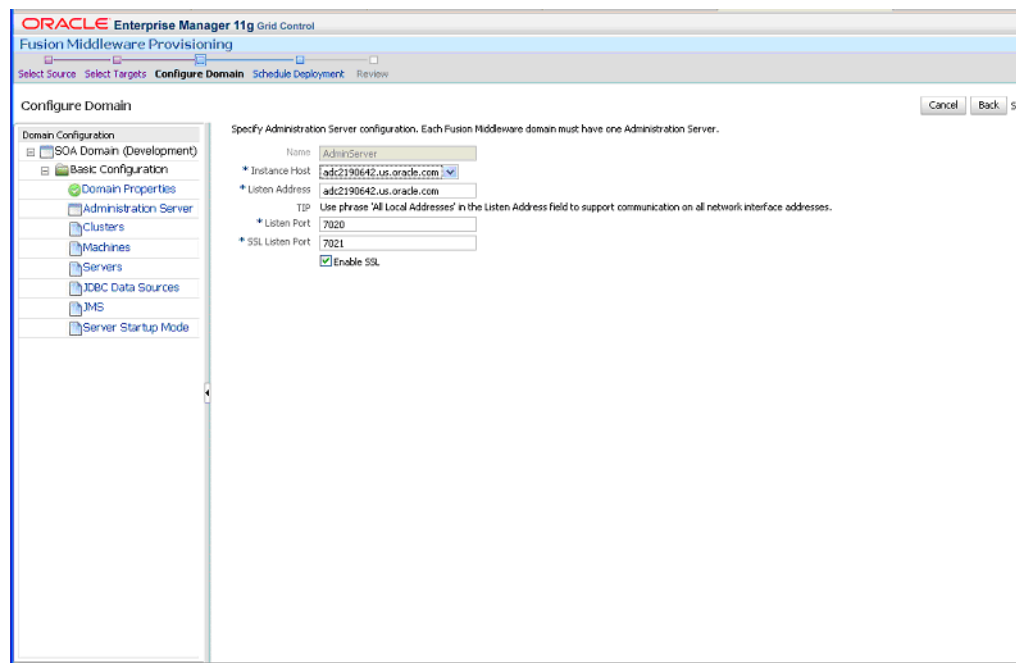
c. In the Clusters page, enter a name for the cluster, the cluster address that identifies the Managed Servers in the cluster, and the messaging mode (multicast or unicast). If you selected Multicast as the Messaging Mode, enter the multicast address and port.

Figure 13–8   Clusters Page
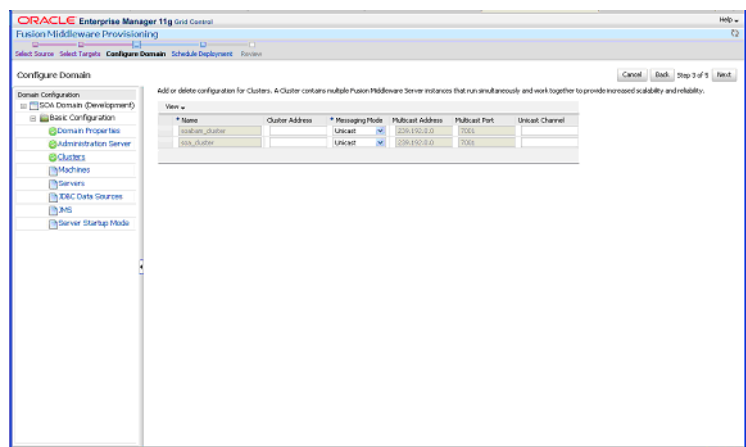
d. In the Machines page, enter the configuration information for the machines in the domain. You can add or delete configuration for the machines on this page.

Figure 13–9   Machines Page
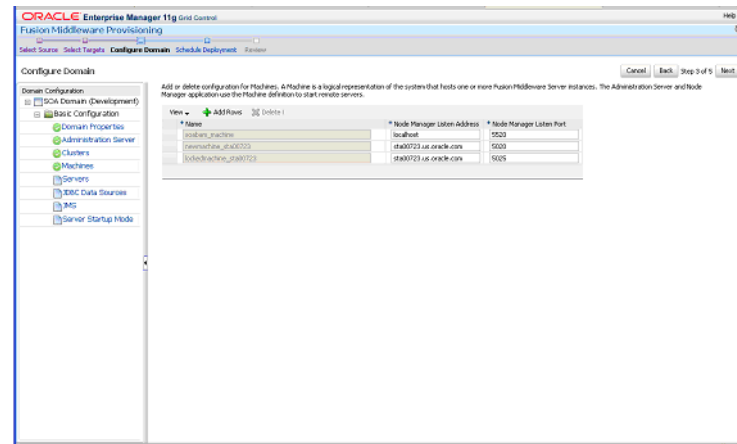
e. In the Servers page, enter the configuration information for the managed server.

Figure 13–10   Servers Page
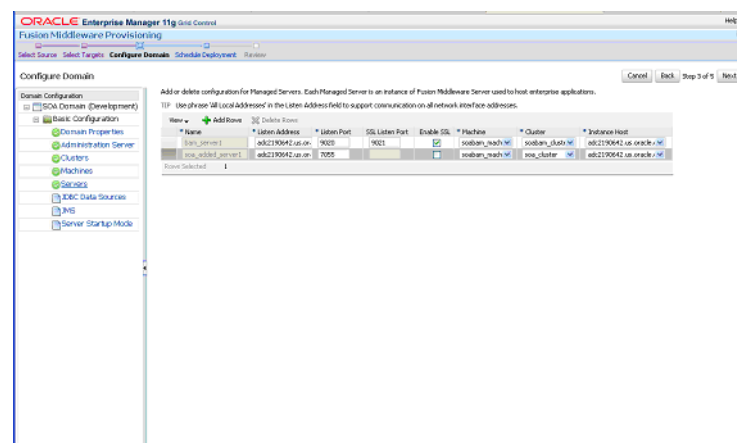
Tip: The port numbers you specify for the Listen Port and the SSL Listen Port on the Machines and Servers pages must be available on the target machine.

f. In the JDBC Data Sources page, enter the configuration information for the data source. In the Driver field, select the correct driver from the drop down list. Based on the driver you have selected, enter the URL in the correct format. Select the target from the drop down list and specify the database user name and password.

Figure 13–11   JDBC Data Sources Page
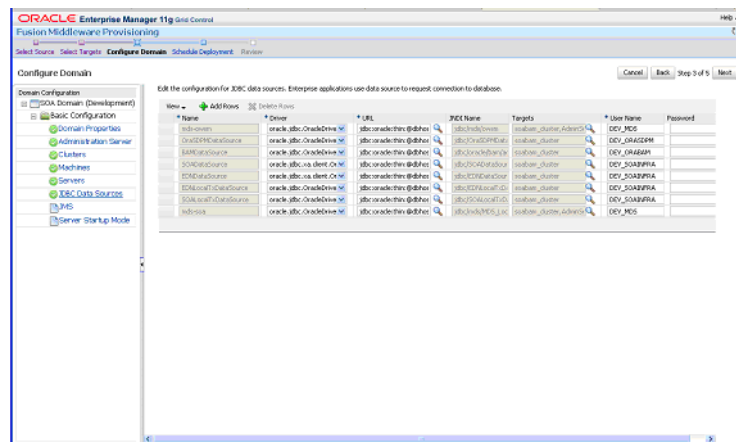
g. In the JMS page, you can add new JMS persistent stores and JMS servers.

Figure 13–12   JMS Page
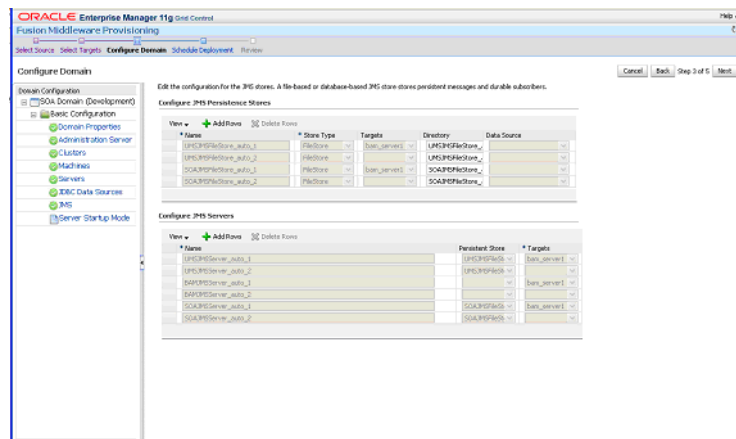
h. In the Server Startup Mode page, you can optionally start up the Managed Servers and the Admin Server. You can choose the following options:

- Start Administration and Managed Servers

- Start Administration Server

- Do not start Administration Server or Managed Server

Figure 13–13   Server Startup Mode Page
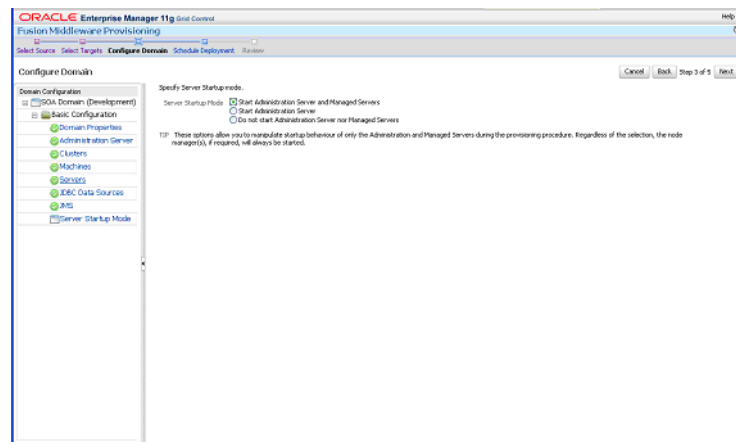
7. On the Schedule Deployment page, you can schedule the provisioning of the domain by specifying the time zone, time at which the deployment process should be started and other details.

8. On the Review page, review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Submit** to run the Deployment Procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes.

    After submitting the Deployment Procedure, you should see a link on top of the page. Click the link to view its progress. Alternatively, from the Deployments page, in the Deployment Procedure Manager section, click **Procedure Completion Status** to view the status of the Deployment Procedure.

9. After the Deployment Procedure has been successfully completed, the newly cloned domain will be added as a new target in Enterprise Manager Grid Control and can be monitored along with the other Fusion Middleware targets.

13.6 Fusion Middleware Domain Scale Up

The Fusion Middleware Domain Scale Up Deployment Procedure allows you to increase the capacity of a domain so that the performance of the applications is improved. This section describes how to scale up a WebLogic Domain by adding a new managed server to an existing cluster or cloning a managed server to an existing host machine that is already running a managed server in the domain/cluster. Or you can scale out the domain by adding a managed server or cloning a managed server to a remote, new host machine that does not yet exist in the domain. It contains the following sections:

- Prerequisites
- Running the Fusion Middleware Domain Scale Up Deployment Procedure

13.6.1 Prerequisites

Before running the Deployment Procedure, meet the prerequisites listed in Section 13.5.1, "Prerequisites".

13.6.2 Running the Fusion Middleware Domain Scale Up Deployment Procedure

To scale up / scale out a WebLogic Domain, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Fusion Middleware Domain Scale Up.** Then click **Schedule Deployment**.

    Enterprise Manager Grid Control displays the Select Source page of the Oracle Fusion Middleware Domain Scale Up Procedure.

    Figure 13–14   Select Source Page (Scale Up)

4. On the Select Source page, you can do the following:

    a. Click the torch icon in the **Select Fusion Middleware Domain** field and select a domain from the list. In the Source Information section, the details of the source domain including the **Host**, **Fusion Middleware Home**, **WebLogic Server Home**, and the **Fusion Middleware Domain Location** are displayed.

        > **Note:**   While running the deployment procedure, ensure that the Administration Server and Managed Server are running. You must also select a WebLogic domain that is supported from the list. If you select an unsupported version, the deployment procedure will fail.

    b. In the Working Directory field, specify the directory on the Administration Server machine on which the domain scale up related files are temporarily stored.   If this directory is not present, it will be created. When the scale up operation has been completed, the contents of this directory will be deleted.

    c. In the Provisioning Mode section, enter the credentials for the Administration Server Console and the Administration Server Host. For the Administration Server Host, you can choose Preferred Credentials or choose Override Preferred Credentials and specify a user name and password for the host on which the server is running. The Preferred Credentials option will be available only if it has already been defined in Enterprise Manager Grid Control. For

more information on setting up Preferred Credentials, see the *Enterprise Manager Security* chapter in the *Enterprise Manager Administration Guide*.

d. Click **Next**.

Enterprise Manager Grid Control displays the Configure Domain page.

Figure 13–15   Configure Domain Page - Scale Up

5. On the Configure Domain page, you can add a new managed server, clone an existing one, or delete a managed server you have created. You can do the following:

a. Click the **Add Server** tab in the left panel. A new server is added to the list and you need to enter the configuration details on the right panel.

b. Select a server from the list and click the **Clone Server** tab in the left panel. A new server which is a copy of the one you selected will be created. You must configure the new server to proceed with the provisioning.

c. In the Configure Managed Server section, enter the name of the host on which the managed server is to be added, the name of the new server, listen address, port, and the SSL port if applicable. If you checked the **Enable SSL** checkbox, enter the number of the SSL port for secure requests.

> **Tip:**   The port numbers you specify for the Listen Port and the SSL Listen Port must be available on the target machine.

d. In the Configure Machines section, you can configure the Machine and the Node Manager associated with the managed server instance.

> **Note:**   To perform administrative operations such as start and stop from the Enterprise Manager Grid Control, the Node Manager must be configured and running on the machine.

e. In the Software Installation section, in the Working Directory field, enter the full path to the directory on which the files required for scale up will be staged.

f. In the Managed Server Credentials section, enter the user name and password for the host on which the managed server is to be created.

    g. Check the **Agent Credentials** checkbox to specify the agent credentials. This field is applicable if the user who has installed the Management Agent is different from the user installing Oracle Fusion Middleware.

    h. Click **Next** to proceed to the Schedule page.

6. On the Schedule Deployment page, you can schedule the domain scale up by specifying the time zone, time at which the deployment process should be started and other details.

7. On the Review page, review the details you have provided for the Deployment Procedure. If you are satisfied with the details, then click **Submit** to run the Deployment Procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes.

After submitting the Deployment Procedure, you should see a link on top of the page. Click the link to view its progress. Alternatively, from the Deployments page, in the Deployment Procedure Manager section, click **Procedure Completion Status** to view the status of the Deployment Procedure.

8. After the Deployment Procedure has been successfully completed, the newly added Managed Servers will be added to the domain and the domain will be refreshed in the Enterprise Manager Grid Control console.

14

Provisioning SOA Artifacts and Composites

This chapter explains how you can provision SOA Artifacts and Composites. In particular, this chapter covers the following:

- Deployment Procedures, Supported Releases, and Core Components Deployed
- Provisioning SOA Artifacts
- Deploying SOA Composites

14.1 Getting Started

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

Table 14–1   Getting Started with Provisioning SOA Artifacts and Composites

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the two Deployment Procedures that are offered by Enterprise Manager Grid Control for provisioning SOA Artifacts and Composites. Know how the two Deployment Procedures function, what use cases they cover, what releases they support, and what core components they provision.	To learn about the Deployment Procedure, see Section 14.2, "Deployment Procedures, Supported Releases, and Core Components Deployed".
Step 2	**Selecting the Deployment Procedure to Provision** This chapter covers use cases for different SOA components. Identify the component you want to provision and understand the use cases that are covered.	■ To learn about provisioning SOA Artifacts, see Section 14.3, "Provisioning SOA Artifacts". ■ To learn about provisioning SOA Composites, see Section 14.4, "Deploying SOA Composites".

Table 14–1   (Cont.) Getting Started with Provisioning SOA Artifacts and Composites

Step	Description	Reference Links
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for provisioning SOA artifacts and composites, access the reference links provided for Step (2) and navigate to the *Prerequisites* subsection.
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision SOA Artifacts and Composites.	To provision SOA Artifacts and Composites, access the reference links provided for Step (2) and navigate to the *Provisioning Procedure* subsection.

14.2 Deployment Procedures, Supported Releases, and Core Components Deployed

Enterprise Manager Grid Control offers the following Deployment Procedures for provisioning SOA Artifacts and Composites:

Deployment Procedure	Supported Releases	Artifacts Migrated
SOA Artifacts Provisioning	Oracle SOA Suite 11g (11.1.1.2.0)	▪ SOA Composites ▪ Oracle WebLogic Server Policies ▪ Assertion Templates ▪ JPS Policy and Credential Stores
Deploy SOA Composites	Oracle SOA Suite 11g (11.1.1.2.0)	▪ SOA Composites

> **Note:**   Provisioning of a gold image from the Software Library is not supported for Microsoft Windows Vista.

> **Note:**   Cloning of human workflow artifacts and B2B artifacts are not supported. For information about cloning human workflow artifacts, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*. For information about cloning B2B artifacts, see the *Oracle Fusion Middleware User's Guide for Oracle B2B*. These guides are available at:
>
> http://download.oracle.com/docs/cd/E14571_01/index.htm

14.3 Provisioning SOA Artifacts

This section describes how you can provision SOA artifacts. In particular, this section covers the following:

- Provisioning from a Reference Installation
- Provisioning SOA Artifacts from Gold Image

14.3.1 Provisioning from a Reference Installation

This section describes how you can provision SOA artifacts from one soa-infra domain to another.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

14.3.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Ensure that you have already provisioned Oracle SOA Suite 11g and its underlying Oracle WebLogic Server Domain.
- Ensure that all the components (not only the soa-infra domain) within the source and target Oracle WebLogic Server Domains are up and running.
- Ensure that the source and the destination soa-infra domains are of the same version.

14.3.1.2 Provisioning Procedure

To provision SOA artifacts (composites, web service policies, JPS configuration) from a reference installation, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **SOA Artifacts Provisioning** deployment procedure. Then click **Schedule Deployment**. Enterprise Manager Grid Control displays the Source Selection page of the Deployment Procedure.

> **Note:** You can also access this deployment procedure as follows:
>
> ■ From the SOA Infrastructure Home page:
>
> > 1. From the Grid Home page, navigate to the Targets page.
> >
> > 2. In the Targets page, click **Middleware**.
> >
> > 3. In the Middleware page, click on a target of type **SOA Infrastructure**.
> >
> > 4. In the SOA Infrastructure home page, from the SOA infrastructure-specific menu, select **SOA Artifacts Provisioning**.

4. On the Source Selection page, in the Source Selection section, do the following:

   a. Retain the default selection, that is, **Provision from reference environment**.

   b. Click on the torch icon against the **Domain Name** field. Search for the Oracle WebLogic Server Domain that you want to deploy the SOA artifacts from and select it. Ensure that the source Oracle WebLogic Server Domain is up and running.

   c. In the **Credentials** section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

   To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. You are prompted to specify the Oracle WebLogic Server Domain credentials and the Oracle WebLogic Administration Server host credentials. In the Oracle WebLogic Server Domain Credentials section, specify the administrator credentials that can be used to access the WebLogic Server Administration Console. In the Oracle WebLogic Administration Server Host Credentials section, specify the operating system credentials of the user who installed the Admin Server.

   d. Optionally, if you want to save the SOA artifacts as an image in the Software Library, select **Save SOA Artifacts Gold Image in Software Library**.

   For example, in future, if you want to provision this particular version to other Oracle WebLogic Server Domains, then instead of using the reference installation, which could potentially be down, you can use the gold image you saved in the Software Library.

   e. Click **Next**.

5. On the Select Destination page, do the following:

   a. Click on the torch icon against the **Domain Name** field. Search for the Oracle WebLogic Server Domain that you want to deploy the SOA artifacts to and select it. Ensure that the destination Oracle WebLogic Server Domain is up and running.

   b. In the **Credentials** section, retain the default selection, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

   To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. You are prompted to specify the Oracle WebLogic Server Domain credentials and the Oracle WebLogic

Administration Server host credentials. In the Oracle WebLogic Server Domain Credentials section, specify the administrator credentials that can be used to access the WebLogic Server Administration Console. In the Oracle WebLogic Administration Server Host Credentials section, specify the operating system credentials of the user who installed the Admin Server.

   c. Click **Next**.

6. On the Select Artifacts page, do the following:

   a. In the **Choose the type of SOA artifacts to provision** section, select **SOA Composites**, **Web Services Policies**, and **Java Platform Security Configuration**.

   b. Click **Next**.

7. On the SOA Composites page, do the following:

   a. Select the composites you want to provision and specify a configuration plan from the Software Library or a directory.

      It is recommended that you place the configuration plan in a directory on destination machine. Alternatively, you can also place the configuration plan in any other shared location which is accessible from the destination machine.

      If the composite already exists on the destination host, then select **Overwrite** to overwrite that existing composite with the composite from the source domain.

   b. Click **Next**.

8. On the Web Services Policies page, do the following:

   a. In the Assertion Templates section, select the assertion templates to migrate.

   b. In the Web Services Policies section, select the policies to migrate.

   c. Click **Next**.

9. On the Java Platform Security page, do the following:

   a. In the Migrate Policy Store and Credential Store section, select **Migrate Policy Store** and **Migrate Credential Store** check boxes.

      To view a list of providers for the source and target, click **Provider details** link.

   b. Click **Next**.

10. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

11. On the Review page, review the details you have provided for provisioning SOA artifacts, and click **Submit**.

14.3.2 Provisioning SOA Artifacts from Gold Image

This section describes how you can provision SOA artifacts (composites, web service policies, JPS configuration) from a gold image stored in the Software Library. In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

14.3.2.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Ensure that you have already provisioned Oracle SOA Suite 11g and its underlying Oracle WebLogic Server Domain.
- Ensure that the source and the destination soa-infra domains are of the same version.
- Ensure that you have already saved the gold image in the Software Library while provisioning the SOA artifacts from a reference installation.

14.3.2.2 Provisioning Procedure

To provision SOA artifacts from a gold image, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **SOA Artifacts Provisioning** deployment procedure. Then click **Schedule Deployment**. Enterprise Manager Grid Control displays the Source Selection page of the Deployment Procedure.

   > **Note:** You can also access this deployment procedure as follows:
   >
   > - From the SOA Infrastructure Home page:
   >   1. From the Grid Home page, navigate to the Targets page.
   >   2. In the Targets page, click **Middleware**.
   >   3. In the Middleware page, click on a target of type **SOA Infrastructure**.
   >   4. In the SOA Infrastructure home page, from the SOA infrastructure-specific menu, select **SOA Artifacts Provisioning**.

4. On the Source Selection page, do the following:

   a. Select **Provision from Gold Image**.

   b. Click on the torch icon against the **Gold Image Name** field. Search for the gold image you want to provision the SOA artifacts from and select it.

   c. Click **Next**.

5. On the Select Destination page, do the following:

   a. Click on the torch icon against the **Domain Name** field. Search for the Oracle WebLogic Server Domain that you want to deploy the SOA artifacts to and select it.

b. In the **Credentials** section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. You are prompted to specify the Oracle WebLogic Server Domain credentials and the Oracle WebLogic Administration Server host credentials. In the Oracle WebLogic Server Domain Credentials section, specify the administrator credentials that can be used to access the WebLogic Server Administration Console. In the Oracle WebLogic Administration Server Host Credentials section, specify the operating system credentials of the user who installed the Admin Server.

c. Click **Next**.

6. On the Select Artifacts page, do the following:

a. In the **Choose the type of SOA artifacts to provision** section, select **SOA Composites**, **Web Services Policies**, and **Java Platform Security Configuration**.

b. Click **Next**.

7. On the SOA Composites page, do the following:

a. Select the composites you want to provision and specify a configuration plan from the Software Library or a directory.

If the composite already exists on the destination host, then select **Overwrite** to overwrite that existing composite with the composite from the source domain.

b. Click **Next**.

8. On the Web Services Policies page, do the following:

a. In the Assertion Templates section, select the assertion templates to migrate.

b. In the Web Services Policies section, select the policy assertions to migrate.

c. Click **Next**.

9. On the Java Platform Security page, do the following:

a. In the Migrate Policy Store and Credential Store section, select **Migrate Policy Store** and **Migrate Credential Store** checkboxes.

To view a list of providers for the target, click **Provider details** link.

ORACLE Enterprise Manager 11g Grid Control

SOA Artifacts Provisioning

SOA Composites   Web Services Policies   **Java Platform Security**   Schedule   Review

Java Platform Security

Use this page to migrate policy and credential store. Click the link to view current providers. Provider details

Migrate Policy Store and Credential Store

Select to migrate policy store and credential store.

☐ Migrate Policy Store
☐ Migrate Credential Store

b. Click **Next**.

10. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

11. On the Review page, review the details you have provided for provisioning SOA artifacts, and click **Submit**.

14.4 Deploying SOA Composites

This section explains how you can deploy SOA composites. In particular, this section contains:

- Prerequisites
- Provisioning Procedure

14.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that you have already provisioned Oracle SOA Suite 11g and its underlying Oracle WebLogic Server Domain.

- Ensure that the source and the destination soa-infra domains are of the same version.

  The domain should have at least one managed server with the SOA Infrastructure application running. In the case of a SOA Cluster, the composites will be deployed to any one managed server in the cluster.

- Ensure that you have the SOA Composites either in the Software Library or in a file system accessible from the Admin Server host.

14.4.2 Provisioning Procedure

To provision SOA composites, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Deploy SOA Composites** deployment procedure. Then click **Schedule Deployment**. Enterprise Manager Grid Control displays the Destination Selection page of the Deployment Procedure.

4. On the Destination page, do the following:

   a. Click on the torch icon against the **Destination Domain Name** field. Search for the Oracle WebLogic domain that you want to deploy the SOA composites to, and select it.

   b. In the **Credentials** section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. You are prompted to specify the Oracle WebLogic Server Domain credentials and the Oracle WebLogic Administration Server host credentials. In the Oracle WebLogic Server Domain Credentials section, specify the administrator credentials that can be used to access the WebLogic Server Administration Console. In the Oracle WebLogic Administration Server Host Credentials section, specify the operating system credentials of the user who installed the Admin Server.

   c. Click **Next**.

5. On the Source page, do the following:

   a. In the Composites section, click **Add**. Select Composites Source as Software Library or File System depending on where the composites are located. Select the Plan Source location and path.

     If the composite already exists on the destination host, then select **Overwrite** to overwrite that existing composite with the composite from the source domain. If you want the composite that you are deploying now to be set as the default component, then retain the **Force Default** selection.

   b. In the Options section, select **Verify adapter dependencies** if you want to ignore the missing adapters in the destination domain and proceed with the provisioning operation. Each composite may refer to one or more adapters, and the composites may not run properly if the depending adapters are missing in the destination domain. However, if you select this option, you can ignore all such missing adapters.

   c. Click **Next**.

6. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

7. On the Review page, review the details you have provided for provisioning SOA composites, and click **Submit**.

15

Provisioning Coherence Nodes and Clusters

This chapter explains how you can provision Coherence nodes or clusters across multiple targets in a farm. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedure

- Supported Releases

- Deploying a Coherence Node

- Troubleshooting

15.1 Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Coherence nodes and clusters. The Coherence Deployment Procedure allows you to do the following:

- Add one or more nodes to a new cluster and add this cluster as an Enterprise Manager monitored target.

- Add a management node to an existing cluster and add this cluster as an Enterprise Manager monitored target.

- Add one or more nodes to a cluster that is already being monitored by Enterprise Manager.

Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully deploy a Coherence node. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 15–1   Getting Started with Deploying a Coherence Node

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for deploying a Coherence node. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 15.2, "Deployment Procedure".

Table 15–1 (Cont.) Getting Started with Deploying a Coherence Node

Step	Description	Reference Links
Step 2	**Knowing About the Supported Releases** Know what releases of Oracle Coherence are supported.	To learn about the releases supported by the Deployment Procedure, see Section 15.3, "Supported Releases".
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as adding the Coherence node as a target, and setting up of Oracle Software Library.	To learn about the prerequisites for deploying a Coherence node, see Section 15.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully deploy a Coherence node.	To deploy a Coherence node, follow the steps explained in Section 15.4.2, "Procedure".

15.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure:

- *Coherence Node Provisioning*

15.2.1 Coherence Architecture

Before you can monitor a coherence node, you must identify the management node in the cluster and add its details in Enterprise Manager. The following diagram shows the Coherence architecture:

Figure 15–1 Architecture Diagram: Monitoring the Coherence Cluster
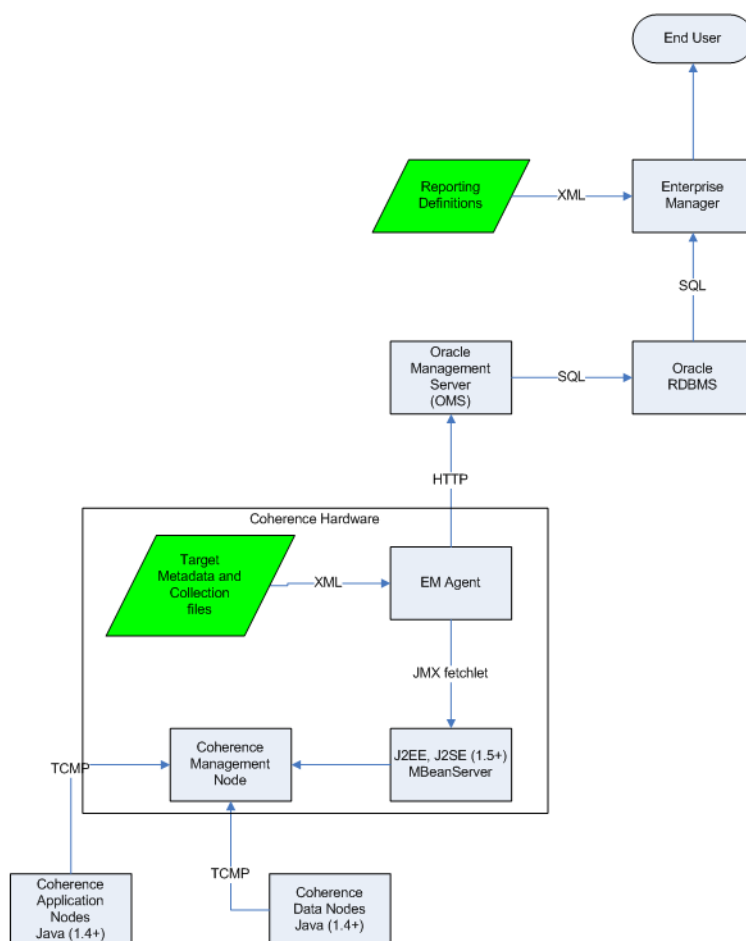
The management node hosts the Coherence MBeanServer. All Coherence MBeans, Enterprise Manager MBeans for Coherence, and customer specific MBeans are registered in this MBeanServer. Since the number of MBeans may be large, we recommend that you use the BulkOperationsMBean to upload data from all the MBeans quickly and efficiently.

When the Enterprise Manager Agent is installed, the $ORACLE_ HOME/sysman/jlib/coherenceEMIntg.jar and the $ORACLE_ HOME/modules/bulkoperationsmbean_11.1.1.jar files are also installed. The bulkoperationsmbean_11.1.1.jar contains all required classes needed for the BulkOperationsMBean. The coherenceEMIntg.jar contains the oracle.sysman.integration.coherence.EMIntegrationServer class that starts a new Coherence management node with the MBeanServer, registers all the MBeans in this MBeanServer along with the BulkOperationsMBean. We recommend that you add this new management node using the BulkOperationsMBean. You can monitor a cluster without the BulkOperationsMBean but data collection may be slow and inefficient.

> **Note:** You can deploy a node on remote hosts only if these remote hosts are being monitored by Enterprise Manager. If a host is not monitored by Enterprise Manager, you cannot create new nodes on that host.

15.3 Supported Releases

This section lists the releases of Oracle Coherence supported by this Deployment Procedure:

- Oracle Coherence 3.3, 3.4, 3.5, and 3.6

15.4 Deploying a Coherence Node

This section describes how you can deploy a Coherence node using the Coherence Node Provisioning procedure.

This section covers the following:

- Prerequisites
- Procedure

15.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- The host on which a Coherence node needs to be deployed should be a monitored target in Oracle Enterprise Manager.

- A software library with the necessary components that include Coherence jars, configuration files, and application specific jars must be created. We recommend that you add all software components with the product name **Coherence**.

- A zip file with the Coherence software, default configuration files and start scripts must be created. This zip file must be added as a software component to the Enterprise Manager software library.

- To facilitate the creation of the management node, you must add `coherenceEMIntg.jar` and `bulkoperationsmbean_11.1.1.jar` as software components to the library.

15.4.2 Procedure

To deploy a Coherence cluster, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployments page, click **Deployment Procedures**.

3. All the available deployment procedures are listed on the Deployment Procedure Manager page. Click the **Coherence Node Provisioning** deployment procedure on the list.

4. The View Procedure page is displayed. Click **Schedule Deployment**.

5. The Source Selection page which is the first page of the Coherence Node Provisioning wizard is displayed. On this page, do the following:

Figure 15–2   Source Selection Page

1. Click **Add** to add a software component. The Select Source popup is displayed. By default all software components with the product name Coherence are displayed. Select required components and click **Select**.

2. All selected components appear in Sources table in the Source Selection page. Specify the destination directory information for all selected components. The destination directory must be an absolute or relative path from COHERENCE_ HOME (contains all the Coherence related files).

> **Note:**   You can skip this step if all the required software components are available on the target machine.

---

3. Click **Next** to continue with the next step in the wizard.

6. The Target Selection page is displayed. On this page, you can add one or more nodes that need to be deployed. You can add nodes to:

   1. **EM Monitored Coherence Cluster Target**: Click the **Search** icon and select a target from the list.

   2. **New EM Coherence Cluster Target:** This option allows you to add a new Coherence target. Enter the following details:

Figure 15–3   Target Selection Page

Table 15–2   Target Selection Page

Field Name	Description
Cluster Name	Enter the name of the cluster. This is a mandatory field.
Cluster Port / Address	If you do not specify the values for the Cluster Port and Cluster Address, the values in the `tangosol_coherence.xml` configuration file are used.
Coherence Version	The default version used is 3.3 but you can specify 3.4 or later.

3.   Click **Add** to add nodes to the cluster. The Add Coherence Node page is displayed. Enter the following details:

Figure 15–4   Add Coherence Node Page

Table 15–3   Add Coherence Node Page

Field Name	Description
Host Details	
Host Name	Select the host on which the node is to be added.You may have more than one node on the host depending on the machine configuration and the node configuration.
	Note: The Host Name you select here is used to set two start up parameters. Use the `tangosol_coherence_machine` environment variable to set the `tangosol.coherence.machine` parameter and `oracle_coherence_machine` to set the `oracle.coherence.machine` parameter.
Number of Nodes	Specify the number of nodes that need to be added. By default, this field has the value of 1 but you can add as many nodes as required. If the value is more than 1, then all the nodes will have the following properties:
	■ Each node will use the same `COHERENCE_HOME` and start script.
	■ The Node Name value will be added as the prefix and a number will be appended to each node. For example `<node_name>_1`, `<node_name>_2` and so on. Each name should be a unique one in the cluster.
	■ The JMX Remote Port value will be increased by 1 for each additional node. For example, if the value of the JMX Remote Port for the first node is 8088, the value for the second node will be 8089 and so on.

Table 15–3   (Cont.) Add Coherence Node Page

Field Name	Description
Credentials	
Use Preferred Credentials	Select this option, if you already entered the credentials on the Preferred Credentials page under Preferences.
Override Preferred Credentials	Select this option to explicitly specify the user name and password and override the preferred credentials.
Node Details	
Node Name	Enter a unique name for the node
Site Name	This is the location of the Coherence node. This is geographical physical site name which identifies the racks and machines on which the node is running.
Rack Name	The name of the rack in the site on which the machine is located.
Role Name	The role could be storage/data, application/process, proxy or management node. **Note:** The Node Name, Site Name, Rack Name, and Role Name cannot exceed 32 characters.
Do not Copy Software Components	If you are creating more than one node with same software components on a host, you need not copy these components again. By selecting this option, you can skip copying of software components.
Management Node Details	
Management Node with MBeanServer	You can define multiple management nodes in the cluster but only one management node can be marked as the Primary Management Node. We recommend that you add at least two management nodes preferably running on different hosts / machines to support fail over.
JMX Remote Port	The port number of the EMIntegration Mbean server.
JMX User Name	The user name for the JMX server if authentication is enabled.
JMX Password	The password for the JMX server if authentication is enabled. **Note:** To enable the JMX authentication, you need to set `com.sun.management.jmxremote.authenticate=true`. The JMX User name and JMX Password need to be set in the `$JDK_HOME/jre/lib/management/jmxremote.password` and `$JDK_HOME/jre/lib/management/jmxremote.access` files.
Primary Management Node used for Monitoring	Select this checkbox to mark the management node you are adding as the Primary Management Node used for Monitoring. If several nodes are being added to a cluster, only one management node can be marked as the primary one. If the primary management node fails, you can configure any of the other management nodes for monitoring. If no other management node is available, you can add a new primary management node to an existing cluster and this node can be used to monitoring. **Note**: To configure the management node, navigate to the Monitoring Configuration by clicking this link under the Related Links in the Cluster Home page.
Use Bulk Operations MBean	This checkbox is selected by default. When this option is selected and `bulkoperationsmbean_11.1.1.jar` and the `coherenceEMIntg.jar` files have been added to classpath, a new management node with BulkOperationsMBean will be started.

Table 15–3   (Cont.) Add Coherence Node Page

Field Name	Description
Environment Details	
Coherence Home	Enter the absolute path to the folder under which the Coherence software components reside. The path specified here will be used as the Destination Directory specified on the Coherence Node Provisioning: Source Selection page. This value could be different for each node or the same for one or more nodes.
Start Script	This script is used to bring up the Coherence node. This script is operating system specific and sets the proper environment required for the node by specifying the relevant system parameters. See sample script for an example.

The following table lists the environment variables and their values that can be used to set the Coherence System Parameters in the start script. The deployment procedure sets these environment variables with the inputs you specify in the interview wizard. You can execute a separate script to set different values for these environment variables in the start script. You can use these environment variables to set start parameters for the Coherence node or override them by specifying appropriate values.

Table 15–4   Environment Variables

UI Parameter	Environment Variable Set	Coherence System Parameter
COHERENCE_ HOME	coherence_home	oracle.coherence.home
JAVA_HOME	JAVA_HOME	JAVA_HOME
APP_HOME	APP_HOME	APP_HOME
ORACLE_HOME	ORACLE_HOME	ORACLE_HOME
CLUSTER_ NAME	tangosol_coherence_ cluster	tangosol.coherence.cluster
CLUSTER_ ADDRESS	tangosol_coherence_ clusteraddress	tangosol.coherence.clusteraddr ess
CLUSTER_ PORT	tangosol_coherence_ clusterport	tangosol.coherence.clusterport
NODE_NAME (32 chars)	tangosol_coherence_ member	tangosol.coherence.member
SITE_NAME (32 chars)	tangosol_coherence_site	tangosol.coherence.site
RACK_NAME (32 chars)	tangosol_coherence_rack	tangosol.coherence.rack
MACHINE_ NAME	tangosol_coherence_ machine	tangosol.coherence.machine
ROLE_NAME (32 chars)	tangosol_coherence_role	tangosol.coherence.role
JMX_REMOTE_ PORT	jmx_remote_port	com.sun.management.jmxremote.p ort
PROCESS_ NAME	tangosol_coherence_ process	tangosol.coherence.process
BULK_MBEAN	bulk_mbean	

Table 15–4   (Cont.) Environment Variables

UI Parameter	Environment Variable Set	Coherence System Parameter
JMX_ENABLE_ AUTH	jmx_enable_auth	com.sun.management.jmxremote.a uthenticate
START_ SCRIPT	start_script	oracle.coherence.startscript

4.   After adding the node, click **Continue** to return to the Coherence Node Provisioning: Target Selection page.

5.   Click **Next** to go to the next step in the wizard.

7.   The Schedule page is displayed. On this page, you can specify the schedule for deploying the node. You can choose to deploy the node immediately or at a later date.

Figure 15–5   Schedule Page

> **Note:** If you set the Grace Period as **Indefinite**, Enterprise Manager will keep trying to deploy the node for an indefinite period. If you specify a date / time in this field, the deployment process will be aborted after this period.

8.   Click **Next** to go to the next step in the wizard.

9.   The Review page is displayed. On this page, review the details you have provided for deploying the node. Click **Finish** to deploy the node.

> **Note:** Depending on the schedule you have defined, the deployment will take place either immediately or at a later date as specified.

15.4.2.1 Sample Script

A sample script to start a coherence management node is shown below. When this script is executed from Deployment Procedure, it sets some environment variables. If you want to override those variables you may run a local script as part of this script and override those settings. If the JMX Remote Port is defined, then this would create

a Coherence management node with a MBeanServer. A BulkOperationsMBean also
will be registered.

Sample Script (Unix): You can use the `default-start-script.pl` script:

```perl
#!/usr/local/bin/perl
#
# Sample script to demonstrate starting of following Coherence nodes.
# When this script is passed in as a start script in Coherence Node Provisioning
# Deployment Procedure, while executing start node step, deployment procedure sets
# all user entered options as environment variables. Based on the values of these
# environment variables, you can start different types of Coherence nodes.
#
# - Management Node with Oracle Bulk Operation MBean is started when
#   "bulk_mbean" and "jmx_remote_port" variables are set. For this option,
#   oracle.sysman.integration.coherence.EMIntegrationServer Java class is executed
#     that starts a MBeanServer in this node and registers Oracle Bulk Operation
#     MBean. You need coherenceEMIntg.jar and bulkoperationsmbean_11.1.1.jar in the
#     classpath. Oracle recommends the usage of Bulk Operation MBean as it improves
#     performance of metric collection.
#
# - Management Node is started when "jmx_remote_port" is set, but "bulk_mbean" is
#     NOT set.
#
# - Managed node when "jmx_remote_port" is not set.
#
#
# Following variables are set from the deployment procedure. Use these values to
# define required system parameters to override Coherence default settings.
#

my $coherence_home=$ENV{'COHERENCE_HOME'};
my $start_script=$ENV{'START_SCRIPT'};
my $java_home=$ENV{'JAVA_HOME'};
my $app_home=$ENV{'APP_HOME'};
my $oracle_home=$ENV{'ORACLE_HOME'};

my $member=$ENV{'tangosol_coherence_member'};
my $site=$ENV{'tangosol_coherence_site'};
my $rack=$ENV{'tangosol_coherence_rack'};
my $machine=$ENV{'tangosol_coherence_machine'};

# tangosol.coherence.machine has a limitation of 32 chars. As a workaround, use
# oracle.coherence.machine to set machine name. This parameter is used to identify
# hosts for cluster management features
my $oracle_coherence_machine=$ENV{'oracle_coherence_machine'};
my $role=$ENV{'tangosol_coherence_role'};

my $jmxport=$ENV{'jmx_remote_port'};
my $cluster=$ENV{'tangosol_coherence_cluster'};
my $clusteraddr=$ENV{'tangosol_coherence_clusteraddress'};
my $clusterport=$ENV{'tangosol_coherence_clusterport'};
my $bulkmbean=$ENV{'bulk_mbean'};
my $jmx_auth=$ENV{'jmx_enable_auth'};

my $SYS_OPT="";
my $JVM_OPT="";

my $psep="";
my $dsep="";
if ( !&IsWindows() ) {
```

```
        $psep=":";
        $dsep="/";
}
else
{
        $psep=";";
        $dsep="\\";
}

my $cohlogfile=$coherence_home.$dsep.$member.".cohout";

my $logfile=$coherence_home.$dsep.$member.".log";
print "$logfile\n";
open(LOGFL,"> $logfile");

print LOGFL "Starting Node : $member\n";
print LOGFL "Coherence Home : $coherence_home \n";
print LOGFL "Start Script : $start_script \n";
print LOGFL "Java Home : $java_home \n";
print LOGFL "Oracle Home : $oracle_home \n";
print LOGFL "APP Home : $app_home \n";

print LOGFL "Site Name : $site \n";
print LOGFL "Rack Name : $rack \n";
print LOGFL "Machine Name : $machine \n";
print LOGFL "Oracle Coherence Machine Name : $oracle_coherence_machine \n";
print LOGFL "Role Name : $role \n";

print LOGFL "Cluster Name : $cluster \n";
print LOGFL "Cluster Addr : $clusteraddr \n";
print LOGFL "Cluster Port : $clusterport \n";
print LOGFL "JMX Port : $jmxport \n";
print LOGFL "Bulk MBean : $bulkmbean \n";
print LOGFL "JMX Auth Enabled : $jmx_auth \n";

#
# you may run a local script as part of this script and override those
# settings.
# Override JAVA_HOME variable by setting it locally
#
#. ./set-env.sh
#echo "After setting JAVA_HOME locally, JAVA_HOME: $JAVA_HOME"

# Options for Java Virtual Machine.
$JVM_OPT="-server -Xms512m -Xmx512m -Xincgc -verbose:gc";

#
# Set system parameters to Coherence node
$SYS_OPT="-Djava.net.preferIPv4Stack=true";

# This param allows the mbeans on this node to be registered to mbean servers
running on management nodes
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.management.remote=true";

$SYS_OPT="$SYS_OPT -Dcom.sun.management.jmxremote.ssl=false";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.cluster=$cluster";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.member=$member";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.site=$site";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.rack=$rack";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.machine=$machine";
```

```
# set this if machine name > 32 chars, tangosol.coherence.machine has a limitation
of 32 chars
$SYS_OPT="$SYS_OPT -Doracle.coherence.machine=$oracle_coherence_machine";
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.role=$role";

# Set coh home and start script so they will be part of input args
$SYS_OPT="$SYS_OPT -Doracle.coherence.home=$coherence_home";
$SYS_OPT="$SYS_OPT -Doracle.coherence.startscript=$start_script";

# set jmxremote.authenticate=true if $jmx_enable_auth is present.
# username/password needs to be set in $JDK_
HOME/jre/lib/management/jmxremote.password and
# $JDK_HOME/jre/lib/management/jmxremote.access files. Uncomment the following
block after adding these files.
#
#if ("$jmx_auth" ne "") {
#       $SYS_OPT="$SYS_OPT -Dcom.sun.management.jmxremote.authenticate=$jmx_auth";
#} else {
#       $SYS_OPT="$SYS_OPT -Dcom.sun.management.jmxremote.authenticate=false";
#}

# Default is true, so make sure to set it to false if not using authentication
        $SYS_OPT="$SYS_OPT -Dcom.sun.management.jmxremote.authenticate=false";


# set jmxremote.port only for management nodes, if user passes in.
# It enables monitoring from remote systems through this port.
if ($jmxport ne  "") {
        $SYS_OPT="$SYS_OPT -Dcom.sun.management.jmxremote.port=$jmxport";
}


#
# Define clusteraddress and clusterport if we have valid values.
#
if($clusteraddr ne "") {
        $SYS_OPT="$SYS_OPT -Dtangosol.coherence.clusteraddress=$clusteraddr";
}

if("$clusterport" ne "") {
        $SYS_OPT="$SYS_OPT -Dtangosol.coherence.clusterport=$clusterport";
}

my $startup_class="";
my $cmd="";

# Note that Coherence lib is under $COHERENCE_HOME/coherence. Add any application
specific jars to this classpath, if needed.

my $CLASSPATH=$coherence_
home.$dsep."coherence".$dsep."lib".$dsep."coherence.jar".$psep.$coherence_
home.$dsep."coherence".$dsep."lib".$dsep."reporter.jar";

print LOGFL "CLASSPATH: $CLASSPATH\n";

if ($bulkmbean ne "" && $jmxport ne "") {
        # Management node with Bulk Operation MBean.
        # add Oracle supplied jars for Bulk Operation MBean

$CLASSPATH=$CLASSPATH.$psep.$coherence_
```

```
home.$dsep."coherenceEMIntg.jar".$psep.$coherence_home.$dsep."bulkoperationsmbean_
11.1.1.jar";
# Start MBeanServer
     $SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=all";
      print LOGFL "Starting a management node with Bulk Operation MBean \n";
     $startup_class="oracle.sysman.integration.coherence.EMIntegrationServer";
     $cmd=$java_home.$dsep."bin".$dsep."java -cp $CLASSPATH $JVM_OPT $SYS_OPT
$startup_class > $cohlogfile 2<&1";
} elsif ($jmxport ne "") {
        # Management Node with out Bulk Operation MBean
        # Start MBeanServer
        $SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=all";
        print LOGFL "Starting a management node ...\n";
        $startup_class="com.tangosol.net.DefaultCacheServer";
$cmd=$java_home.$dsep."bin".$dsep."java -cp $CLASSPATH $JVM_OPT $SYS_OPT $startup_
class > $cohlogfile 2<&1";
} else {
        # A simple managed node. Do not start MBeanServer.
        $SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=none";
        print LOGFL "Starting a simple managed node ...\n";
        $startup_class="com.tangosol.net.DefaultCacheServer";
$cmd=$java_home.$dsep."bin".$dsep."java -cp $CLASSPATH $JVM_OPT $SYS_OPT $startup_
class > $cohlogfile 2<&1";
}


if ( !&IsWindows() ) {
    if (fork() == 0) {
        print LOGFL "Executing start script from child process... $cmd \n";
        exec("$cmd") or die "Could not execute $cmd\n";
    }
} else {
        print LOGFL "Command used to start node = $cmd\n";
        exec($cmd);
}
close(LOGFL);
exit 0;

sub lsWindows {
    $osname = $^O;
    if (   $osname eq "Windows_NT"
        || $osname eq "MSWin32"
        || $osname eq "MSWin64" )
    {
        return 1;
    }
    else {
        return 0;
    }
}
```

Sample Script (Windows): You need to run the start-node.vbs script which in turn launches the default-start-script.pl script.

start-node.vbs

```
Set WshShell = WScript.CreateObject("WScript.Shell")
Set objEnv = WshShell.Environment("Process")
Dim run_cmd
run_cmd=objEnv("ORACLE_HOME") & "\perl\bin\perl.exe " & objEnv("ORACLE_HOME") &
"\sysman\admin\scripts\coherence\default-start-script.pl"
```

```
obj = WshShell.Run(run_cmd, 0)
set WshShell = Nothing
```

15.5 Troubleshooting

We recommend that you have at least two management nodes running on different machines in a Coherence cluster. If a monitoring failure occurs, the second management node can be used. Some of the common failure scenarios are listed below:

- **Error Condition**: Loss of Management Node

  Solution: If the primary management node fails, you need to change the monitoring configuration of the cluster to point to another management node in the cluster. If no other management node is present in the cluster, you can use the Coherence Node Provisioning deployment procedure, select an existing cluster and add a new management node. This process will update the monitoring configuration for the cluster.

- **Error Condition**: Loss of Agent Monitoring the Cluster

  Solution: If the Agent is not available, you need to use another Enterprise Manager Agent to point to the management node of Coherence cluster.

- **Error Condition**: Loss of Host with EM Agent and Management Node

  Solution: If the Host is not available, you need to switch to another management node that is running on a different machine.

16

Provisioning Oracle Service Bus Resources

Oracle Service Bus is an enterprise-class service bus that connects, manages, and mediates interactions between heterogeneous services. Oracle Service Bus accelerates service configuration, integration, and deployment, thus simplifying management of shared services across the Service-Oriented Architecture (SOA). For information about discovering and monitoring Oracle Service Bus, see the *Enterprise Manager Getting Started with Management Pack for SOA Guide* available at:

http://www.oracle.com/technology/documentation/oem.html

The resources of Oracle Service Bus can be organized into individual projects. Projects are non-hierarchical, disjointed, top-level grouping constructs. All resources (such as business services, proxy services, WS-Policies, WSDLs, schemas, XQuery transformations, JARs, and so on) reside in exactly one non-overlapping project. Resources can be created directly under a project or be further organized into folders. Folders may be created inside projects or inside other folders, and the folders are similar to directories in a file system, with the project level being the root directory.

While Enterprise Manager Grid Control allows you to discover and monitor these Oracle Service Bus targets, it also provides Deployment Procedures that help you provision Oracle Service Bus resources.

This chapter explains how you can provision Oracle Service Bus resources. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedure

- Supported Releases

- Provisioning Oracle Service Bus Resources from Oracle Service Bus Domain

- Provisioning Oracle Service Bus Resources from Oracle Software Library

16.1 Getting Started

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

Table 16–1   Getting Started with Provisioning Oracle Service Bus Resources

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Service Bus resources. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 16.2, "Deployment Procedure".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle Service Bus can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 16.3, "Supported Releases".
Step 3	**Selecting the Use Case** This chapter covers a few use cases for provisioning Oracle Service Bus resources. Select the use case that best matches your requirement.	■ To learn about provisioning Oracle Service Bus resources from the an Oracle Service Bus domain, see Section 16.4, "Provisioning Oracle Service Bus Resources from Oracle Service Bus Domain". ■ To learn about provisioning Oracle Service Bus resources from the Software Library, see Section 16.5, "Provisioning Oracle Service Bus Resources from Oracle Software Library".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	■ To learn about the prerequisites for provisioning Oracle Service Bus resources from the an Oracle Service Bus domain, see Section 16.4.1, "Prerequisites". ■ To learn the prerequisites for provisioning Oracle Service Bus resources from the Software Library, see Section 16.5.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Service Bus resources.	■ To provision Oracle Service Bus resources from the an Oracle Service Bus domain, follow the steps explained in Section 16.4.2, "Provisioning Procedure". ■ To provision Oracle Service Bus resources from the Software Library, follow the steps explained in Section 16.5.2, "Provisioning Procedure".

16.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for provisioning Oracle Service Bus resources:

■   *Oracle Service Bus Resource Provisioning*

16.3 Supported Releases

Using this Deployment Procedure, you can provision the resources for Oracle Service Bus 2.6, 2.6.1, 3.0, and 10gR3 (3.1).

16.4 Provisioning Oracle Service Bus Resources from Oracle Service Bus Domain

This section describes how you can provision Oracle Service Bus resources directly from an Oracle Service Bus domain.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

16.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that the source Oracle Service Bus (from where you want to export the resources) is already discovered and monitored in Enterprise Manager Grid Control.

- If you want to use a customization file to customize the environment variables in the changed (target) environment, then you must ensure that the customization file is available as a generic component in Oracle Software Library. For instructions to create generic components, see Section 3.1.2, "Setting Up Oracle Software Library".

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

16.4.2 Provisioning Procedure

To provision Oracle Service Bus resources from a source Oracle Service Bus domain, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Service Bus Resource Provisioning**. Then click **Schedule Deployment**.

4. On the Select Source page, in the Source section, select **Oracle Service Bus Domain**.

   a. For **Domain**, click the torch icon and select the Oracle Service Bus domain from where the resources can be exported and deployed to a target Oracle Service Bus domain. In the following page of the wizard, you will be allowed to select the domain's projects that you want to export.

   b. For **BEA Home Directory**, specify the full path to the BEA home directory where all BEA product-related files are stored. For example, `/home/mark/bea`.

   c. Click **Next**.

5. On the Select Projects page, do the following:

   a. In the Resource Summary section, select the projects you want to export and deploy to the target OSB domain. The selected projects are exported to a JAR file, and the JAR file is moved to the host where the target OSB domain is running.

   Note that the resources of the selected projects that exist in the target OSB domain but not in the exported JAR file will be deleted.

   b. In the Export Mode section, do one of the following:

   Select **Export Projects** if you want to export the resources at project level. While deploying the exported JAR file to the target host, the entire project is deployed. This may add, overwrite, or delete resources depending on the availability of resources on the target host.

Select **Export Resources** if you want to export the resources at resource level. While deploying the exported JAR file to the target host, only the resources are deployed. This may add or overwrite resources depending on the availability of resources on the target host.

To understand these options better, read the use cases described inSection 16.4.2.1, "Understanding Export Modes".

c. (Optional) In the Security Options section, if the projects you want to export contain any resources with sensitive data, then specify a pass-phrase to protect them. The same pass-phrase will be used to import the protected resources during deployment.

d. (Optional) In the Save Projects to Software Library section, select **Save Projects to Software Library** and specify a component name and location if you want to save the exported project JAR file as a generic component in the Software Library.

By default, the projects you select here are exported to a JAR file and moved to the host where the Administration server of the target Oracle Service Bus domain is running. However, the JAR files are not saved in the Software Library for future use. Using this option, you can save them as a component in the Software Library.

6. On the Select Target page, do the following:

a. In the Target section, specify the following:

For **Domain**, click the torch icon and select the Oracle Service Bus domain where you want to deploy the selected resources.

For **BEA Home Directory**, specify the full path to the BEA home directory where all BEA product-related files are stored.

b. (Optional) In the Advanced Options section, select the settings you want to retain if you have done some customization to the resources selected for deployment, and if you want to preserve those changes in the target Oracle Service Bus domain.

Note that for Oracle Service Bus 2.6.x, Security and Policy Configuration, Credentials, and Access Control Policies cannot be preserved.

c. In the Customization section, provide details about the customization file that can be used to modify the environment settings in the target OSB domain.

If you do not want to use a customization file, select **None**.

If you are using a customization file and if it is available on the host where the target OSB domain is running, then select **Use the Customization file on the target host** and specify the full path to the location where the file is present.

If the customization file is stored as a generic component in Oracle Software Library, then select **Select the customization file from the Software Library** and specify the full path to the location in Oracle Software Library where the generic component is stored.

d. Click **Next**.

7. On the Set Credentials page, specify the following and click **Next**.

a. Specify the login credentials of the source and target Oracle Service Bus (OSB) domains.

b. Specify the credentials of the hosts where the Management Agents, which are monitoring the administration servers of the OSB domains, are running

8. On the Schedule page, schedule the deployment procedure to run immediately or later, and specify a unique name for the deployment procedure so that it can be tracked. Click **Next**. Enterprise Manager Grid Control displays the Review page.

9. On the Review page, review the details you have provided for provisioning Oracle Service Bus resources, and click **Submit**.

16.4.2.1 Understanding Export Modes

The following describes the different use cases and explains how the export modes will work for those circumstances.

While the first column shows the project selected from the source domain and the resources contained in that selected project, the second column shows the availability of that project in the target domain. And, while the third column shows how Export at Project Level work, the fourth column shows how Export at Resource Level works.

Table 16–2   Understanding Export Modes

Source Domain	Target Domain	Export at Project Level	Export at Resource Level
You have selected Project_1 from the source domain, and this project has Resource_1, Resource_2, and Resource_3.	The target domain has no projects at all.	The entire Project_1 will be deployed to the target domain.	The entire Project_1 will be deployed to the target domain.
You have selected Project_1 from the source domain, and this project has Resource_1, Resource_2, and Resource_3.	The target domain has Project_1, and this project has Resource_1.	The entire Project_1 will be deployed to the target domain, wherein, Resource_1 will be overwritten because it is already available in the target domain, and Resource_2 and Resource_3 will be ADDED.	Only the resources of Project_1 will be deployed to the target domain, wherein, Resource_1 will be overwritten because it is already available in the target domain, and Resource_2 and Resource_3 will be ADDED.
You have selected Project_1 from the source domain, and this project has Resource_1.	The target domain has Project_1, and this project has Resource_1, Resource_2, and Resource_3.	The entire Project_1 will be deployed to the target domain, wherein, Resource_1 will be overwritten because it is already available in the target domain, and Resource_2 and Resource_3 will be DELETED.	Only the resources of Project_1 will be deployed to the target domain, wherein, only Resource_1 will be overwritten because it is already available in the target domain. The other two resources already available in the target domain, that is, Resource_2 and Resource_3 will NOT be affected.

16.5 Provisioning Oracle Service Bus Resources from Oracle Software Library

This section describes how you can provision Oracle Service Bus resources from the Software Library.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

16.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

■ Ensure that you meet the prerequisites described in the following sections:

– Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

– Section 3.2, "Optional Infrastructure Requirements"

■ Export the resources of an Oracle Service Bus domain as a JAR file. Use Oracle Service Bus console for this.

■ Ensure that the JAR file is available as a generic component in Oracle Software Library. For instructions to create generic components, see Section 3.1.2, "Setting Up Oracle Software Library".

■ If you want to use a customization file to customize the environment variables in the changed (target) environment, then you must ensure that the customization file is available as a generic component in Oracle Software Library. For instructions to create generic components, see Section 3.1.2, "Setting Up Oracle Software Library".

Prerequisites for Operators

■ If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

16.5.2 Provisioning Procedure

To provision Oracle Service Bus resources from a source Oracle Service Bus domain, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Oracle Service Bus Resource Provisioning**. Then click **Schedule Deployment**.

4. On the Select Source page, in the Source section, select **Oracle Software Library**.

   a. For **Component**, click the torch icon and select the generic component that contains the resources to be deployed to a target Oracle Service Bus domain.

b. (Optional) For **Pass Phrase**, specify a pass-phrase if any of the resources in the JAR file contain sensitive data and are protected. The same pass-phrase is used while importing these resources to the target domain.

c. Click **Next**. Enterprise Manager Grid Control displays Select Target page.

5. On the Select Target page, do the following:

a. In the Target section, specify the following:

For **Domain**, click the torch icon and select the Oracle Service Bus domain where you want to deploy the selected resources.

For **BEA Home Directory**, specify the full path to the BEA home directory where all BEA product-related files are stored.

b. (Optional) In the Options section, select the settings you want to retain if you have done some customization to the resources selected for deployment, and if you want to preserve those changes in the target Oracle Service Bus domain.

Note that for Oracle Service Bus 2.6.x, Security and Policy Configuration, Credentials, and Access Control Policies cannot be preserved.

c. In the Customization section, provide details about the customization file that can be used to modify the environment settings in the target OSB domain.

If you do not want to use a customization file, select **None**.

If you are using a customization file and if it is available on the host where the target OSB domain is running, then select **Use the Customization file on the target host** and specify the full path to the location where the file is present.

If the customization file is stored as a generic component in Oracle Software Library, then select **Select the customization file from the Software Library** and specify the full path to the location in Oracle Software Library where the generic component is stored.

d. Click **Next**.

6. On the Set Credentials page, specify the following and click **Next**.

   a. Specify the login credentials of the source and target Oracle Service Bus (OSB) domains.

   b. Specify the credentials of the hosts where the Management Agents, which are monitoring the administration servers of the OSB domains, are running

7. On the Schedule page, schedule the deployment procedure to run immediately or later, and specify a unique name for the deployment procedure so that it can be tracked. Click **Next**. Enterprise Manager Grid Control displays the Review page.

8. On the Review page, review the details you have provided for provisioning Oracle Service Bus resources, and click **Submit**.

17

Provisioning Oracle BPEL Processes

Business Process Execution Language (BPEL) is an XML-based language for enabling task sharing across multiple enterprises using a combination of Web services. BPEL is based on the XML schema, simple object access protocol (SOAP), and Web services description language (WSDL). BPEL provides enterprises with an industry standard for business process orchestration and execution.

Oracle BPEL Process Manager (BPEL Process Manager) provides a framework for easily designing, deploying, monitoring, and administering processes based on BPEL standards. For information about discovering and monitoring these SOA applications, see the *Enterprise Manager Getting Started with Management Pack for SOA Guide* available at:

http://www.oracle.com/technology/documentation/oem.html

While Enterprise Manager Grid Control allows you to discover and monitor these BPEL Process Managers, it also provides Deployment Procedures that help you provision BPEL processes on BPEL Process Managers.

This chapter explains how you can provision BPEL processes on BPEL Process Managers. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedure
- Supported Releases
- Provisioning Oracle BPEL Processes

17.1 Getting Started

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

Table 17–1   Getting Started with Provisioning Oracle BPEL Processes

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle BPEL processes. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 17.2, "Deployment Procedure".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle BPEL Process Manager are supported by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 17.3, "Supported Releases".
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for provisioning Oracle BPEL processes, see Section 17.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle BPEL processes.	To provision Oracle BPEL processes, follow the steps explained in Section 17.4.2, "Provisioning Procedure".

17.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for provisioning BPEL processes on BPEL Process Managers:

- *BPEL Process Provisioning*

17.3 Supported Releases

Using this Deployment Procedure, you can provision BPEL processes for Oracle BPEL Process Manager 10.1.3.1, 10.1.3.3, and 10.1.3.4.

17.4 Provisioning Oracle BPEL Processes

This section describes how you can provision BPEL processes for Oracle BPEL Process Manager.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

17.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Ensure that BPEL Process Manager on which the process suitcase files have to be deployed is already discovered and monitored in Enterprise Manager Grid Control.

- Store the BPEL process suitcase files as generic components in the Software Library. For instructions to create generic components, see Section 3.1.2, "Setting Up Oracle Software Library".

> **Note:** While adding a generic component for BPEL processes, on the Create Component: Describe page of the Software Library Wizard, select **Generic Component** from the **Type** list, provide a name for the parent folder, and navigate to the Upload File page to upload the files. You DO NOT have to provide details for Customize and Set Directives page.

- If you want to use a deployment plan that can be associated with a BPEL process suitcase file (JAR file), then store this deployment file as a generic component in the Software Library. For instructions to create generic components, see Section 3.1.2, "Setting Up Oracle Software Library".

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

17.4.2 Provisioning Procedure

To provision BPEL processes to a target BPEL Process Manager, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **BPEL Process Provisioning**. Then click **Schedule Deployment**.

4. On the Source Selection page, do the following:

    a. In the Source section, click **Add** and select the BPEL Process suitcase files that you want to deploy to a target BPEL Process Manager. The table is populated based on the selection made.

If you have selected multiple suitcase files, then from the table, select the BPEL process suitcase file and use **Move Up** and **Move Down** to order the components the way they should be deployed by the Deployment Procedure.

    b. In the Select Deployment Plan, select a deployment plan that can be associated with a BPEL process suitcase file (JAR file).

The deployment plan helps you modify the configuration details and partner link binding properties, which have been set for a particular environment, at run time. You can also use it to search and replace strings and URLs that have been set for a particular environment. This way, you can deploy the same BPEL processes on BPEL Process Managers that are in development, test, and production environments; without having to reconfigure your settings across these environments.

    c. Click **Next**.

5. On the Target Selection page, in the Target section, click **Add** and select the BPEL Process Managers on which you want to deploy the BPEL processes. If there are multiple domains available for a BPEL target, then you can select an appropriate domain from the **BPEL Domain** list, on which the suitcase files can be deployed. Click **Next**.

> **Note:** When you click **Next**, Enterprise Manager Grid Control
> internally checks to see if the Context Provider URL is captured. If this
> URL is not captured, then you may see some errors. To resolve this
> issue, set the URL in the Monitoring Configuration page of the BPEL
> target.

6. On the Credentials page, specify the following:

   a. Credentials of application server instances on which the selected BPEL Process
      Managers are running.

   Application Server Credentials
   Provide credentials for the application server instances on which the selected Oracle BPEL Process Manag

Application Server	Username	Password
soa_demo.stapp04.us.oracle.com		

   b. BPEL administrator credentials (and RMI credentials for 10.1.2 BPEL targets)
      for the selected Oracle BPEL Process Managers.

   BPEL Process Manager Credentials
   Provide BPEL administrator credentials for the selected Oracle BPEL Process Managers

   BPEL 10.1.3 Targets

BPEL Process Manager	Username	Password
soa_demo.stapp04.us.oracle.com_bpel		

   The credentials required for BPEL Process Managers vary according to the
   supported BPEL Process Manager version. For BPEL 10.1.3 targets, you need
   to provide only one set of credentials that will be used for accessing the BPEL
   Process Manager. However, for BPEL 10.1.2 targets, you need to provide the
   BPEL administrator password and another set of OC4J RMI Access credentials
   for remote access.

   If the preferred credentials are already set and stored in the Management
   Repository, then by default, they are prefilled on this page. You can choose to
   either use these prefilled preferred credentials or edit them to use the changed
   credentials. If the preferred credentials are not stored, then the fields are
   blank. In this case, you have to specify the credentials. The credentials
   specified here apply only to the current deployment procedure session and do
   not get stored in the Management Repository for future use.

   Note that if you change the credentials, the change applies only to the current
   deployment procedure session and does not override the preferred credentials
   stored in the Management Repository.

   c. Click **Next**.

7. On the Schedule page, schedule the deployment procedure to run immediately or
   later, and specify a unique name for the deployment procedure so that it can be
   tracked. Click **Next**.

8. On the Review page, review the details you have provided for provisioning BPEL
   processes, and click **Submit**.

18

Provisioning Oracle Application Server

This chapter explains how you can provision Oracle Application Servers. In particular, this chapter covers the following:

- Deployment Procedures, Supported Releases, and Core Components Deployed

- Provisioning Oracle Application Server 10g Release 1 (10.1.3)

- Provisioning Oracle Application Server 10g Release 1 (10.1.2.0.2)

- Provisioning Oracle SOA Suite 10g (10.1.3.x)

18.1 Getting Started

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

Table 18–1   Getting Started with Provisioning Oracle Application Server

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for provisioning Oracle Application Server. Know how the Deployment Procedure functions, what use cases it covers, what releases it supports, and what core components it provisions.	To learn about the Deployment Procedure, see Section 18.2, "Deployment Procedures, Supported Releases, and Core Components Deployed".

Table 18–1 (Cont.) Getting Started with Provisioning Oracle Application Server

Step	Description	Reference Links
Step 2	**Selecting the Release to Provision** This chapter covers use cases for different releases of Oracle Application Server. Identify the release you want to provision and understand the use cases that are covered for each release.	■ To learn about provisioning Oracle Application Server 10g Release 1 (10.1.3), see Section 18.3, "Provisioning Oracle Application Server 10g Release 1 (10.1.3)". ■ To learn about provisioning Oracle Application Server 10g Release 1 (10.1.2.0.2), see Section 18.4, "Provisioning Oracle Application Server 10g Release 1 (10.1.2.0.2)". ■ To learn about provisioning Oracle SOA Suite 10g (10.1.3.x), see Section 18.5, "Provisioning Oracle SOA Suite 10g (10.1.3.x)".
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for provisioning Oracle Application Server, access the reference links provided for Step (2) and navigate to the *Prerequisites* subsection.
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully provision Oracle Application Server.	To provision Oracle Application Server, access the reference links provided for Step (2) and navigate to the *Provisioning Procedure* subsection.

18.2 Deployment Procedures, Supported Releases, and Core Components Deployed

Enterprise Manager Grid Control offers the following Deployment Procedures for provisioning Oracle Application Server:

Deployment Procedure	Supported Releases	Core Components Deployed
Application Server Deployment (myJ2EE) 10.1.2.0.2	Oracle Application Server 10g Release 2 (10.1.2.0.2)	■ Application tier ■ Web tier
Application Server Deployment 10.1.3	Oracle Application Server 10g Release 3 (10.1.3.x)	■ Application tier ■ Web tier
Application Server Deployment 10.1.3.xSOA	Oracle SOA Suite 10g (10.1.3.x)	Deploys Oracle SOA Suite 10g (10.1.3.x) with the application tier and Web tier of Oracle Application Server 10g Release 3 (10.1.3.x)

> **Note:** These Deployment Procedures do not install the database tier. However, they facilitate the configuration of Java Authentication and Authorization Service (JAAS) provider (also called as JAZN) with an existing data tier.

Oracle Application Server you want to provision may be available in different formats such as an installation kit in the installation medium, a running instance on a host monitored by Enterprise Manager Grid Control, or even a gold image in Oracle Software Library (Software Library).

Using the *Application Server Deployment 10.1.3* Deployment Procedure, you can provision any of these formats. However, using the *Application Server Deployment (myJ2EE) 10.1.2.0.2* Deployment Procedure, you can provision only a new installation using the software binaries stored in the Software Library.

For example, if you want to have a new installation of Oracle Application Server 10g Release 3 (10.1.3.x), the recommended option is to store the software binaries of the application server from an installation medium onto the Software Library, and then provision it from the Software Library. However, if you want to have a copy of a running instance that is stable and has all the latest patches applied, the recommended option is to clone that existing instance so that you retain the same configuration settings. Similarly, if you have created a gold image of this stable, well-patched application server instance in the Software Library, then you can use it to deploy a similar instance in your enterprise configuration.

> **Note:** Provisioning of a gold image from the Software Library is not supported for Microsoft Windows Vista.

18.3 Provisioning Oracle Application Server 10g Release 1 (10.1.3)

This section describes how you can provision Oracle Application Server 10g Release 1 (10.1.3). In particular, this section covers the following:

- Cloning a Running Oracle Application Server Instance
- Provisioning a New Installation or Gold Image of Oracle Application Server

18.3.1 Cloning a Running Oracle Application Server Instance

This section describes how you can clone an existing Oracle Application Server instance that is running on a host monitored by Enterprise Manager Grid Control.

This option is best suited when you have a running instance of Oracle Application Server that is stable and has all the latest patches applied, and you want to make identical copies of it on multiple hosts.

However, the risk involved in using a existing instance is that the instance may be deleted or deinstalled anytime without prior notice, and as a result, the Deployment Procedure may fail. Therefore, use this option when you know that the running instance is available for cloning.

This section covers the following:

- Cloning from an Existing Cluster, Scaling Up the Existing Cluster, and Using the Same Internet Directory
- Provisioning from an Existing Cluster and Creating a New Cluster Without Internet Directory
- Provisioning from an Existing Cluster and Creating a New Cluster With Internet Directory

18.3.1.1 Cloning from an Existing Cluster, Scaling Up the Existing Cluster, and Using the Same Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3) from an existing cluster, extend the cluster to include the new node, and by using the same Internet Directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.1.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

  ```
  Directives/Oracle Directives/Loadbalancer
  ```

- Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.1.1.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**. Enterprise Manager Grid Control displays the Source Selection page of the Deployment Procedure.

4. On the Source Selection page, in the Source Selection section, in the Select the source Oracle Homes from the Installed Cluster Environment subsection, do the following:

   a. Retain the default selection, that is, **Select from Existing Cluster Installation**.

   b. From the **Select Cluster** list, select an existing cluster from where the Web tier and application tier can be deployed. On selection of an existing cluster, the table gets populated with information about the Web tier and application tier that are part of it.

   c. In the table, for **Files to Exclude**, specify the files in the Oracle home that you do not want to include while running this Deployment Procedure. You can specify either the file names or the file extensions, and separate multiple entries with a comma.

   And for **Working Directory**, retain the default value or specify another location that already exists on the source host and can be used to stage the file for cloning.

   d. If you want to deploy only one of the two products (application server or web server) from the selected source cluster, then click **Reset** to delete all values in the table. Once the values are reset, click the torch icon corresponding to the product that you want to deploy, and select the required target.

   e. If you want to save the Oracle homes as an image in the Software Library, select **Save to Software Library**.

   f. Click **Next**.

   Note: When you click **Next**, if there are no prerequisites in the Software Library for that particular platform, then you will see **Next, Ignore Warnings** option. In this case, you can do one of the following:

   - Upload the components to the Software Library under *Components/Oracle Components/Prerequisite-fixup components/<Platform name>/Prerequisite-Fixup component* and click **Next**.

   - Click **Next, Ignore Warnings** to skip the prerequisite check while running the Deployment Procedure.

   ---

5. On the Target List page, do the following:

   a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

   b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

   c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

   a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

   To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

   b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

   Note: You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

   - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

   - The Deployment Procedures has been customized to support locked environments.

   For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

   From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

   c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

   To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make,

specify the credentials. Ensure that the users belong to the same operating system group.

d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

e. Click **Next**.

7. On the Application and Web Tier page, do the following:

a. In the Cluster Details section, retain the default selection, that is, **Extend existing cluster**.

b. In the Instance Details section, retain the default instance names, and specify the OC4J administrator password for the source host and the target host.

You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

IMPORTANT:

- Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

- If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).

- The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

---

c. In the Port Details section, retain the default port number (7777) that is displayed if Oracle Web Cache is not configured.

d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

NOTE: Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

---

e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

IMPORTANT: If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.

This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.

However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

---

    f. In the Identify Management Configuration section, retain the default selection.

    g. Click **Next**.

8. On the Configure Oracle Home page, do the following:

    a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

    b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

Note: You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

---

    c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded,

see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

d. Click **Next**.

9. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.1.2 Provisioning from an Existing Cluster and Creating a New Cluster Without Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3) from an existing cluster, and create a new cluster without Internet directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.1.2.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  – Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  – Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

  ```
  Directives/Oracle Directives/Loadbalancer
  ```

- Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.1.2.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the source Oracle Homes from the Installed Cluster Environment subsection, do the following:

   a. Retain the default selection, that is, **Select from Existing Cluster Installation**.

   b. From the **Select Cluster** list, select an existing cluster from where the Web tier and application tier can be deployed. On selection of an existing cluster, the table gets populated with information about the Web tier and application tier that are part of it

   c. In the table, for **Files to Exclude**, specify the files in the Oracle home that you do not want to include while running this Deployment Procedure. You can specify either the file names or the file extensions, and separate multiple entries with a comma.

   And for **Working Directory**, retain the default value or specify another location that already exists on the source host and can be used to stage the file for cloning.

   d. If you want to deploy only one of the two products (application server or web server) from the selected source cluster, then click **Reset** to delete all values in the table. Once the values are reset, click the torch icon corresponding to the product that you want to deploy, and select the required target.

   e. If you want to save the Oracle homes as an image in the Software Library, select **Save to Software Library**.

   f. Click **Next**.

5. On the Target List page, do the following:

   a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

Note: You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

■ The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

■ The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

    e. Click **Next**.

7. On the Application and Web Tier page, do the following:

    a. In the Cluster Details section, select **Create new Cluster**. By default, Grid Control prefills the details based on an existing cluster. You can either use the default values or specify a new cluster name, installation directory, and multicast address and port. Also ensure that the multicast address and port are different from the ones configured for the source cluster.

    For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively.

    For **Multicast Address**, ensure that the address is within the range of 224.0.0.0 and 239.255.255.255. This is a single IP address for a set of nodes that are joined in a multicasting group.

    b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

    You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

IMPORTANT:

- Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

- If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).

- The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

---

    c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

    d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

> **NOTE:** Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

    e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

    f. In the Identify Management Configuration section, select **None**.

    g. Click **Next**.

8. On the Configure Oracle Home page, do the following:

    a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

    An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

    If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

    b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

> **Note:** You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

    c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

    In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

    d. Click **Next**.

9. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.1.3 Provisioning from an Existing Cluster and Creating a New Cluster With Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3) from an existing cluster, and create a new cluster with Internet directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.1.3.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For

other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

```
Directives/Oracle Directives/Loadbalancer
```

- Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.1.3.2   Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the source Oracle Homes from the Installed Cluster Environment subsection, do the following:

   a. Retain the default selection, that is, **Select from Existing Cluster Installation**.

   b. From the **Select Cluster** list, select an existing cluster from where the Web tier and application tier can be deployed. On selection of an existing cluster, the table gets populated with information about the Web tier and application tier that are part of it

   c. In the table, for **Files to Exclude**, specify the files in the Oracle home that you do not want to include while running this Deployment Procedure. You can specify either the file names or the file extensions, and separate multiple entries with a comma.

      And for **Working Directory**, retain the default value or specify another location that already exists on the source host and can be used to stage the file for cloning.

      d. If you want to deploy only one of the two products (application server or web server) from the selected source cluster, then click **Reset** to delete all values in the table. Once the values are reset, click the torch icon corresponding to the product that you want to deploy, and select the required target.

      e. If you want to save the Oracle homes as an image in the Software Library, select **Save to Software Library**.

      f. Click **Next**. Enterprise Manager Grid Control displays the Target List page.

5. On the Target List page, do the following:

      a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

      b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

      c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

      a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

> **Note:** You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

      From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

      b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

      To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

    c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

    To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

    d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

7. On the Application and Web Tier page, do the following:

    a. In the Cluster Details section, select **Create new Cluster**. By default, Grid Control prefills the details based on an existing cluster. You can either use the default values or specify a new cluster name, installation directory, and multicast address and port. Also ensure that the multicast address and port are different from the ones configured for the source cluster.

    For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively.

    For **Multicast Address**, ensure that the address is within the range of 224.0.0.0 and 239.255.255.255. This is a single IP address for a set of nodes that are joined in a multicasting group.

    b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

    You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

IMPORTANT:

- Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

- If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).

- The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

---

c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

> **NOTE:** Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

f. In the Identify Management Configuration section, select **Configure Java Authentication and Authorization Service (JAZN) with a LDAP-based provider**.

Using this option, you can set up JAZN LDAP-based provider for authentication and authorization for OC4J application.

Java Authentication and Authorization Service (JAAS) is a Java package that enables services and applications to authenticate and enforce access controls upon users. Oracle Application Server 10g Containers for J2EE (OC4J) supports JAAS by implementing a JAAS provider (also called as JAZN).

The JAAS provider provides application developers with user authentication, authorization, and delegation services to integrate into their application environments. It also supports JAAS policies. Policies contain the rules (permissions) that authorize a user to use resources, such as reading a file and so on.

g. Click **Next**. Enterprise Manager Grid Control displays the Identity Management page.

8. On the Identity Management page, do the following:

a. In the Identity Management Host Details section, specify the connection information for the Internet Directory to be used for Identity Management of the Oracle Application Server users and groups. If you do not have an Internet Directory installed, install it using the OracleAS Infrastructure component.

b. In the Internet Directory Login Details section, specify credentials of the user who belongs to the iASAdmin group in the Internet Directory.

c. Click **Next**. Enterprise Manager Grid Control displays the Configure Oracle Home page.

9. On the Configure Oracle Home page, do the following:

a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

> **Note:** You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

10. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.2 Provisioning a New Installation or Gold Image of Oracle Application Server

This section describes how you can provision a new installation or a gold image of Oracle Application Server. In particular, this section covers the following:

- Provisioning and Creating a New Cluster Without Internet Directory
- Provisioning and Creating a New Cluster With Internet Directory
- Provisioning and Treating Oracle Application Server as a Standalone Instance Without Internet Directory
- Provisioning and Treating Oracle Application Server as a Standalone Instance With Internet Directory

18.3.2.1 Provisioning and Creating a New Cluster Without Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical either to the ones available on the installation medium or a gold image available in the Software Library, and create a new cluster without Internet directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.2.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

  ```
  Directives/Oracle Directives/Loadbalancer
  ```

- Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.2.1.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the Gold Image of an Oracle Home or a shiphome from the Software Library subsection, do the following:

   a. Select **Select from Software Library**.

   b. In the **Source for Web Tier** section, click the torch icon and select the generic component that contains the Web tier.

   c. In the **Source for App Tier** section, click the torch icon and select the generic component that contains the application tier.

   Note: When you click the torch icon to search for components, the Select Component page displays the components available in the Software Library.

   - If you want to deploy a fresh installation, then select the shiphome component that is in "Ready" status. By default, the Select Component page does not display components with "Active" status.

   - If you want to deploy a gold image, then select the gold image component that is either in "Ready" or "Active" status.

   ---

   d. Click **Next**.

5. On the Target List page, do the following:

      a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

      b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

      c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

      a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

      To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

      b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

> **Note:** You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

      From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

      c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

      To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make,

specify the credentials. Ensure that the users belong to the same operating system group.

   d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

   e. Click **Next**.

7. On the Application and Web Tier page, do the following:

   a. In the Cluster Details section, select **Create new Cluster**. By default, Grid Control prefills the details based on an existing cluster. You can either use the default values or specify a new cluster name, installation directory, and multicast address and port. Also ensure that the multicast address and port are different from the ones configured for the source cluster.

     For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively.

     For **Multicast Address**, ensure that the address is within the range of 224.0.0.0 and 239.255.255.255. This is a single IP address for a set of nodes that are joined in a multicasting group.

   b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

     You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

IMPORTANT:

- Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

- If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).

- The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

---

   c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

   d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

> **NOTE:** Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

    e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

    f. In the Identify Management Configuration section, select **None**.

    g. Click **Next**.

8. On the Configure Oracle Home page, do the following:

    a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

    An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

    If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

    b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

> **Note:** You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

    c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

    In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

    d. Click **Next**.

9. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.2.2 Provisioning and Creating a New Cluster With Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical either to the ones available on the installation medium or a gold image available in the Software Library, and create a new cluster with Internet directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.2.2.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For

other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

```
Directives/Oracle Directives/Loadbalancer
```

■ Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

■ If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.2.2.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the Gold Image of an Oracle Home or a shiphome from the Software Library subsection, do the following:

   a. Select **Select from Software Library**.

   b. In the **Source for Web Tier** section, click the torch icon and select the generic component that contains the Web tier.

   c. In the **Source for App Tier** section, click the torch icon and select the generic component that contains the application tier.

> **Note:** When you click the torch icon to search for components, the Select Component page displays the components available in the Software Library.
>
> - If you want to deploy a fresh installation, then select the shiphome component that is in "Ready" status. By default, the Select Component page does not display components with "Active" status.
>
> - If you want to deploy a gold image, then select the gold image component that is either in "Ready" or "Active" status.

    d. Click **Next**.

5. On the Target List page, do the following:

    a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

    b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

    c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

    a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

> **Note:** You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

    From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

    b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

    To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same**

    for all Oracle Homes if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

  c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

    To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

  d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

  e. Click **Next**.

7. On the Application and Web Tier page, do the following:

  a. In the Cluster Details section, select **Create new Cluster**. By default, Grid Control prefills the details based on an existing cluster. You can either use the default values or specify a new cluster name, installation directory, and multicast address and port. Also ensure that the multicast address and port are different from the ones configured for the source cluster.

    For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively.

    For **Multicast Address**, ensure that the address is within the range of 224.0.0.0 and 239.255.255.255. This is a single IP address for a set of nodes that are joined in a multicasting group.

  b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

    You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

> **IMPORTANT:**
>
> ■ Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.
>
> ■ If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).
>
> ■ The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

> **NOTE:** Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

f. In the Identify Management Configuration section, select **Configure Java Authentication and Authorization Service (JAZN) with a LDAP-based provider**.

Using this option, you can set up JAZN LDAP-based provider for authentication and authorization for OC4J application.

Java Authentication and Authorization Service (JAAS) is a Java package that enables services and applications to authenticate and enforce access controls upon users. Oracle Application Server 10g Containers for J2EE (OC4J) supports JAAS by implementing a JAAS provider (also called as JAZN).

The JAAS provider provides application developers with user authentication, authorization, and delegation services to integrate into their application environments. It also supports JAAS policies. Policies contain the rules (permissions) that authorize a user to use resources, such as reading a file and so on.

g. Click **Next**.

8. On the Identity Management page, do the following:

a. In the Identity Management Host Details section, specify the connection information for the Internet Directory to be used for Identity Management of the Oracle Application Server users and groups. If you do not have an Internet Directory installed, install it using the OracleAS Infrastructure component.

b. In the Internet Directory Login Details section, specify credentials of the user who belongs to the iASAdmin group in the Internet Directory.

c. Click **Next**.

9. On the Configure Oracle Home page, do the following:

a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

Note: You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

---

c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

d. Click **Next**.

10. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.2.3 Provisioning and Treating Oracle Application Server as a Standalone Instance Without Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical either to the ones available on the installation medium or a gold image available in the Software Library, and treat that as a standalone instance without Internet directory.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.3.2.3.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

  To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

  ```
  Directives/Oracle Directives/Loadbalancer
  ```

■ Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

■ If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.2.3.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the Gold Image of an Oracle Home or a shiphome from the Software Library subsection, do the following:

   a. Select **Select from Software Library**.

   b. In the **Source for Web Tier** section, click the torch icon and select the generic component that contains the Web tier.

   c. In the **Source for App Tier** section, click the torch icon and select the generic component that contains the application tier.

   Note: When you click the torch icon to search for components, the Select Component page displays the components available in the Software Library.

   ■ If you want to deploy a fresh installation, then select the shiphome component that is in "Ready" status. By default, the Select Component page does not display components with "Active" status.

   ■ If you want to deploy a gold image, then select the gold image component that is either in "Ready" or "Active" status.

   ---

     d. Click **Next**.

5. On the Target List page, do the following:

     a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

     b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

     c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

     a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

> **Note:** You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> ■ The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> ■ The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

       From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

     b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

       To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

     c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

       To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials

across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

   d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

   e. Click **Next**.

7. On the Application and Web Tier page, do the following:

   a. In the Cluster Details section, select **Standalone AS**.

   b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

   You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

   IMPORTANT:

   - Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

   - If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).

   - The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

   ---

   c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

   d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

   NOTE: Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

   ---

   e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment

Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

   f. In the Identify Management Configuration section, select **None**.

   g. Click **Next**.

8. On the Configure Oracle Home page, do the following:

   a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

   An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

   If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

   b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

   > **Note:** You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

   c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

d. Click **Next**.

9. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.3.2.4 Provisioning and Treating Oracle Application Server as a Standalone Instance With Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical either to the ones available on the installation medium or a gold image available in the Software Library, and treat that as a standalone instance with Internet directory.

In particular, this section covers the following:

■ Prerequisites

■ Provisioning Procedure

18.3.2.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

■ Ensure that you meet the prerequisites described in the following sections:

– Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

– Section 3.2, "Optional Infrastructure Requirements"

■ Compare the configuration of the source and target hosts and ensure that they have the same configuration. If the configurations are different, then contact your system administrator and fix the inconsistencies before running the Deployment Procedure.

To compare the configuration of the hosts, in Grid Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then click the Configuration tab. On the Configuration page, click **Compare Configuration** and select the target host.

■ If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

```
Directives/Oracle Directives/Loadbalancer
```

■ Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.3.2.4.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Source Selection page, in the Source Selection section, in the Select the Gold Image of an Oracle Home or a shiphome from the Software Library subsection, do the following:

   a. Select **Select from Software Library**.

   b. In the **Source for Web Tier** section, click the torch icon and select the generic component that contains the Web tier.

   c. In the **Source for App Tier** section, click the torch icon and select the generic component that contains the application tier.

   Note: When you click the torch icon to search for components, the Select Component page displays the components available in the Software Library.

   - If you want to deploy a fresh installation, then select the shiphome component that is in "Ready" status. By default, the Select Component page does not display components with "Active" status.

   - If you want to deploy a gold image, then select the gold image component that is either in "Ready" or "Active" status.
   ---

   d. Click **Next**.

5. On the Target List page, do the following:

    a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

    b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

    c. Click **Next**.

6. On the Credentials/Schedule page, do the following:

    a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

    To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

    b. In the Agent Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

Note: You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

    From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

    c. In the Source Oracle Home Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

    To override the preferred credentials with another set of credentials, select **Override Preferred Credentials**. From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make,

specify the credentials. Ensure that the users belong to the same operating system group.

   d. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

   e. Click **Next**.

7. On the Application and Web Tier page, do the following:

   a. In the Cluster Details section, select **Standalone As**.

   b. In the Instance Details section, specify unique instance names for the Web tier instance and application tier instance, and specify the OC4J administrator password for the source host and the target host.

You can always change the default instance names with any other custom names. However, ensure that the names you specify are unique because, by default, the host name and domain name of that host are appended to the instance name you specify here.

> **IMPORTANT:**
>
> - Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.
>
> - If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in the Additional Parameters section. For more information, see Step 7 (e).
>
> - The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

   c. In the Port Details section, specify the HTTP load balancer host and listener ports to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777. Select **Enable SSL** if you want to secure the communications.

   d. In the Load Balancer Details section, select **Configure Load Balancer** and provide the required information if you have an F5 BIG-IP Local Traffic Manager load balancer configured.

> **NOTE:** Enterprise Manager Grid Control supports only F5 BIG-IP Local Traffic Manager. If you have any other load balancer, then deselect **Configure Load Balancer**, and manually configure that load balancer and the HTTP servers associated with it. Do not use this section in this case.

   e. In the Additional Parameters section, specify any additional Web tier-specific or application tier-specific parameters you want to pass to the Deployment Procedure. For example, -debug. You can specify any other Oracle Universal Installer (OUI) parameter that can be in this provisioning operation.

> **IMPORTANT:** If you are cloning an application server with multiple OC4J instances that have different passwords, then ensure that you use the `-force` command as an additional parameter in this section.
>
> This is to ensure that the password you specify in the Instance Details section is uniformly propagated to all OC4J instances that are being provisioned.
>
> However, if you still want to maintain different passwords for the provisioned OC4J instances, then after the deployment procedure ends successfully, access the application server console and change the passwords for the individual OC4J instances.

   f. In the Identify Management Configuration section, select **Configure Java Authentication and Authorization Service (JAZN) with a LDAP-based provider**.

   Using this option, you can set up JAZN LDAP-based provider for authentication and authorization for OC4J application.

   Java Authentication and Authorization Service (JAAS) is a Java package that enables services and applications to authenticate and enforce access controls upon users. Oracle Application Server 10g Containers for J2EE (OC4J) supports JAAS by implementing a JAAS provider (also called as JAZN).

   The JAAS provider provides application developers with user authentication, authorization, and delegation services to integrate into their application environments. It also supports JAAS policies. Policies contain the rules (permissions) that authorize a user to use resources, such as reading a file and so on.

   g. Click **Next**.

8. On the Identity Management page, do the following:

   a. In the Identity Management Host Details section, specify the connection information for the Internet Directory to be used for Identity Management of the Oracle Application Server users and groups. If you do not have an Internet Directory installed, install it using the OracleAS Infrastructure component.

   b. In the Internet Directory Login Details section, specify credentials of the user who belongs to the iASAdmin group in the Internet Directory.

   c. Click **Next**. Enterprise Manager Grid Control displays the Configure Oracle Home page.

9. On the Configure Oracle Home page, do the following:

   a. If the hosts where the Web tier and application tier are being provisioned have a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

   An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

   If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle

Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

b. If the hosts where the Web tier and application tier are being provisioned have an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

Note: You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned application server.

---

c. If the hosts where the Web tier and application tier are being provisioned do not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

d. Click **Next**.

10. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Submit**.

18.4 Provisioning Oracle Application Server 10g Release 1 (10.1.2.0.2)

This section describes how you can provision Oracle Application Server 10g Release 1 (10.1.2.0.2). In particular, this section covers the following:

- Provisioning Oracle Application Server and Creating a New Cluster Without Internet Directory

- Provisioning Oracle Application Server and Creating a New Cluster With Internet Directory

---

Note: You can provision Oracle Application Server 10g Release 1 (10.1.2.0.2) only using the software binaries archived from the installation medium and stored in the Software Library.

18.4.1 Provisioning Oracle Application Server and Creating a New Cluster Without Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical to the ones available on the installation medium, and create a new cluster without Internet directory.

This option is best suited when you want a completely new installation to be provisioned across multiple hosts. Of course, understandably, this is a fresh installation and you will have to update it with all the latest patches that have been released so far.

In particular, this section covers the following:

- Prerequisites
- Provisioning Procedure

18.4.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

  The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

  ```
  Directives/Oracle Directives/Loadbalancer
  ```

- Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

- If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.4.1.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Target List page, do the following:

   a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

   b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

   c. Click **Next**.

5. On the Upload page, from the Shiphomes from Software Library section, click **Upload** to select a component from the software library that contains the gold images of the application tier and Web tier that you want to deploy.

6. On the Application and Web Tier page, do the following:

   a. In the Cluster Details section, for **Cluster Name**, specify a name for the cluster. Alternatively, you can retain the default name provided.

   For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively. For example, if you specify `/scratch/appserver/` as the Web tier install base directory, then the Web tier is installed in `/scratch/appserver/10gWebTier`.

   b. In the Instance Details section, for **Instance Name**, specify a unique name and the ias_admin password. Alternatively, you can retain the default name provided.

> **IMPORTANT:**
>
> - Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.
>
> - The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

   c. In the Port Details section, specify the HTTP load balancer host, the listener port, and the Application Server Control port to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777 and 1810, for Oracle HTTP Server Port and Application Server Control port, respectively. Select **Enable SSL** if you want to secure the communications.

   d. In the Identify Management Configuration section, select **None**.

   e. Click **Next**.

7. On the Library Step Properties page, specify property values for the custom library steps, and click **Next**.

8. On the Credentials/Schedule page, do the following:

   a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

> **Note:** You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:
>
> - The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.
>
> - The Deployment Procedures has been customized to support locked environments.
>
> For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

   From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

   b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

   c. Click **Next**.

9. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Finish**.

18.4.2 Provisioning Oracle Application Server and Creating a New Cluster With Internet Directory

This section describes how you can provision the application tier and Web tier of Oracle Application Server 10g Release 1 (10.1.3), which are identical to the ones available on the installation medium, and create a new cluster with Internet directory.

This option is best suited when you want a completely new installation to be provisioned across multiple hosts. Of course, understandably, this is a fresh installation and you will have to update it with all the latest patches that have been released so far.

In particular, this section covers the following:

■ Prerequisites

■ Provisioning Procedure

18.4.2.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

■ Ensure that you meet the prerequisites described in the following sections:

– Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

– Section 3.2, "Optional Infrastructure Requirements"

■ If you are deploying multiple Web tiers, then ensure that a Software Load Balancer is set up and is accessible from the Web tier's hosts.

The Software Library provides a script to configure the F5 Big IP Application Switch (Software Version 4.5 PTF.5) load balancer on the selected Web tiers. For other systems such as Cisco CSM 3.1, overwrite the script with your specific configuration. The script is available in the following location of the Software Library:

```
Directives/Oracle Directives/Loadbalancer
```

■ Ensure that the installation base directory, where the Web tier and application tier will be installed, is accessible (but not shared) on all target hosts.

Prerequisites for Operators

■ If you have PAM/LDAP enabled in your environment, then ensure that the target agents are configured with PAM/LDAP. For more information, see My Oracle Support note 422073.1.

■ Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

18.4.2.2 Provisioning Procedure

To provision a Web tier and an application tier, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Application Server Provisioning Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Application Server Deployment 10.1.3**. Then click **Schedule Deployment**.

4. On the Target List page, do the following:

   a. In the Web Tier Hosts section, click **Add** to add hosts on which you want to deploy the Web tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the application tier.

   b. In the Application Tier Hosts section, click **Add** to add hosts on which you want to deploy the application tier. Ensure that the platform of this host is the same as the platform of the host on which you want to deploy the Web tier.

   c. Click **Next**.

5. On the Upload page, from the Shiphomes from Software Library section, click **Upload** to select a component from the software library that contains the gold images of the application tier and Web tier that you want to deploy.

6. On the Application and Web Tier page, do the following:

   a. In the Cluster Details section, for **Cluster Name**, specify a name for the cluster. Alternatively, you can retain the default name provided.

   For **Web Tier Install Base Directory** and **Application Tier Install Base Directory**, ensure that you specify the absolute path to the directory where you want to deploy the Web tier and application tier, respectively. For example, if you specify `/scratch/appserver/` as the Web tier install base directory, then the Web tier is installed in `/scratch/appserver/10gWebTier`.

   b. In the Instance Details section, for **Instance Name**, specify a unique name and the ias_admin password. Alternatively, you can retain the default name provided.

IMPORTANT:

- Each Oracle Application Server 10g instance has its own password, regardless of which user performed the installation. Passwords are not shared across instances, even if the instances were installed by the same user.

- The passwords you specify here must contain a minimum of 5 and a maximum of 30 alphanumeric characters. It can include underscore (_), dollar ($), or pound (#) characters. It must start with an alphabet and must contain at least one numeric value.

---

   c. In the Port Details section, specify the HTTP load balancer host, the listener port, and the Application Server Control port to manage HTTP connections made by client applications. Alternatively, you can retain the default value, that is, 7777 and 1810, for Oracle HTTP Server Port and Application Server Control port, respectively. Select **Enable SSL** if you want to secure the communications.

   d. In the Identify Management Configuration section, select **Configure Java Authentication and Authorization Service (JAZN) with a LDAP-based provider**.

   e. Click **Next**.

7. On the Identity Management page, do the following:

   a. In the Identity Management Host Details section, specify the connection information for the Internet Directory to be used for Identity Management of the Oracle Application Server users and groups. If you do not have an Internet Directory installed, install it using the OracleAS Infrastructure component.

   b. In the Internet Directory Login Details section, specify credentials of the user who belongs to the iASAdmin group in the Internet Directory.

   c. Click **Next**.

8. On the Library Step Properties page, specify property values for the custom library steps, and click **Next**.

9. On the Credentials/Schedule page, do the following:

   a. In the Target Host Credentials section, retain the default section, that is, **Preferred Credentials** so that the preferred credentials stored in the Management Repository can be used.

Note: You can optionally override these preferred credentials. The credentials you specify here are used by the Deployment Procedure to run the provisioning operation. If this environment is secure and has locked accounts, then make sure that:

- The credentials you specify here have the necessary privileges to switch to the locked account for performing the provisioning operation.

- The Deployment Procedures has been customized to support locked environments.

For more information, see Chapter 31.7, "Customizing Steps to Run as Another User".

From the **Host Credentials** list, select **Same for all Oracle Homes** if you want to use the same operating system credentials across hosts, or select **Different for each Oracle Home** if you want to use different credentials for each host. According to the selection you make, specify the credentials. Ensure that the users belong to the same operating system group.

b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

c. Click **Next**.

10. On the Review page, review the details you have provided for provisioning a Web tier and application tier, and click **Finish**.

18.5 Provisioning Oracle SOA Suite 10g (10.1.3.x)

The use cases and the prerequisites and deployment instructions for each of the use cases for provisioning Oracle SOA Suite 10g (10.1.3.x) are almost the same as the ones described in Section 18.3, "Provisioning Oracle Application Server 10g Release 1 (10.1.3)".

The only differences are:

- You cannot provision Oracle SOA Suite using the software binaries from an installation medium.

- In Step (3), on the Deployment Procedure Manager page, in the Procedures tab, from the table, select **Application Server Deployment 10.1.3.xSOA**. Then click **Schedule Deployment**.

Part VI

Patching

This part describes all patching-related Deployment Procedures. In particular, this part describes the Deployment Procedures for:

19

Getting Started with Patching

This chapter describes the best practices for running patching Deployment Procedures and helps you get started with the patching operation. In particular, this chapter covers the following:

- Getting Started with Patching

- Best Practices

- Recommendations for Testing and Implementation

19.1 Getting Started with Patching

Table 19 shows the sequence of steps you must follow to successfully patch a target.

Table 19–1 Getting Started with Patching

Sequence	Action	More Information
Step 1	**Meeting Prerequisites** Ensure that you meet all the prerequisites, mainly creating and configuring Oracle Software Library (Software Library), applying mandatory patches, uploading OPatches to the Software Library, and meeting all other prerequisites for online and offline patching.	For information about the prerequisites to be met, see the prerequisites section of the chapters in Part VI, "Patching".
Step 2	**Identifying Patchability of an Environment by Running EM Target Patchability Report** Run the Patchability Report to analyze the database environment and identify whether the database targets you want to patch are suitable for a patching operation . If they are not suitable, then you can diagnose and find out whether they are not suitable due to missing properties or unsupported configuration, and also resolve them based on the recommendations provided for those issues.	For information about the patchability report, see Appendix H, "About Patching Automation Reports".

Table 19–1   (Cont.) Getting Started with Patching

Sequence	Action	More Information
Step 3	**Creating "Patch Plans", accumulating patches, and checking for conflicts** Create a *Patch Plan* and accumulate patches from various sources, such as patch recommendations, service requests, or even simple patch search results. *Patch Plan* is a feature that has resulted from the integration of My Oracle Support and Enterprise Manager 11g Grid Control. The feature helps you roll out patches as a group using My Oracle Support. Using a *Patch Plan*, accumulate patches, validate applicability of the patches, identify patch conflicts, and resolve them beforehand. Then, select a Deployment Procedure to apply the patches. **Note:** This requires connectivity to *My Oracle Support*.	For information about creating a *Patch Plan*, Section 20.2, "Creating a Patch Plan".
Step 4	**Running Pre-Flight Checks through Deployment Procedure in Analyze Mode** Simulate the patch execution by running the Deployment Procedures in *Analyze Mode*, and check whether the prerequisites and all other necessary requirements for patching are met. This way, you can resolve all operational conflicts and issues beforehand. **Note:** If you are using *Patch Plans* to roll out patches, then first run the Deployment Procedure in *Analyze* mode, then reselect the *Patch Plan* and associate it with the Deployment Procedure, and then run the Deployment Procedure again in *Deploy* mode. The direct flow will be offered in the future release of Enterprise Manager Grid Control.	To understand how you can run the patching Deployment Procedures in *Analyze Mode*, see Chapter 21, "Simulating Patching Deployment Procedures to Check Prerequisites".
Step 5	**Deploying Patch Using Deployment Procedures** After ensuring that the prerequisites are met, go ahead and run the Deployment Procedure to patch a target.	To understand how you can run the patching Deployment Procedure, follow the instructions outlined in the chapters of Part VI, "Patching".

19.2 Best Practices

The following are the best practices you can implemented to make the Patch Management process simpler and efficient using Deployment Procedures:

19.2.1 Customizing a Deployment Procedure

The Deployment Procedures offered by Enterprise Manager Grid Control are default procedures that have been created considering all the best practices in the industry. You can, of course, use them with the default settings to patch your targets in the environment, however, you also have the choice of customizing them to include additional custom steps, disable unwanted steps, and use authentication tools to run some steps as another user.

By customizing Deployment Procedures, you can also implement different error handling methods. For example, in a patching operation where multiple hosts are patched in parallel, it may be wise to skip the host on which a failure occurs. However, failure on a device creation could render the remaining provisioning operation ineffective. Therefore, it may be necessary to abort the entire procedure for failure of such a step.

For more information about customizing a Deployment Procedure, see Chapter 31, "Customizing Deployment Procedures".

19.2.2 Enabling Secure Authentication

Occasionally, due to secure practices, the direct access to an account of a target (for example, an `Oracle` user) might not be available. You might have deployed secure authentication software such as SUDO or Powerbroker to enable logging into the systems and carrying out the operations as a secure user (for example, `Oracle`). The Deployment Procedures support enabling secure authentication to carry out orchestration. You can do one-time customization to enable SUDO, Power broker, or any other secure authentication modules.

For information about using Privilege Delegation to set SUDO/PAM, see Section 31.7.1, "Using Privilege Delegation". Also see *My Oracle Support* note 603108.1.

19.2.3 Eliminating Patch Staging Requirement During Maintenance Downtimes

Typically, during patch rollouts using Deployment Procedures, multiple targets across hosts can be patched. In the standard mode, the Deployment Procedure stages the patch to every Oracle home and applies the patch from there.

Using the Deployment Procedure extensibility framework, you can opt to prestage the patch to a shared location amongst the servers being patched in a local hub (for example, if you have targets worldwide and if you are managing them from a central OMS, then you can prestage the patch to the regional location before patching the targets local to that location). This can be carried out as a pre-maintenance activity.

This approach eliminates the requirement for staging a patch during the maintenance window, and saves time and network bandwidth during the patching process.

The following are the two parts to this mode:

- Prestage the Patch While Running the Deployment Procedure in Analyze Mode
- Customize the Deployment Procedure to Make It 'Apply Only'

19.2.3.1 Prestage the Patch While Running the Deployment Procedure in Analyze Mode

Select a Deployment Procedure (for example, Patch Oracle Database Procedure)

In the first step of the wizard, in the **Select Stage Location** section, specify a shared stage location or central stage location (see Figure 19–1). Save the location by clicking **Save** for further usage. The location can be any NFS location or can be on a Net App filer, which is accessible by the OMS and the targets to be patched.

Figure 19–1   Specifying a Shared Stage Location

Provide the required details in the wizard, and run the Deployment Procedure in *Analyze* mode so that you can check for prerequisites and also stage the patch in the specified location.

19.2.3.2 Customize the Deployment Procedure to Make It 'Apply Only'

1. Do a *Create Like* of the out-of-box patching Deployment Procedures. For example, select the *Patch Oracle Database* procedure. Create a Deployment Procedure, which has the staging steps in the procedure disabled as shown in Figure 19–2. Save the procedure with a new name such as *Patch Oracle Database Procedure Apply Only*.

Figure 19–2   Disabling Staging Steps

2. Run the customized apply-only procedure that you saved in Step (1). In the first step of the run, select the stage location you saved or used when you ran the Deployment Procedure in Analyze mode.

Figure 19–3   Selecting Stage Location

3. Select the patch, select the targets, and then submit the procedure for execution. During the run, the patch will not get staged again, instead it will be selected from the specified pre-staged location.

19.2.4 Enabling Unattended Patching Using Remote Notifications

This section covers the following:

- Enabling Notifications
- Configuring Notification Rules

19.2.4.1 Enabling Notifications

The Deployment Procedures can be tied up with the Enteprise Manager notification systems to get notifications on the status of the procedure run. Deployment Procedures use the standard *PAF Status Notification Rule* to send out notifications. The prerequisite is to set the methods of notifications and their setup (such as SMTP server for emails) in Enterprise Manager Grid Control. Also, the standard rule can be customized to extend notifications to specific procedure runs and various methods of notification based on the requirement.

The following steps describe how you can enable notifications and get notified on the status of Deployment Procedures. This involves enabling of notifications in the required Deployment Procedures and by configuring the notification rule for specific jobs and methods of notification.

1. To enable the notification, do a *Create Like* of the out-of-box Deployment Procedure. Figure 19–4 illustrates the 'XYZ Corp Lab Database Patch Procedure', which is a copy of the out-of-box procedure "Patch Oracle Database".

Figure 19–4   Enable Notification for Deployment Procedures

2. Select **Enable Notification**. Provide a name for the "Procedure Status Notifications Job Tag". This is to create notification rules specific to a Deployment Procedure. Also, the job name can be associated to create a rule specific to this procedure and set notification methods as required. Figure 19–4 illustrates the Job tag specified as "LabSys_DBPatching_Notification", which is unique to the Deployment Procedure used to patch lab systems at XYZ Corp.

3. Select a status from **Status for which Notification is to be Sent**. The list displays the status for which you would want to be notified for this particular Deployment Procedure. In Figure 19–4, you will see that **Action Required**, **Failed**, and **Succeeded** are selected to receive the notifications.

19.2.4.2 Configuring Notification Rules

The Provisioning Framework uses the out-of-box generic "PAF Status Notification" rule to send notifications for the status of the procedure run. The standard rule can be used to send the basic e-mail notifications.

To access the notification rules, Click **Preferences**. On the Preferences page, from the left menu bar, under Notification, click **Rules**. Figure 19–5 shows how to enable *Subscribe* options for the **PAF Status Notification** rule to send e-mails. Click **Apply** after subscribing.

Figure 19–5   Subscribe to Emails for the Standard 'PAF Status Notification' Rule
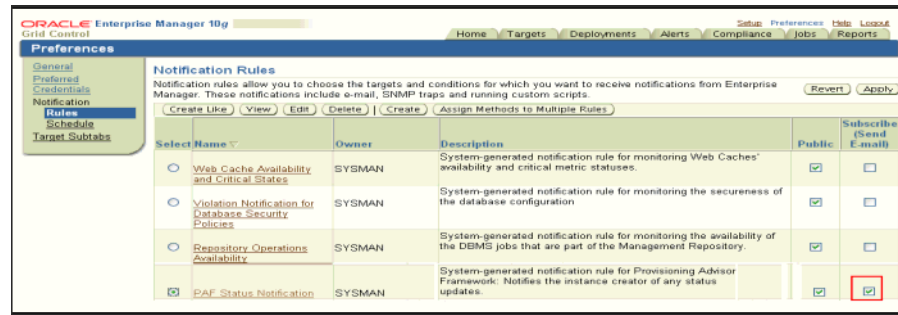
Advanced users can customize the standard **PAF Status Notification** rule to receive notifications in required ways for specific Deployment Procedures. For example, you might want to be notified by e-mail for a test system procedure, but for a production run, you might want to be informed on the status through SMS alerts.

To incorporate specific requirements and enable different methods of notification, you need to edit the standard out-of-box notification rule and the job with a specific name, and associate a specific method of notification from the predefined notification methods. The following steps explain how you can do this:

1. Do a *Create Like* of the standard **PAF Status Notification Rule**. Provide a name. For example, based on the procedure you are configuring this for, you can name it 'Lab system databases patching procedure notification' rule. See Figure 19–6.

Figure 19–6   "Create Like" of Standard PAF Notification Rule and incorporating changes

2. Select the **Jobs** subtab and edit the job to associate a job name (LabSys_ DBPatching_Notification) with this new notification rule. Provide a name as the job name tag used during the customization of the Deployment Procedure for which notification is enabled.

Figure 19–7   Edit the Job of the new PAF Status Notification Rule

Figure 19–8   Associate the job name of the Custom Deployment Procedure

3. Select the **Methods** subtab. Set the required notification method for the new rule. You can choose the basic "Send Email" option or select from various advanced notification methods available or incorporated by you. shows the notification methods, which have been created and available with Enterprise Manger Grid Control.

Figure 19–9   Set Notification Methods -Basic Email Notification and Advanced Notification Methods

4. Save the new rule. This completes the process of enabling notifications and configuring rule for a specific Deployment Procedure.

    On running the Deployment Procedure, based on the status selected under "Send Notification For Status" list, the users registered with the OMS will receive notifications on various notification methods deployed.

    > **Note:**   The notification is sent to all registered e-mail addresses under the Enterprise Manager user account. If you want this only to a specific group or individual (for example, to the DBA group or DBA who is patching), you need to create a shared Enterprise Manager user or a separate Enterprise Manager account for the DBA, and register the specific e-mail address to notification setup by logging into Enterprise Manager with the specific Enterprise Manager user credentials.

19.2.5 Using OPatch Options to Handle Issues Identified with Patches

When the Deployment Procedures are run in *Analyze* mode, they identify issues related to the patches selected. They can either be patch conflicts *(can use the plans and*

go through the merge patch process or can also handle during the execution) or other issues such as failures associated with missing components.

By default, the patching Deployment Procedures (for example, *Patch Oracle Database*) trigger OPatch in *napply* mode with the option **skip_subset**. This option skips the molecules from applying if it is a subset of any patch in the Oracle home.

The standard patching Deployment Procedures support passing of additional OPatch options from the Select Patches screen of the patching Deployment Procedure to work around the conflicts.

Figure 19–10   Advanced OPatch options UI field for additional flexibility with patching operation

However, this method is dynamic and requires users to input every time the procedure is run. If you want to incorporate these options as a part of the procedures for rollouts, you can hard-code the options in the directive step *Apply Patches*.

Select a patching Deployment Procedure, and do a *Create Like* of it. Select and edit the step *Apply Patches*, and in the Map Properties step, specify the required OPatch options. Save the procedure. See .

Figure 19–11   OPatch options as part of the Directive step

19.2.6 Handling Patch Selection During Rollouts Involving Multiple Patches

Typical patch rollouts consist of two or more patches. In order to avoid errors while selecting the patches during a patching process, you can define the patches in a predefined patch list. After doing so, you can select the patch list instead of selecting the individual patches, and thereby avoid errors.

To create a patch list and use it for rollouts, follow these steps:

1. Create a patch list file with the details of the patches you want to roll out simultaneously. The patch list file must contain the patch ID, the platform ID, and the release number of the target on which you want to apply the patch. Ensure that you separate these values using a comma.

Figure 19–12   Creating Patch List

Table 19–2 shows the platform IDs for different platforms.

Table 19–2   Platform IDs

Platform Name	Platform ID
Linux x86	46
Linux x86_64	226
Windows_NT_32 bit	912
Solaris Sparc	23
HP UX Itanium	197
HP UX PARISC	59
IBM AIX 5L	212

2. Store the patch list file either on your local host or upload it to the Software Library.

   If you want to upload it to the Software Library, then follow these steps:

   a. In Grid Control, click **Deployments**, then **Provisioning**.

   b. On the Components page, click **Create Component**.

   c. In the Describe page, select **Generic Component** from **Type** list, and specify a name for the component. Click **Next**.

   d. On the Customize page, click **Next**.

   e. On the Upload File page, select an appropriate option and upload the patch list file.

   f. On the Set Directives page, click **Next**.

   g. On the Review page, click **Finish**.

3. In Grid Control, Click **Deployments**.

4. In the Patching section, click **Patching through Deployment Procedures**.

5. On the Deployment Procedure Manager page, select a patching Deployment Procedure you want to run. Click **Schedule Deployment**.

6. On the Software Updates page, from the Standalone Database Updates section, click **Upload From File**.

7. On the Upload File page, select a source where the patch list file is available, and click **Upload**.

Figure 19–13   Selecting the Patch List

8. On the Software Update page, in the Standalone Database Updates section, you should see the patches automatically selected using the patch list file.

9. Provide the required details for the rest of the Deployment Procedure and submit it.

19.2.7 Rolling Back a Patch

Deployment Procedures, by default, do not support rolling back of a patch. You can handle it by customization, either through a simple host command step or by associating a directive step to rollback the patches. For more information, see *My Oracle Support* note 577557.1.

> **Note:**   If there are failures while applying a patch, then OPatch automatically rolls back the patch and restores the system. However, patchset or patset failure cannot be rolled back, so you must follow the best practice of taking a backup of the Oracle home before applying patchsets.

19.3 Recommendations for Testing and Implementation

The following are the recommendations for testing and implementation:

- Use Oracle Management Service 11g Release 1 (11.1) that is updated with all the latest recommended patches

- Select a canonical set of targets based on the production version

- Select standalone databases to begin with and then move to RAC configuration

- Run the Deployment Procedure in *Analyze* mode to check for prerequisites. After checking for prerequisites, run the Deployment Procedure in *Deploy* mode.

- If you have secure authentication requirement, first use direct Oracle home credentials and then configure the Deployment Procedure for SUDO

- Increase the coverage on the configurations gradually

- Introduce best practices as applicable for your rollouts

20

Patching Using My Oracle Support

My Oracle Support patch and update functionality is now tightly integrated with Enterprise Manager Grid Control. This greatly simplifies patching and provisioning by permitting easy access to patch and support information during the patch planning phase, thus allowing you to utilize the wealth of information from Oracle Support to implement the best possible patch rollout for your organization.

This chapter covers the following:

- My Oracle Support Integration with Enterprise Manager
- Creating a Patch Plan
- Deploying Patches in a Patch Plan
- Patch Information and Community Reviews

20.1 My Oracle Support Integration with Enterprise Manager

My Oracle Support provides instant access to the latest support information, forums, communities and product alerts. Beginning with Enterprise Manager Grid Control 11g, My Oracle Support patching and update functionality becomes an integral part of the console. Leveraging My Oracle Support functionality as part of your patching work flow, provides you with timely and invaluable information from users, knowledge base, and Oracle Support. Not only does this help you make important decisions regarding your patch and deployment process, it also helps you resolve conflicts quickly.

My Oracle Support patch and update functionality allows you to perform key operations that are crucial to the patching work flow. With My Oracle Support, you can:

- Review patches for applicability to your environment.
- Add patches to a plan - Each selected patch has an "Add to plan" option on the Patch toolbar.
- Validate the plan and automatically receive patches to resolve validation issues.
- Automate the deployment of the patches through Deployment Procedures. Note: Deployment Procedure functionality is only available with the Enterprise Manager Provisioning Pack.

As mentioned in Chapter 2, "Things to Know", there are some key concepts you need to be aware of when using My Oracle Support as part of your patching work flow:

- Patches and Updates
- Patch Plans

- Patch Recommendations

- Validation

- Merge Requests

For more information on these concepts, see Chapter 2, "Things to Know".

20.2 Creating a Patch Plan

Patch Plans are a convenient way to do a patch rollout using My Oracle Support.

The Patch Plan is the foundation of your patch and deployment strategy when using My Oracle Support. Briefly, a Patch Plan is a collection of patches that you want to apply as a group. You use Plans to accumulate patches and validate applicability of the patches. Depending on your patching needs, a Patch Plan might correlate with a maintenance window (you can set a date for the plan), contain patches for a single target, a collection of similarly configured targets (the same key patches applied to all of your databases, for example), or all of the patches for a system or group of systems, such as all patches for the current month for production databases.

Enterprise Manager/My Oracle Support integration provides several ways for you to find and accumulate patches to be included in a Plan. The following sections illustrate three possible ways to compile patches to add to a Patch Plan:

- From Patch Recommendations

- By Monitoring Service Requests and Selectively Adding Patches

- By Searching for Specific Patches

20.2.1 From Patch Recommendations

Patch recommendations are proactive notifications of potential system issues and recommendations that help you improve system performance and avert outages. My Oracle Support's Patch Recommendations region provides a portal to all recommended patches. From the bar graph, you can drill down to a list of recommended patch, view details about those patches, download the patches, or add them to a patch plan. A bar graph summarizes the number of issues found (one issue = one recommendation for one target). The Recommended Patches region appears by default on the Patch and Updates page and can also be added to the Dashboard. You can edit this region to filter its contents.

1. From the Oracle Enterprise Manager console, go to the **My Oracle Support** tab and select **Patches & Updates**.

2. From the Patch Recommendations region you can change the display either by classification or by target type. To view recommended patches, click on the bar graph pertaining to the desired patches, or click **All Recommendations** to display all recommendations in the **Patch Recommendation** page. The page displays all the recommendations currently available for the Enterprise Manager targets.
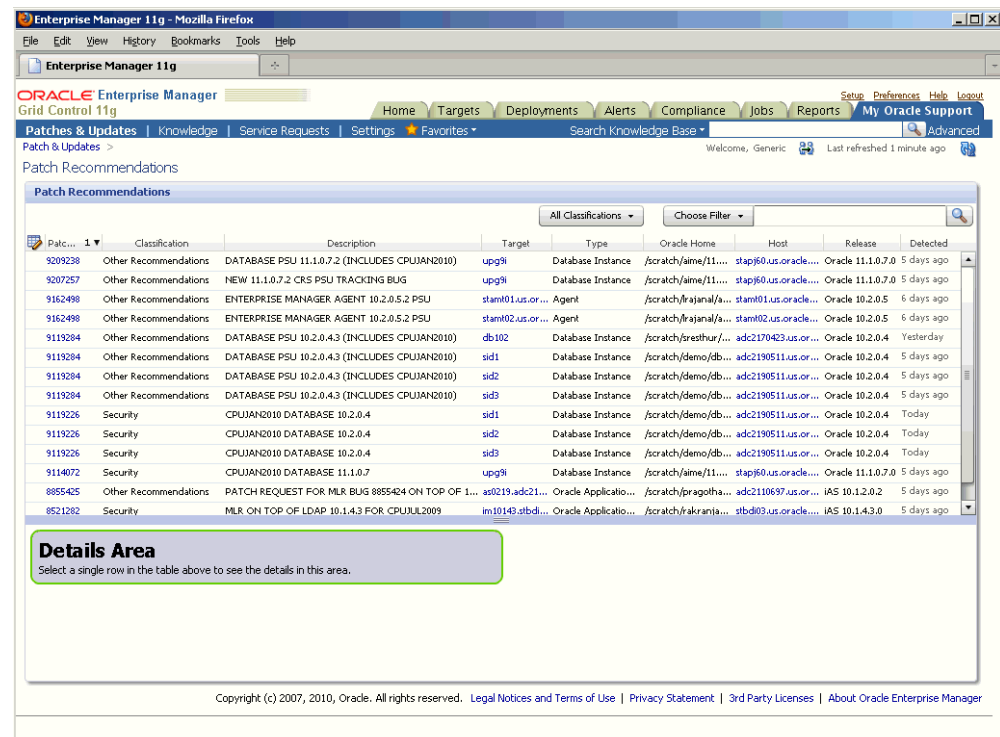
The Patch Recommendations are determined from within Enterprise Manager by directly downloading the metadata XML from My Oracle Support. User environments that are OFFLINE or DISCONNECTED can download the XML and upload to Grid Control to compute recommendations in offline mode.

3. To display information about a patch, select a patch and click 'Full Screen' on the action menu drop down.
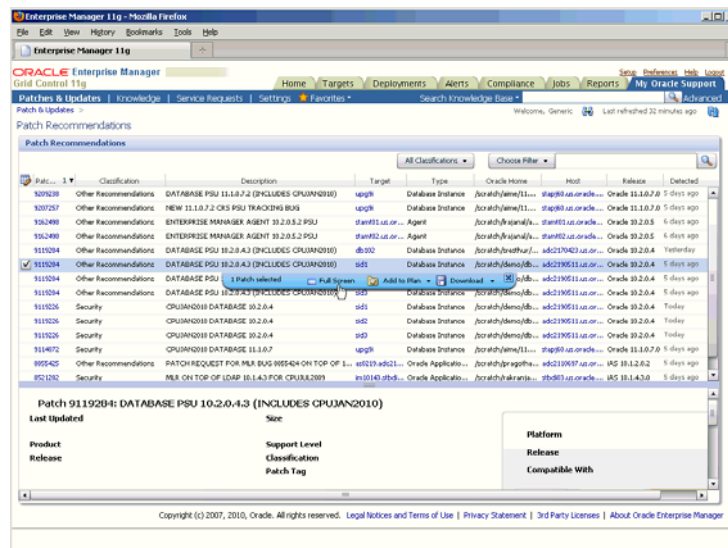
Comprehensive patch-related information such as Patch attributes, bugs fixed, related Knowledge articles are displayed.
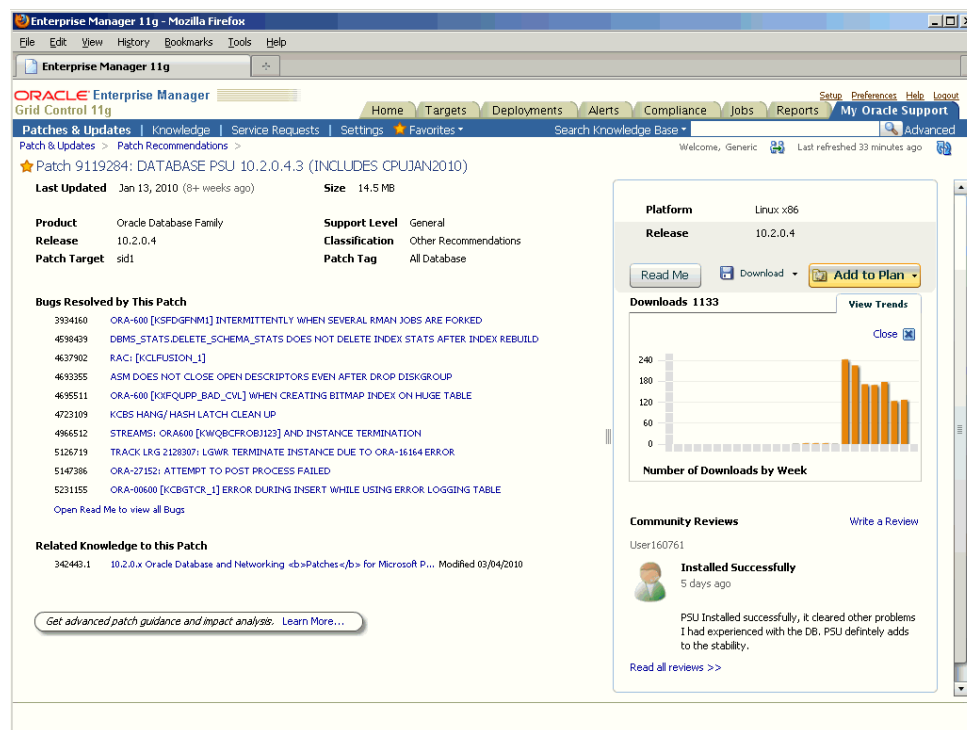
You can also obtain information regarding patch popularity with the number of downloads and also view the trend of downloads. Additionally, you can view and comment on reviews made by other members of the community.

> **Note:** patch information (such as popularity, trending, member reviews) is only available for non-Security patches.

From this page, you can add the patch to the Patch Plan.

4. Click **Add to Plan** and select **Add to New...** if you want to add the patch to a new Patch Plan, or Add to Existing to the patch to an existing Plan.

When creating a new Patch Plan, you must provide a unique name to the Plan.

5. Click **Create Plan**.

 My Oracle Support displays a verification message.

Click **View Plan** to open the Create Plan wizard and shows the selected patch under the patches section.
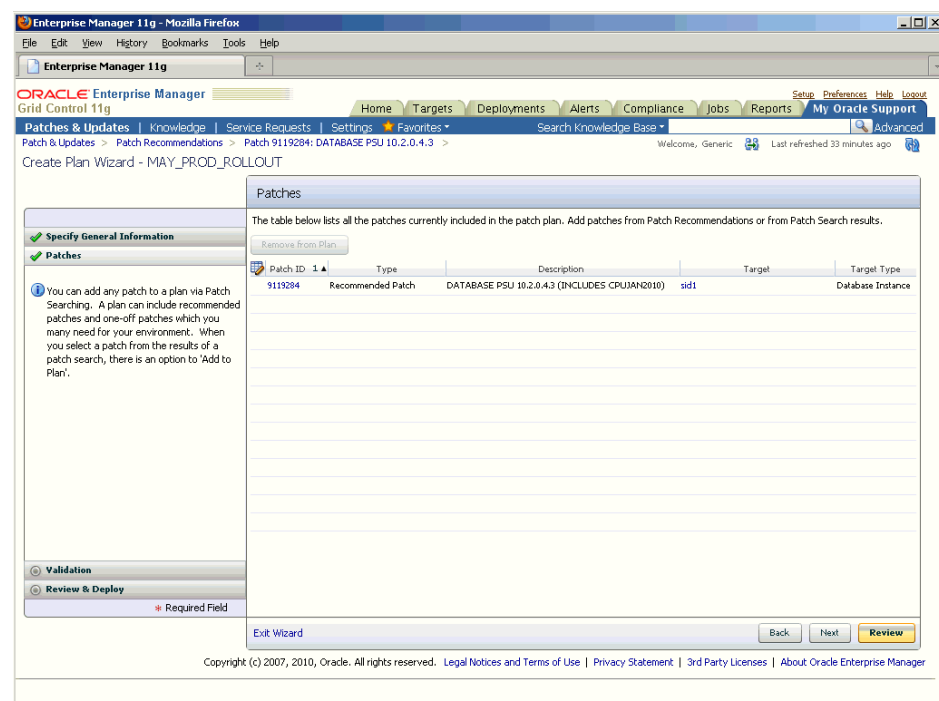
20.2.2 By Monitoring Service Requests and Selectively Adding Patches

My Oracle Support integration with Enterprise Manager provides convenient access to Service Requests and Knowledge articles. This allows you to carry out typical actions of rollouts by accumulating patches and effectively scheduling the rollout during the a designated maintenance window.

Enterprise Manager/My Oracle Support integration makes it easy to address this situation by allowing you to:

■ Create / Update / Track Service Requests and immediately associate patches with a Patch Plan and automate deployment to resolve issues.

■ Browse through Knowledge Articles providing relevant information about issues relating to problems fixed by the patch.

Example Situation

An issue is reported by a customer with an Oracle E- Business Suite 11i installation. Oracle Support verifies the customer's configuration information, made available through Oracle Configuration Manager, and promptly determines that the issue relates to the ORA-600 error on the underlying Database. As a resolution, Oracle Support recommends that a patch be applied. The lead database administrator for the customer recommends applying the patch as a part of the upcoming June rollout cycle.

1. From the **My Oracle Support** tab, click **Service Requests**.

   You can conveniently access closely monitored Service Requests by clicking the star icon associated with a particular SR.
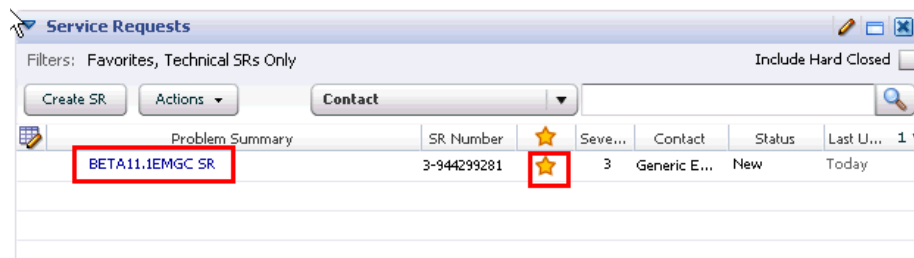
2. To view the most recent updates to a Service Request (as might be made by Oracle Support), click on the SR's Problem Summary. Comprehensive information about that SR is displayed.

3. To view the most recent updates to the SR, see the History region. This region logs all SR activity, including any patches recommended by Oracle Support.
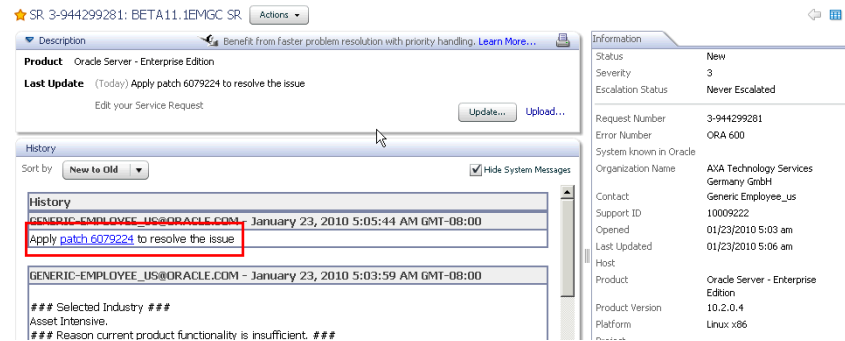
4. To view specific patch information, click on the specific patch link in the History table. Comprehensive information about the selected patch is displayed.
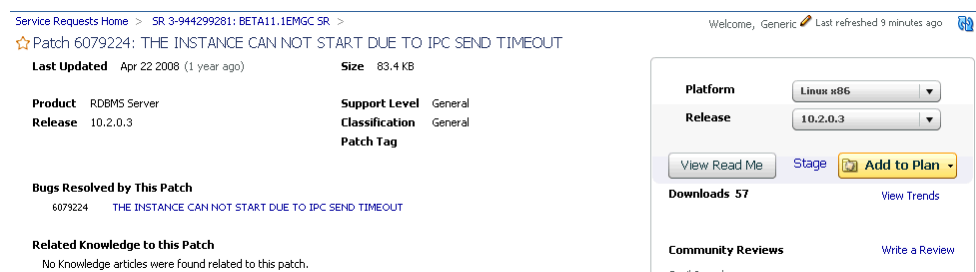
From this page, you have access to a explicit patch information such as customer download trends and Community Reviews that provide useful information about patch issues and resolutions from other customers.

5. When ready, you can add the patch to your Patch Plan by clicking **Add to Plan** and then selecting **Add to Existing** if you are adding the patch to an existing Plan, or **Add to New...** if you are creating an entirely new Patch Plan. For the sake of procedural flow, we will assume that you already have a Patch Plan and that you have selected the **Add to Existing** option. The **Add to Plan** window appears.
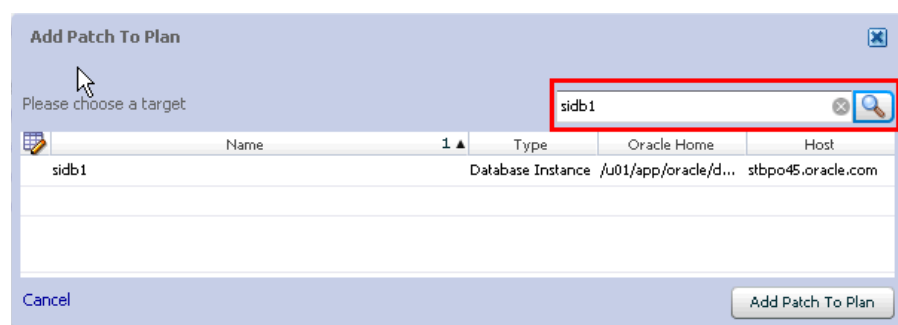
You are now ready to associate a target with the patch.

6. In the **Add to Plan** window, choose a target. You can also search for a specific target if you do not know the target name.

7. Once you have associated the desired targets with the patch, you are now ready to add the patch to the Patch Plan. Click **Add Patch to Plan**.

20.2.3 By Searching for Specific Patches

The most common method for finding patches to be added to a Patch Plan is by simply searching for them. My Oracle Support's advanced patch search functionality allows you to perform more flexible and advanced searches. It also provides capabilities such as saving a search that's used routinely, and searching based on existing 'saved' searches. All of this lets you perform searches more quickly and efficiently.
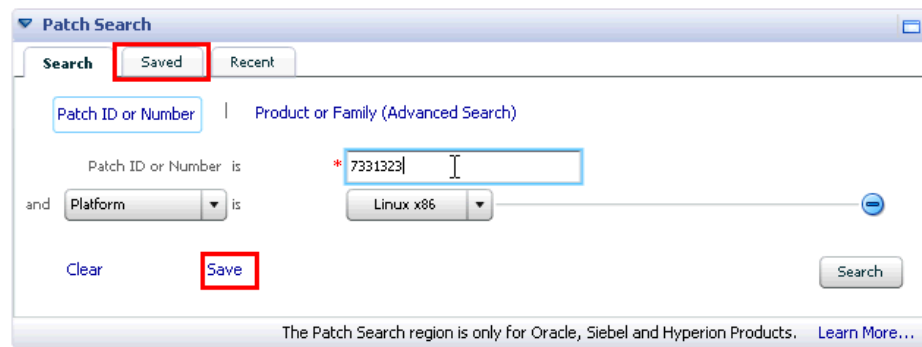
Typical patch search scenarios would be where a DBA learns the existence of a patch from external sources such as blogs, Oracle technology forums, or from colleagues. My Oracle Support patch search and download functionality is tightly integrated with Patch Plan creation/update functionality. The DBA can now search and add the patches directly to the Patch Plan for further Validation in preparation for rollout, thus creating a seamless workflow.

1. From the **My Oracle Support** tab, click **Patches & Updates**.
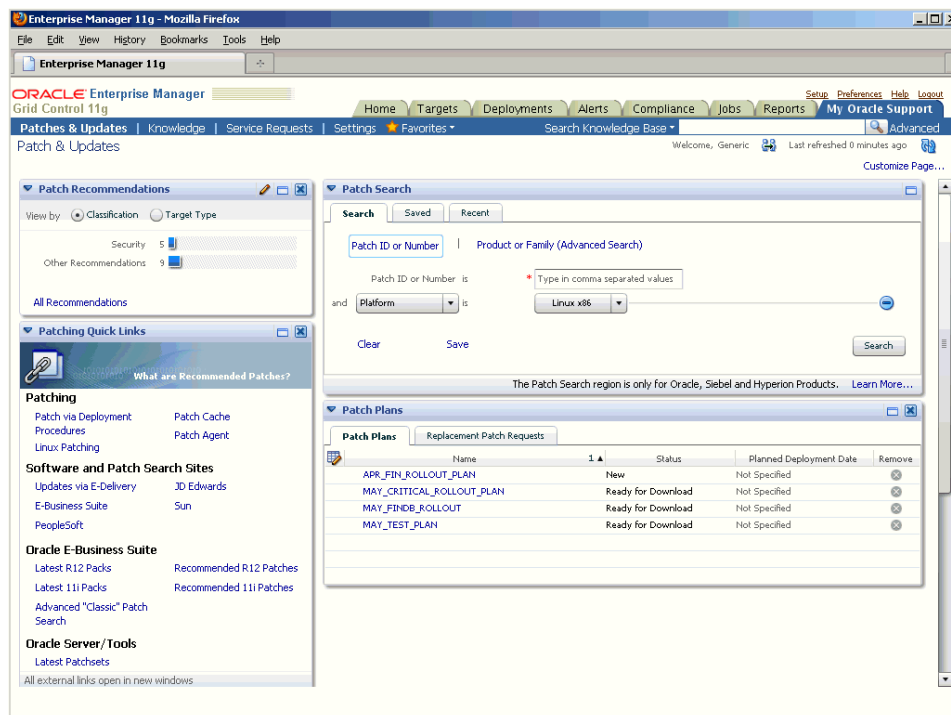
2. From the **Patch Search** region, enter the search parameters you want to use and click **Search**.

Alternatively, you can search using any previously save searches. My Oracle Support search functionality maintains a list of the most recent searches. If you know that you recently performed an identical (or similar) search, you can access the list of searches on the **Recent** tab of the **Patch Search** region.

3. Once the patch search is complete, the search results are displayed in the **Patch Search Results** region.

4. Select one or more patches and click **Add to Plan** from the context menu and choose either to add them to a new plan or to an existing plan.

For the sake of task flow, we will assume you are adding the patches) to an existing Patch Plan.

5. From the **Add to Plan** dialog, associate the patch with the intended target.

6. Click **Add Patch to Plan**. My Oracle Support will display a message verifying that the patch has been added to a Patch Plan.

7. Click **View Plan** to access the Patch Plan wizard. All patches accumulated through searches, Oracle Support Recommendations, and/or Service Requests for a specific rollout will be displayed

20.2.4 Resolving Patch Conflicts

Ultimately, when attempting to accumulate patches to rollout to heterogeneous IT environments, conflicts may arise between patches that are the result of such issues as version incompatibility between patches, or perhaps incompatibility with existing host software. My Oracle Support simplifies and automates patch conflict resolution.

As covered previously, the Patch Plan wizard allows you to view or modify any aspect of the Patch Plan.

The most important function of the wizard is to validate the Patch Plan. When My Oracle Support validates a plan, it checks for conflicts using patch information from Oracle, the inventory of patches on your system (gathered by the Configuration Manager) and information from the candidate patches. In addition to checking patch merge conflicts between patches listed in the plan, the validation function also checks for conflicts between the listed plan patches and patches that have already been applied.

1. From the Patch Plan wizard, click **Validation** or **Next**. When new patches are added to an existing Patch Plan, click **Revalidate**.

When validating, a request to check if the selected patch can be applied to targets is sent to My Oracle Support. Upon processing, there are three possible outcomes:

- There are no conflicts, then the selected patches are "Ready for Deployment"

- Conflict exists: You can directly file a Merge Request to obtain a conflict free patch. If there is an open request filed by another customer, you will be automatically added to the existin request.

- Upon Conflicts, if there is a Merge Patch already available, you can directly opt to replace the conflicting patches with the Merge Patch

For the sake of task flow, we will assume there is a conflict.

2. Upon validation, My Oracle Support finds that there are conflicts between the two patches. Click **Request Replacement Patches**.
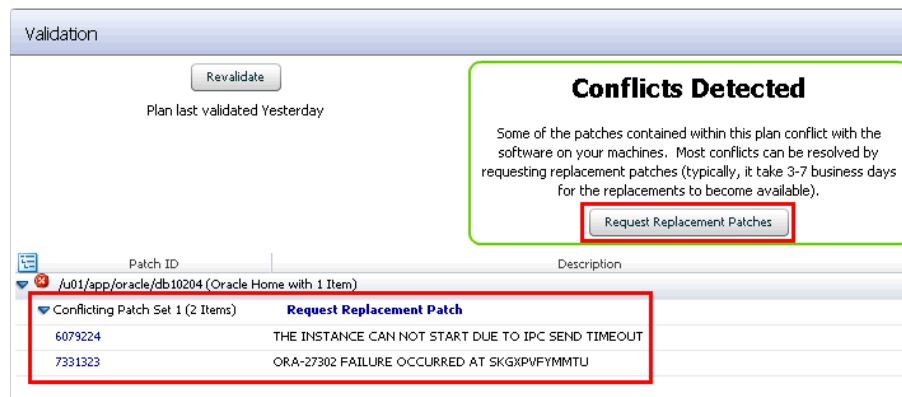
3. In this situation, a Merge Patch is already available. Click **Replace Conflicting Patches**.

Conflicts Found, Replacements Available

Replacement patches are available for some or all the conflicts in your patching plan.

[Replace Conflicting Patches]

> **Note:** Upon filing a Merge Request, a numeric request ID is automatically generated. This ID can be tracked in the Patch Plans section under **Replacement Patch Requests**.

20.3 Deploying Patches in a Patch Plan

In the previous sections, you created patch plans by obtaining patches from a variety of sources and strategies. Once you have a valid patch plan containing both patches and desired targets, you are ready to deploy the patches throughout your IT environment.

The following procedure assumes that your Patch Plan contains both patches and their associated targets.

1. From the Patch Plan region, click on a plan containing the patches you want to deploy.

The Patch Plan wizard appears.

2. Click **Review and Deploy**.

3. Select the patches you want to deploy.

4. Click **Run Procedure** from the context menu.

The Patch Plan Deployment verification dialog appears.
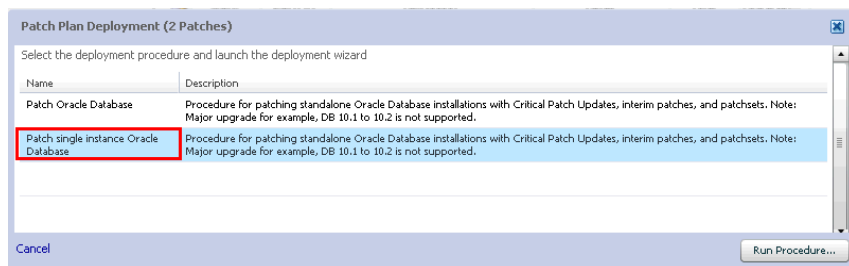
5. Select a procedure and click **Run Procedure**. This will trigger the automated deployment procedure to deploy the patches. You will be taken to the Software Updates page prepopulated with the selected patches.
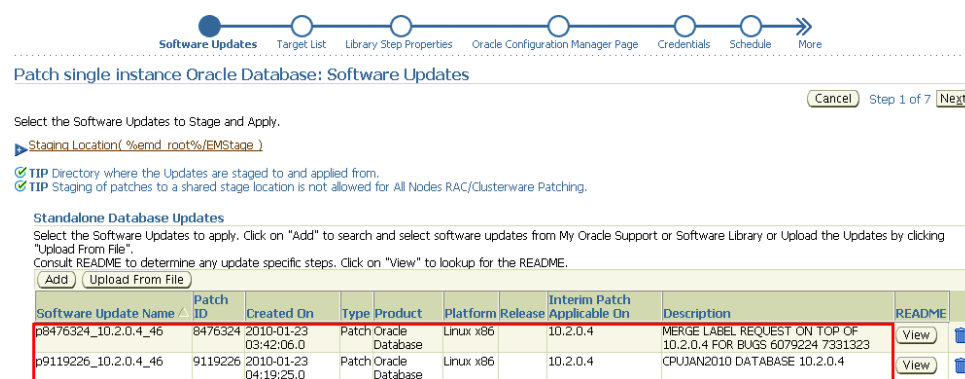
6. Click Next to go to the Target List page. Enterprise Manager will display the following message:

   Currently Targets associated with the patches in the PLAN are not auto selected. Search and re-select the targets. Target that were in the plan are -'sid1','sid2'.

   From this point forward, you follow the normal deployment procedure as outlined in Section 22.4, "Patching Standalone Oracle Database in Online Mode".

20.4 Patch Information and Community Reviews

My Oracle Support provides ready access to comprehensive information about patches and related support information, thus allowing you to make an accurate determination of your patching requirements and simplifying patch deployment decisions. An invaluable source if information is from the experiences from other customers who have already applied patches you are planning to deploy.

The patch home page for a given patch provides you with access to this patch community information.
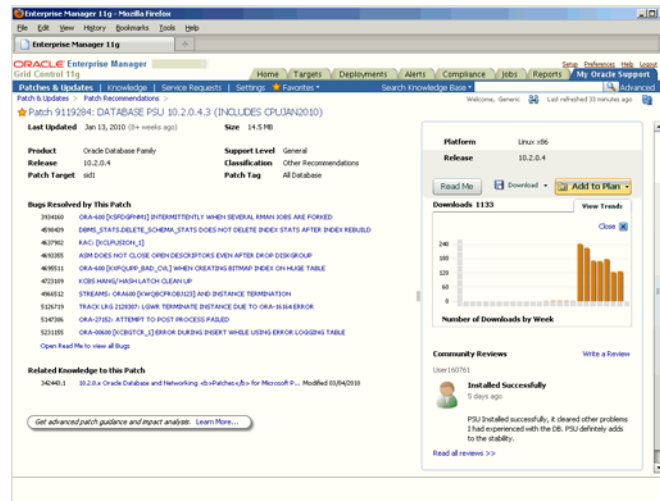
The patch home page displays the number of downloads. Click **View Trends** to see the customer download trends in a bar graph.
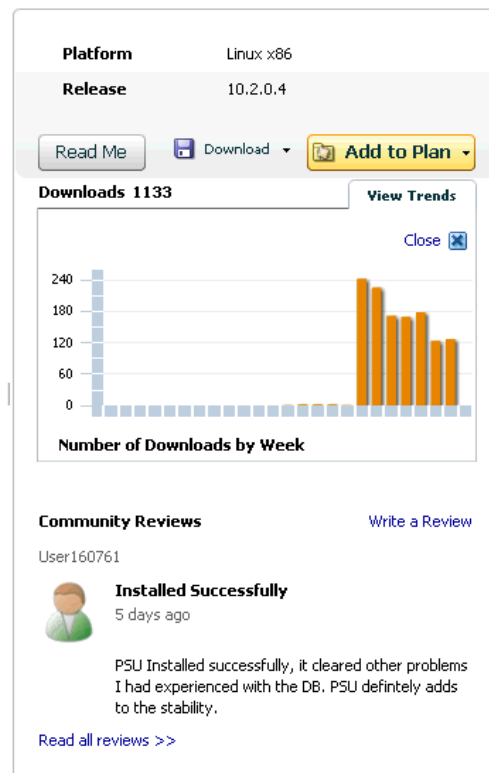
As shown in the previous graphic, Community Reviews provide pertinent information from other users about the patch and their comments about the issues and resolutions.

21

Simulating Patching Deployment Procedures to Check Prerequisites

Enterprise Manager Grid Control offers patching Deployment Procedures that help you patch targets such as Oracle Database, Oracle RAC, Oracle ASM, and Oracle Clusterware. However, before you run the actual patching Deployment Procedures, Oracle strongly recommends you to simulate the Deployment Procedures so that you check the prerequisites and ensure that you meet all the necessary requirements for patching, and also resolve all operational conflicts and issues beforehand.

Table 21–1 lists the patching Deployment Procedures that support the simulation mode, and describes how you must run them in that mode.

Table 21–1   Simulating Patching Deployment Procedures for Checking Prerequisites

Deployment Procedure Name	Support or Not Supported	Instructions to Simulate
Patch Oracle Database	Yes	Run the Deployment Procedure as described in Chapter 22, "Patching Oracle Database", but on the Review page (final page of the wizard), click **Analyze**.
Patch Oracle RAC Database - Rolling	Yes	Run the Deployment Procedure as described in Chapter 23, "Patching Oracle RAC", but on the Review page (final page of the wizard), click **Analyze**.
Patch Oracle RAC Database - All Nodes	Yes	Run the Deployment Procedure as described in Chapter 23, "Patching Oracle RAC", but on the Review page (final page of the wizard), click **Analyze**.
Patch Oracle Clusterware - Rolling	Yes	Run the Deployment Procedure as described in Chapter 24, "Patching Oracle Clusterware", but on the Review page (final page of the wizard), click **Analyze**.
Patch Oracle Clusterware - All Nodes	Yes	Run the Deployment Procedure as described in Chapter 24, "Patching Oracle Clusterware", but on the Review page (final page of the wizard), click **Analyze**.
Patch Standalone Oracle ASM	Yes	Run the Deployment Procedure as described in Chapter 25, "Patching Oracle Automatic Storage Management (ASM)", but on the Review page (final page of the wizard), click **Analyze**.
Patch Application Server	No	

Table 21–1 (Cont.) Simulating Patching Deployment Procedures for Checking

Deployment Procedure Name	Support or Not Supported	Instructions to Simulate
Patch Linux Hosts	No	
Patch Solaris Hosts	No	
Patch Windows Hosts	No	

22

Patching Oracle Database

This chapter explains how you can patch Oracle Database. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedures

- Supported Releases

- Patching Standalone Oracle Database in Online Mode

- Patching Standalone Oracle Database in Offline Mode

22.1 Getting Started

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

Table 22–1 Getting Started with Patching Oracle Database

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Oracle Database. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 22.2, "Deployment Procedures".
Step 2	**Knowing About the Supported Releases** Know what releases of Oracle Database can be patched by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 22.3, "Supported Releases".

Table 22–1   (Cont.) Getting Started with Patching Oracle Database

Step	Description	Reference Links
Step 3	**Selecting the Patching Mode** Identify whether the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support. If it has, then choose to patch in online mode. Otherwise, choose to patch in offline mode.	■ To learn about patching a standalone Oracle Database in online mode, see Section 22.4, "Patching Standalone Oracle Database in Online Mode". ■ To learn about patching standalone Oracle Database in offline mode, see Section 22.5, "Patching Standalone Oracle Database in Offline Mode".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the environment for online or offline patching, applying mandatory patches, setting up of Oracle Software Library, and so on.	■ To learn about the prerequisites for patching a standalone Oracle Database in online mode, see Section 22.4.1, "Prerequisites". ■ To learn about the prerequisites for patching a standalone Oracle Database in offline mode, see Section 22.5.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Oracle Database.	■ To patch a standalone Oracle Database in online mode, follow the steps explained in Section 22.4.2, "Patching Procedure". ■ To patch a standalone Oracle Database in offline mode, follow the steps explained in Section 22.5.2, "Patching Procedure".

22.2 Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for patching Oracle Database:

■ *Patch Oracle Database*

> **Note:**
>
> - If the Oracle home of the database you are patching also has an ASM installed, then the Deployment Procedure patches only the database instance, but appropriately shuts down the ASM instance before patching the database and restarts it after the operation is complete.
>
> - The Deployment Procedure leaves several .log and .lst files per database instance in the Oracle home of the database being patched. For example, with a databases named UPGRD1, the following files will be left in the oracle home after patching:
>
>   -rw-r----- 1 oraets92 oinstall 0 Mar 27 12:19   UPGRD1_is_ running.lst
>
>   -rw-r----- 1 oraets92 oinstall 86 Mar 27 12:18   UPGRD1_ shutdown.log
>
>   -rw-r----- 1 oraets92 oinstall 401 Mar 27 12:19 UPGRD1_ shutdown.lst
>
>   You cannot avoid the creation of these log files. If you want to delete them, you can either navigate to the Oracle home directory and manually delete them or customize the Deployment Procedure to include a step to automatically delete them.

22.3 Supported Releases

Using this Deployment Procedure, you can patch the following releases of Oracle Database:

- Oracle Database 9i Release 2 (9.2.0.8)

- Oracle Database 10g Release 1 (10.1.x.x)

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

- Oracle Database 11g Release 2 (11.2.x.x)

22.4 Patching Standalone Oracle Database in Online Mode

This section describes how you can patch a standalone Oracle Database in online mode, that is, by connecting to My Oracle Support. This option is best suited when you have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites

- Patching Procedure

22.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

  - Section 3.3, "Setting Up Infrastructure for Online Patching"

- Ensure that you run the patchability report to understand whether the Oracle Database you want to patch is suitable for this patching operation. In the report, if the Oracle Database you want to patch appears as a target that cannot be patched, then fix the target configuration issues to make it a target that can be patched. For more information about the patchability report, and to understand how you can diagnose the issues with an Oracle Database target and take corrective actions, see Appendix H, "About Patching Automation Reports".

- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

22.4.2 Patching Procedure

To patch a standalone Oracle Database, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle Database**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

   a. In the Select Stage Location section, you can retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

This default value refers to a location on the target host where the database to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched.

b. In the Standalone Database Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search My Oracle Support**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control connects to My Oracle Support, searches the patches, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. *(Optional)* In the Apply SQL Script section, retain the default selection so that the default SQL script that is packaged with the patch is used for modifying the schemas.

> If you want to run any custom script, then select **Enter the script to apply SQL** and specify the full path to the location where the custom script is available. For example, `%oracle_home%/rdbms/admin/catupgrd.sql`.
>
> Enterprise Manager Grid Control automatically runs certain scripts such as Catcpu.sql, utlrp.sql for CPUs; Catpsu.sql, utlrp.sql for PSUs; and catupgrade.slq for Patch Sets. However, some scripts mentioned in the ReadMe of a patch might have to be run manually. To run those scripts, you can select **Enter the script to apply SQL** and specify the full path to the location where the custom scripts are available.
>
> The variables you can use for the directory location are %emd_root% to indicate a location on the target host, `%oracle_home%` to indicate the current home location being patched, or `%perlbin%` to indicate the location of perl binary used by the Management Agent.
>
> You have an option of running both the scripts by selecting both these options or one of the scripts by selecting only one of these options. If you select both these options, then both the scripts are run, that is, the default script runs first and then the custom script runs.

d. *(Optional)* In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the database targets.

> Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

e. *(Optional)* In the Black Out Associated Targets section, retain the default selection so that all targets associated with the database to be patched are blacked out while the patching is in progress.

> For example, if the Oracle Database to be patched is running in an Oracle home location where Oracle ASM is running, then before patching that Oracle Database, all associated targets such as Oracle Listener and Oracle ASM instances will be blacked out.

f. *(Optional)* In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

g. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you identified any patch conflicts, then you can pass options to rectify the patch conflicts.

> **Advanced OPatch Options**
> Specify the advanced options to be passed to the Apply Patch Step.
> Options [   ]

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

`-force_conflict`, applies patches even if there are conflicts

Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

h. Click **Next**.

5. On the Target List page, in the Database Targets to be Patched section, click **Add** to add the database targets you want to patch. Then click **Next**.

> **Database Targets to be patched**
>
> (Add)
>
> | Name △ | Target Type | Host Name | Status | |
> |--------|-------------|-----------|--------|---|
> | horcl_horcl1 (Unpatched) | Database Instance | stamy04.us.oracle.com | ⇧ | 🗑 |
>
> (Remove All)

Note: Selection of multiple targets to be patched in a single operation is supported. Typically, this procedure orchestrates patching in parallel across multiple hosts and does a sequential operation across the Oracle homes within a host. Within a given Oracle home, the procedure patches all of the Databases simultaneously.

---

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre- filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch database Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the databases are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

9. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

- If you want to check for prerequisites, then click **Analyze**.

- If you have already checked for prerequisites or if you do not want to check for prerequisites, then click **Deploy** to patch the selected databases.

22.5 Patching Standalone Oracle Database in Offline Mode

This section describes how you can patch a standalone Oracle Database in offline mode, that is, without connecting to My Oracle Support. This option is best suited when you do not have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites
- Patching Procedure

22.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  – Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  – Section 3.2, "Optional Infrastructure Requirements"

  – Section 3.4, "Setting Up Infrastructure for Offline Patching"

- Ensure that you run the patchability report to understand whether the Oracle Database you want to patch is suitable for this patching operation. In the report, if the Oracle Database you want to patch appears as a target that cannot be patched, then fix the target configuration issues to make it a target that can be patched. For more information about the patchability report, and to understand how you can diagnose the issues with an Oracle Database target and take corrective actions, see Appendix H, "About Patching Automation Reports".

- Ensure that you run the Deployment Procedure in Analyze mode to run prerequisite checks.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".
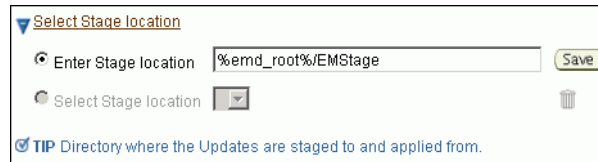
22.5.2 Patching Procedure

To patch a standalone Oracle Database, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle Database**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

This default value refers to a location on the target host where the database to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched.

b. In the Standalone Database Updates section, click **Add** to search and select software updates from Oracle Software Library (Software Library). Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search Software Library**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control searches the patches in the Software Library, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. In the Apply SQL Script section, retain the default selection so that the default SQL script that is packaged with the patch is used for modifying the schemas.

If you want to run any custom script, then select **Enter the script to apply SQL** and specify the full path to the location where the custom script is available. For example, `%oracle_home%/rdbms/admin/catupgrd.sql`.

Enterprise Manager Grid Control automatically runs certain scripts such as Catcpu.sql, utlrp.sql for CPUs; Catpsu.sql, utlrp.sql for PSUs; and

catupgrade.slq for Patch Sets. However, some scripts mentioned in the ReadMe of a patch might have to be run manually. To run those scripts, you can select **Enter the script to apply SQL** and specify the full path to the location where the custom scripts are available.

The variables you can use for the directory location are %emd_root% to indicate a location on the target host, `%oracle_home%` to indicate the current home location being patched, or `%perlbin%` to indicate the location of perl binary used by the Management Agent.

You have an option of running both the scripts by selecting both these options or one of the scripts by selecting only one of these options. If you select both these options, then both the scripts are run, that is, the default script runs first and then the custom script runs.

d. In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the database targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

e. In the Black Out Associated Targets section, retain the default selection so that all targets associated with the database to be patched are blacked out while the patching is in progress.

For example, if the Oracle Database to be patched is running in an Oracle home location where Oracle ASM is running, then before patching that Oracle Database, all associated targets such as Oracle Listener and Oracle ASM instances will be blacked out.

f. In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

g. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you identified any patch conflicts, then you can pass options to rectify the patch conflicts.

The following are the options you can specify:

-skip_duplicate, skips duplicate patches

-skip_conflict, skips conflicting patches

-skip_missing_component, skips patches whose components are missing

-invPtrLoc, the location of oraInst.loc file

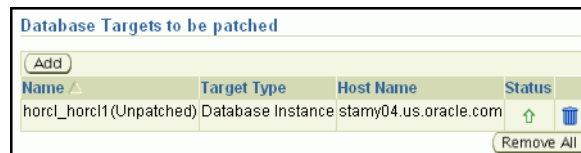-force_conflict, applies patches even if there are conflicts

> **Note:** If you want to specify multiple options, then separate the list with a white space. For example, -skip_duplicate -skip_conflict

   h. Click **Next**.

5. On the Target List page, in the Database Targets to be Patched section, click **Add** to add the database targets you want to patch. Then click **Next**

> **Note:** Selection of multiple targets to be patched in a single operation is supported. Typically, this procedure orchestrates patching in parallel across multiple hosts and does a sequential operation across the Oracle homes within a host. Within a given Oracle home, the procedure patches all of the Databases simultaneously.

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps you customized.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre-filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch database Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the databases are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

9. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

- If you want to check for prerequisites, then click **Analyze**.

- If you have already checked for prerequisites or if you do not want to check for prerequisites, then click **Deploy** to patch the selected databases.

23

Patching Oracle RAC

This chapter explains how you can patch Oracle RAC. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedures
- Core Components Patched
- Supported Releases
- Patching Oracle RAC in Online Mode
- Patching Oracle RAC in Offline Mode

23.1 Getting Started

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

Table 23–1   Getting Started with Patching Oracle RAC

Step	Description	
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Oracle RAC. Know how the Deployment Procedure functions, what use cases it covers, what core components are patched, and so on.	To learn about the Deployment Procedure, see Section 23.2, "Deployment Procedures". To learn about the core components that are patched, see Section 23.3, "Core Components Patched".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle RAC can be patched by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 23.4, "Supported Releases".

Table 23–1 (Cont.) Getting Started with Patching Oracle RAC

Step	Description	
Step 3	**Selecting the Patching Mode** Identify whether the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support. If it has, then choose to patch in online mode. Otherwise, choose to patch in offline mode.	■ To learn about patching Oracle RAC in online mode, see Section 23.5, "Patching Oracle RAC in Online Mode". ■ To learn about patching Oracle RAC in offline mode, see Section 23.6, "Patching Oracle RAC in Offline Mode".
Step 4	**Patching RAC Nodes Individually or Collectively** Decide whether you want to patch the nodes one by one, using the *Rolling* methodology, or all at a time, in parallel, using the *All Nodes* methodology.	**Online Mode:** ■ To learn about patching Oracle RAC nodes individually in online mode, see Section 23.5.1, "Patching Oracle RAC Nodes Individually". ■ To learn about patching Oracle RAC nodes collectively in online mode, see Section 23.5.2, "Patching Oracle RAC Nodes Collectively". **Offline Mode:** ■ To learn about patching Oracle RAC nodes individually in offline mode, see Section 23.6.1, "Patching Oracle RAC Nodes Individually". ■ To learn about patching Oracle RAC nodes collectively in offline mode, see Section 23.6.2, "Patching Oracle RAC Nodes Collectively".
Step 5	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the environment for online or offline patching, applying mandatory patches, setting up of Oracle Software Library, and so on.	**Online Mode:** ■ To learn about the prerequisites for patching Oracle RAC nodes individually in online mode, see Section 23.5.1.1, "Prerequisites". ■ To learn about the prerequisites for patching Oracle RAC nodes collectively in online mode, see Section 23.5.2, "Patching Oracle RAC Nodes Collectively". **Offline Mode:** ■ To learn about the prerequisites for patching Oracle RAC nodes individually in offline mode, see Section 23.6.1.1, "Prerequisites". ■ To learn about the prerequisites for patching Oracle RAC nodes collectively in offline mode, see Section 23.6.2, "Patching Oracle RAC Nodes Collectively".

Table 23–1   (Cont.) Getting Started with Patching Oracle RAC

Step	Description	
Step 6	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Oracle RAC.	**Online Mode:** ■ To patch Oracle RAC nodes individually in online mode, see Section 23.5.1.2, "Patching Procedure". ■ To patch Oracle RAC nodes collectively in online mode, see Section 23.5.2, "Patching Oracle RAC Nodes Collectively". **Offline Mode:** ■ To patch Oracle RAC nodes individually in offline mode, follow the steps explained in Section 23.6.1.2, "Patching Procedure". ■ To patch Oracle RAC nodes collectively in offline mode, follow the steps explained in Section 23.6.2, "Patching Oracle RAC Nodes Collectively".

23.2 Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for patching Oracle RAC:

■ *Patch Oracle RAC Database - Rolling*

■ *Patch Oracle RAC Database - All Nodes*

While the former one allows you to patch the nodes one by one, the latter one allows you to patch all the nodes at a time, in parallel. For more information about these patching methodologies and circumstances under which you must use the *Rolling* methodology or *All Nodes* methodology, see Section 2.4.3, "Understanding Patching Methodologies".

23.3 Core Components Patched

When you patch Oracle RAC, essentially, the Deployment Procedures patch the following core components:

■ Oracle RAC Database

■ Optionally, Automatic Storage Management (ASM)

23.4 Supported Releases

Using these Deployment Procedures, you can patch the following releases of Oracle RAC:

■ Oracle Clusterware 10g Release 1 (10.1.x.x)

■ Oracle Clusterware 10g Release 2 (10.2.x.x)

■ Oracle Clusterware 11g Release 1 (11.1.x.x)

■ Oracle Clusterware 11g Release 2 (11.2.x.x)

IMPORTANT:

- You cannot run this procedure on Oracle RAC installations that are on third-party clusterware (that is, other than Oracle Clusterware).

- You can run this procedure on Oracle RAC installations that have shared Oracle homes, as long as they are not mixed-shared.

  For example, let us consider that the Oracle RAC installation is cluster1 and it has four nodes, mainly node1, node2, node3, and node4. You can run this procedure on this cluster if all the nodes are part of the same shared Oracle home, say /nfs_1/home/oracle/db10g. However, you cannot run this procedure if two of the nodes are from one shared Oracle home, say /nfs_1/home/oracle/db10g, and the other two are from another shared Oracle home, say /nfs_2/home/oracle/db10g.

- This deployment procedure supports patching of shared Oracle home RAC installations only for Oracle Real Application Clusters (RAC) 10g Release 2 (10.2.0.1) or higher.

23.5 Patching Oracle RAC in Online Mode

This section describes how you can patch Oracle RAC in online mode, that is, by connecting to My Oracle Support. This option is best suited when you have an Internet connection on the host where Enterprise Manager Grid Control is running.

In online mode, you can either patch the nodes one by one, or patch all the nodes at a time, in parallel. This section provides instructions for both these methodologies.

In particular, this section provides instructions for the following:

- Patching Oracle RAC Nodes Individually
- Patching Oracle RAC Nodes Collectively

23.5.1 Patching Oracle RAC Nodes Individually

This section describes how you can patch the nodes of Oracle RAC one by one, using the *Rolling* methodology.

This methodology is best suited when you are applying one-off patches that support this methodology, and when you want to maintain high availability of your targets, so when one node is being patched, the other nodes are available for service. For example, if you are patching a clusterware that has five nodes, then the first node is shut down, patched, and restarted, and then the process is rolled over to the next node, and so on, until all the nodes are patched successfully.

This section covers the following:

- Prerequisites
- Patching Procedure

- You cannot use this methodology to patch shared Oracle homes. If you want to patch a shared Oracle home, then use Deployment Procedure that support the *All Nodes* methodology. For more information, see Section 23.5.2, "Patching Oracle RAC Nodes Collectively".

- Although most one-off patches support this methodology, there are some one-off patches that can be applied only using the *All Nodes* methodology. You can identify the supported methodology while downloading the one-off patches from My Oracle Support.

- If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

23.5.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

  - Section 3.3, "Setting Up Infrastructure for Online Patching"

- Ensure that operating system users such as *oracle* and *crsuser* are available on all the nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* or *oinstall*.

- Ensure that you run the patchability report to understand whether the Oracle Database you want to patch is suitable for this patching operation. In the report, if the Oracle Database you want to patch appears as a target that cannot be patched, then fix the target configuration issues to make it a target that can be patched. For more information about the patchability report, and to understand how you can diagnose the issues with an Oracle Database target and take corrective actions, see Appendix H, "About Patching Automation Reports".

- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

■ Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

23.5.1.2 Patching Procedure
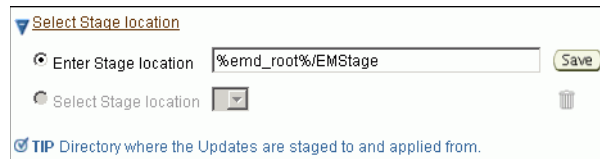
To patch Oracle Real Application Clusters, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle RAC Database - Rolling**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

   a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

      This default value refers to a location on the target host where the Oracle RAC to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

      You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

      Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched.

   b. In the RAC Database Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

      On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search My Oracle Support**, specify the details of the patches you want to search, and click **Go**.

      Enterprise Manager Grid Control connects to My Oracle Support, searches the patches, and displays details about them on the Search and Select Software

Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. *(Optional)* In the Apply SQL Script section, retain the default selection so that the default SQL script that is packaged with the patch is used for modifying the schemas.

If you want to run any custom script, then select **Enter the script to apply SQL** and specify the full path to the location where the custom script is available. For example, `%oracle_home%/rdbms/admin/catupgrd.sql`.

Enterprise Manager Grid Control automatically runs certain scripts such as Catcpu.sql, utlrp.sql for CPUs; Catpsu.sql, utlrp.sql for PSUs; and catupgrade.slq for Patch Sets. However, some scripts mentioned in the ReadMe of a patch might have to be run manually. To run those scripts, you can select **Enter the script to apply SQL** and specify the full path to the location where the custom scripts are available.

The variables you can use for the directory location are %emd_root% to indicate a location on the target host, `%oracle_home%` to indicate the current home location being patched, or `%perlbin%` to indicate the location of perl binary used by the Management Agent.

You have an option of running both the scripts by selecting both these options or one of the scripts by selecting only one of these options. If you select both these options, then both the scripts are run, that is, the default script runs first and then the custom script runs.

d. *(Optional)* In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target host is upgraded before the patches are applied on the database targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

e. *(Optional)* In the Black Out Associated Targets section, retain the default selection so that all targets associated with the clusterware to be patched are blacked out while the patching is in progress.

    f. *(Optional)* In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

    g. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you identified any patch conflicts, then you can pass options to rectify the patch conflicts.

    The following are the options you can specify:

    `-skip_duplicate`, skips duplicate patches

    `-skip_conflict`, skips conflicting patches

    `-skip_missing_component`, skips patches whose components are missing

    `-invPtrLoc`, the location of oraInst.loc file

> **Note:** If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

    h. Click **Next**.

5. On the Target List page, in the RAC Database Targets to be Patched section, click **Add** to add the Oracle RAC database targets you want to patch. If you have multiple Oracle RAC databases in a cluster and if you are not sure of the selection, then search for all RAC databases in that cluster by selecting **Cluster** from the Target Type list. Then click **Next**.

> **Note:** Selection of multiple targets to be patched in a single operation is supported. Typically, this procedure orchestrates patching in parallel across multiple hosts and does a sequential operation across the Oracle homes within a host. Within a given Oracle home, the procedure patches all of the Databases simultaneously.

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre-filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch the Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the Oracle homes are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

9. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

    - If you want to check for prerequisites, then click **Analyze**.

    - If you have already checked for prerequisites or if you do not want to check for prerequisites, then click **Deploy** to patch the selected databases.

23.5.2 Patching Oracle RAC Nodes Collectively

The prerequisites and the deployment instructions for patching all the nodes of Oracle RAC at a time, in parallel, using the *All Nodes* methodology, are the same as the ones described in Section 23.5.1, "Patching Oracle RAC Nodes Individually". However, in Step (3), on the Deployment Procedure Manager page, in the Procedures tab, from the table, select **Patch Oracle RAC Database - All Nodes**. Then click **Schedule Deployment**.

Although the instructions remain the same, when you submit the Deployment Procedure, Enterprise Manager Grid Control internally patches all the nodes of Oracle RAC collectively, that is, all the nodes are shut down and the patch is applied on all of them at the same time.

This methodology is best suited when you are applying a patch set or those one-off patches that support this methodology, or when you want to patch a shared Oracle RAC home. For example, if you are patching a shared Oracle home clusterware that has five nodes, the patch and the scripts within it are run only once on the shared Oracle home, and the other nodes are ignored or skipped. This saves time as the patching operation is performed only once, that is, only on the shared Oracle home, not on all the nodes.

> **Note:**
>
> ■ Patching of shared Oracle homes is supported only for Oracle Database 10g Release 2 (10.2.0.1) or higher.
>
> ■ If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

> **Note:** To view an online demonstration of this feature, access the following URL:
>
> http://www.oracle.com/technology/obe/demos/admin/demos.html

23.6 Patching Oracle RAC in Offline Mode

This section describes how you can patch Oracle RAC in offline mode, that is, without connecting to My Oracle Support. This option is best suited when you do not have an Internet connection on the host where Enterprise Manager Grid Control is running.

In offline mode, you can either patch the nodes one by one, or patch all the nodes at a time, in parallel. This section provides instructions for both these methodologies.

In particular, this section provides instructions for the following:

■ Patching Oracle RAC Nodes Individually

■ Patching Oracle RAC Nodes Collectively

23.6.1 Patching Oracle RAC Nodes Individually

This section describes how you can patch the nodes of Oracle RAC one by one, using the *Rolling* methodology.

This methodology is best suited when you are applying one-off patches that support this methodology, and when you want to maintain high availability of your targets, so when one node is being patched, the other nodes are available for service. For example, if you are patching a clusterware that has five nodes, then the first node is shut down, patched, and restarted, and then the process is rolled over to the next node, and so on, until all the nodes are patched successfully.

■ Prerequisites

■ Patching Procedure

23.6.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

■ Ensure that you meet the prerequisites described in the following sections:

   – Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

   – Section 3.2, "Optional Infrastructure Requirements"

   – Section 3.4, "Setting Up Infrastructure for Offline Patching"

- Ensure that operating system users such as *oracle* and *crsuser* are available on all the nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.

- Ensure that you run the patchability report to understand whether the Oracle Database you want to patch is suitable for this patching operation. In the report, if the Oracle Database you want to patch appears as a target that cannot be patched, then fix the target configuration issues to make it a target that can be patched. For more information about the patchability report, and to understand how you can diagnose the issues with an Oracle Database target and take corrective actions, see Appendix H, "About Patching Automation Reports".

- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

23.6.1.2 Patching Procedure

To patch an Oracle RAC database, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle RAC Database - Rolling**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

   a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

      This default value refers to a location on the target host where the Oracle RAC to be patched is running. The software updates are downloaded from My

Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched.

b. In the RAC Database Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search Software Library**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control searches the patches in Oracle Software Library, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. *(Optional)* In the Apply SQL Script section, retain the default selection so that the default SQL script that is packaged with the patch is used for modifying the schemas.

If you want to run any custom script, then select **Enter the script to apply SQL** and specify the full path to the location where the custom script is available. For example, `%oracle_home%/rdbms/admin/catupgrd.sql`.

Enterprise Manager Grid Control automatically runs certain scripts such as Catcpu.sql, utlrp.sql for CPUs; Catpsu.sql, utlrp.sql for PSUs; and catupgrade.slq for Patch Sets. However, some scripts mentioned in the ReadMe of a patch might have to be run manually. To run those scripts, you

can select **Enter the script to apply SQL** and specify the full path to the location where the custom scripts are available.

The variables you can use for the directory location are %emd_root% to indicate a location on the target host, `%oracle_home%` to indicate the current home location being patched, or `%perlbin%` to indicate the location of perl binary used by the Management Agent.

You have an option of running both the scripts by selecting both these options or one of the scripts by selecting only one of these options. If you select both these options, then both the scripts are run, that is, the default script runs first and then the custom script runs.

d. *(Optional)* In the Black Out Associated Targets section, retain the default selection so that all targets associated with the Oracle RAC database to be patched are blacked out while the patching is in progress.

e. *(Optional)* In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

f. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you identified any patch conflicts, then you can pass options to rectify the patch conflicts.

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

g. Click **Next**.

5. On the Target List page, in the RAC Database Targets to be Patched section, click **Add** to add the Oracle RAC database targets you want to patch. If you have multiple Oracle RAC databases in a cluster and if you are not sure of the selection, then search for all RAC databases in that cluster by selecting **Cluster** from the Target Type list. Then click **Next**.

> **Note:** Selection of multiple targets to be patched in a single operation is supported. Typically, this procedure orchestrates patching in parallel across multiple hosts and does a sequential operation across the Oracle homes within a host. Within a given Oracle home, the procedure patches all of the Databases simultaneously.

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre- filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch the Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the Oracle homes are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home
> Credentials section appears. When there are host-level active steps,
> the Host Credentials section appears. When Oracle home active steps
> and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials
stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of
credentials. If you choose to do so, you can specify either a common set of
credentials to be used across Oracle homes and hosts, or a unique set of credentials
for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click
Save OH Credential if you want to store the new credentials in the Management
Repository.

9. On the Schedule page, schedule the Deployment Procedure to run either
   immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure
   instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

    - If you want to check for prerequisites, then click **Analyze**.

    - If you have already checked for prerequisites or if you do not want to check
      for prerequisites, then click **Deploy** to patch the selected databases.

23.6.2 Patching Oracle RAC Nodes Collectively

The prerequisites and the deployment instructions for patching all the nodes of Oracle
RAC at a time, in parallel, using the *All Nodes* methodology, are the same as the ones
described in Section 23.6.1, "Patching Oracle RAC Nodes Individually". However, in
Step (3), on the Deployment Procedure Manager page, in the Procedures tab, from the
table, select **Patch Oracle RAC Database - All Nodes**. Then click **Schedule
Deployment**.

Although the instructions remain the same, when you submit the Deployment
Procedure, Enterprise Manager Grid Control internally patches all the nodes of Oracle
RAC collectively, that is, all the nodes are shut down and the patch is applied on all of
them at the same time.

This methodology is best suited when you are applying a patch set or those one-off
patches that support this methodology, or when you want to patch a shared Oracle
RAC home. For example, if you are patching a shared Oracle home clusterware that
has five nodes, the patch and the scripts within it are run only once on the shared
Oracle home, and the other nodes are ignored or skipped. This saves time as the
patching operation is performed only once, that is, only on the shared Oracle home,
not on all the nodes.

Note:

- Patching of shared Oracle homes is supported only for Oracle Database 10g Release 2 (10.2.0.1) or higher.

- If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

24

Patching Oracle Clusterware

This chapter explains how you can patch Oracle Clusterware. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedures

- Supported Releases

- Patching Oracle Clusterware in Online Mode

- Patching Oracle Clusterware in Offline Mode

24.1 Getting Started

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

Table 24–1   Getting Started with Patching Oracle Clusterware

Step	Description	
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Oracle Clusterware. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 24.2, "Deployment Procedures".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle Clusterware can be patched by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 24.3, "Supported Releases".

Table 24–1   (Cont.) Getting Started with Patching Oracle Clusterware

Step	Description	
Step 3	**Selecting the Patching Mode** Identify whether the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support. If it has, then choose to patch in online mode. Otherwise, choose to patch in offline mode.	▪ To learn about patching Oracle Clusterware in online mode, see Section 24.4, "Patching Oracle Clusterware in Online Mode". ▪ To learn about patching Oracle Clusterware in offline mode, see Section 24.5, "Patching Oracle Clusterware in Offline Mode".
Step 4	**Patching Clusterware Nodes Individually or Collectively** Decide whether you want to patch the nodes one by one, using the *Rolling* methodology, or all at a time, in parallel, using the *All Nodes* methodology.	**Online Mode:** ▪ To learn about patching clusterware nodes individually in online mode, see Section 24.4.1, "Patching Oracle Clusterware Nodes Individually". ▪ To learn about patching clusterware nodes collectively in online mode, see Section 24.4.2, "Patching Oracle Clusterware Nodes Collectively". **Offline Mode:** ▪ To learn about patching clusterware nodes individually in offline mode, see Section 24.5.1, "Patching Oracle Clusterware Nodes Individually". ▪ To learn about patching clusterware nodes collectively in offline mode, see Section 24.5.2, "Patching Oracle Clusterware Nodes Collectively".
Step 5	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the environment for online or offline patching, applying mandatory patches, setting up of Oracle Software Library, and so on.	**Online Mode:** ▪ To learn about the prerequisites for patching clusterware nodes individually in online mode, see Section 24.4.1.1, "Prerequisites". ▪ To learn about the prerequisites for patching clusterware nodes collectively in online mode, see Section 24.4.2, "Patching Oracle Clusterware Nodes Collectively". **Offline Mode:** ▪ To learn about the prerequisites for patching clusterware nodes individually in offline mode, see Section 24.5.1.1, "Prerequisites". ▪ To learn about the prerequisites for patching clusterware nodes collectively in offline mode, see Section 24.5.2, "Patching Oracle Clusterware Nodes Collectively".

Table 24–1   (Cont.) Getting Started with Patching Oracle Clusterware

Step	Description	
Step 6	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Oracle Clusterware.	**Online Mode:** ■ To patch clusterware nodes individually in online mode, see Section 24.4.1.2, "Patching Procedure". ■ To patch clusterware nodes collectively in online mode, see Section 24.4.2, "Patching Oracle Clusterware Nodes Collectively". **Offline Mode:** ■ To patch clusterware nodes individually in offline mode, follow the steps explained in Section 24.5.1.2, "Patching Procedure". ■ To patch clusterware nodes collectively in offline mode, follow the steps explained in Section 24.5.2, "Patching Oracle Clusterware Nodes Collectively".

24.2 Deployment Procedures

Enterprise Manager Grid Control offers the following Deployment Procedures for patching Oracle Clusterware:

■ *Patch Oracle Clusterware - Rolling*

■ *Patch Oracle Clusterware - All Nodes*

While the former one allows you to patch the nodes one by one, the latter one allows you to patch all the nodes at a time, in parallel. For more information about these patching methodologies and circumstances under which you must use the *Rolling* methodology or *All Nodes* methodology, see Section 2.4.3, "Understanding Patching Methodologies".

24.3 Supported Releases

Using these Deployment Procedures, you can patch the following releases of Oracle Clusterware:

■ Oracle Clusterware 10g Release 1 (10.1.x.x)

■ Oracle Clusterware 10g Release 2 (10.2.x.x)

■ Oracle Clusterware 11g Release 1 (11.1.x.x)

■ Oracle Clusterware 11g Release 2 (11.2.x.x)

> **IMPORTANT:**   These Deployment Procedures are not supported on Microsoft Windows.

24.4 Patching Oracle Clusterware in Online Mode

This section describes how you can patch Oracle Clusterware in online mode, that is, by connecting to My Oracle Support. This option is best suited when you have an Internet connection on the host where Enterprise Manager Grid Control is running.

In online mode, you can either patch the nodes one by one, or patch all the nodes at a time, in parallel. This section provides instructions for both these methodologies.

In particular, this section provides instructions for the following:

- Patching Oracle Clusterware Nodes Individually
- Patching Oracle Clusterware Nodes Collectively

24.4.1 Patching Oracle Clusterware Nodes Individually

This section describes how you can patch the nodes of Oracle Clusterware one by one, using the *Rolling* methodology.

This methodology is best suited when you are applying one-off patches and clusterware bundle patches that support this methodology, and when you want to maintain high availability of your targets, so when one node is being patched, the other nodes are available for service. For example, if you are patching a clusterware that has five nodes, then the first node is shut down, patched, and restarted, and then the process is rolled over to the next node, and so on, until all the nodes are patched successfully.

This section covers the following:

- Prerequisites
- Patching Procedure

> **iMPORTANT:**
>
> - You cannot use this methodology to patch shared Oracle homes. If you want to patch a shared Oracle home, then use Deployment Procedure that support the *All Nodes* methodology. For more information, see Section 24.4.2, "Patching Oracle Clusterware Nodes Collectively".
>
> - Although most one-off patches support this methodology, there are some one-off patches that can be applied only using the *All Nodes* methodology. You can identify the supported methodology while downloading the one-off patches from My Oracle Support.
>
> - If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

24.4.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
  - Section 3.3, "Setting Up Infrastructure for Online Patching"

- Ensure that you have the latest release of the OPatch software on the target host (where the clusterware to be patched is running).

- Ensure that operating system users such as *oracle* and *crsuser* are available on all the nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.

- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

24.4.1.2 Patching Procedure

To patch an Oracle Cluterware, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle Clusterware - Rolling**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

   a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

   This default value refers to a location on the target host where the clusterware to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

b. In the CRS Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search My Oracle Support**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control connects to My Oracle Support, searches the patches, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. *(Optional)* In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the database targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. *(Optional)* In the Black Out Associated Targets section, retain the default selection so that all targets associated with the clusterware to be patched are blacked out while the patching is in progress.

e. *(Optional)* In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

f. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you

identified any patch conflicts, then you can pass options to rectify the patch conflicts.

```
Advanced OPatch Options
Specify the advanced options to be passed to the Apply Patch Step.
Options  [                    ]
```

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

> **Note:** If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

    g. Click **Next**.

5. On the Target List page, in the Clusterware Targets to be Patched section, click **Add** to add the clusterware target you want to patch. Then click **Next**.

> **Note:** You can patch only one cluster at a time.

When you click **Add**, Enterprise Manager Grid Control displays the Select Target page. On this page, search and select the clusterware you want to patch, and click **Continue**.

```
Patch Oracle Clusterware - Rolling: Target List

Select Cluster or Host targets from the list below. RAC database targets having different version from the chosen patch will not be patched.

    Target Type   [Cluster          ▼]
    Target Version [11.1.0.7.0        ]
    Target Name   [                  ]
    On Host       [                  ]
    Platform      [Microsoft Windows (32-bit)  ▼]
                  [Go]

TIP You can use '%' and '*' as wildcard replacements in the text fields, for example, '%value' or '%value%' or '*value%' or similar patterns

Expand All | Collapse All
Select  Name                      Target Type   Select Hosts   Status
        ▼ Cluster Targets
  ⦿     ▼ crspaan(Unpatched)       Cluster                      ⇧
            stbcq18.us.oracle.com  Host          ☑              ⇧
            stbcq17.us.oracle.com  Host          ☑              ⇧
```

> **Note:** Do not deselect any of the hosts from the **Select Hosts** column of the Target List table.

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre- filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

   If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, do the following:

   ■ In the **Clusterware Credentials** section, specify the Oracle home credentials required to patch the clusterware Oracle homes.

   ■ In the RAC Credentials section, specify the Oracle home credentials required to patch the RAC databases associated withthe clusterware.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host. After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

If you are the owner of the clusteware and RAC Oracle homes or if you have the owner credentials, then for RAC Credentials section, you can choose to use the same clusterware credentials you specified in the Clusterware Credentials section. To do so, in the RAC Credentials section, select **Use Clusterware Credentials**.

RAC Credentials

○ Use Oracle Home Preferred Credentials	○ Override Oracle Home Preferred Credentials		◉ Use Clusterware Credentials	
Oracle Home	**Host**	**Target**	**Username**	**Password**
c:\oracle\11.1.0.7\db	stbcq17.us.oracle.com	zeta	oracle	******

9. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

    ■ If you want to check for prerequisites, then click **Analyze**.

    ■ If you have already checked for prerequisites or if you do not want to check for prerequisites, then click **Deploy** to patch the selected cluster.

24.4.2 Patching Oracle Clusterware Nodes Collectively

The prerequisites and the deployment instructions for patching all the nodes of Oracle Clusterware at a time, in parallel, using the *All Nodes* methodology, are the same as the ones described in Section 24.4.1, "Patching Oracle Clusterware Nodes Individually". However, in Step (3), on the Deployment Procedure Manager page, in the Procedures tab, from the table, select **Patch Oracle Clusterware - All Nodes**. Then click **Schedule Deployment**.

Although the instructions remain the same, when you submit the Deployment Procedure, Enterprise Manager Grid Control internally patches all the nodes of the clusterware collectively, that is, all the nodes are shut down and the patch is applied on all of them at the same time.

This methodology is best suited when you want to patch a shared Oracle home clusterware. For example, if you are patching a shared Oracle home clusterware that has five nodes, the patch and the scripts within it are run only once on the shared Oracle home, and the other nodes are ignored or skipped. This saves time as the patching operation is performed only once, that is, only on the shared Oracle home, not on all the nodes.

> **Note:**
>
> ■ Patching of shared Oracle homes is supported only for Oracle Clusterware 10g Release 2 (10.2.0.1) or higher.
>
> ■ If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

24.5 Patching Oracle Clusterware in Offline Mode

This section describes how you can patch Oracle Clusterware in offline mode, that is, without connecting to My Oracle Support. This option is best suited when you do not have an Internet connection on the host where Enterprise Manager Grid Control is running.

In offline mode, you can either patch the nodes one by one, or patch all the nodes at a time, in parallel. This section provides instructions for both these methodologies.

In particular, this section provides instructions for the following:

- Patching Oracle Clusterware Nodes Individually
- Patching Oracle Clusterware Nodes Collectively

24.5.1 Patching Oracle Clusterware Nodes Individually

This section describes how you can patch the nodes of Oracle Clusterware one by one, using the *Rolling* methodology.

This methodology is best suited when you are applying one-off patches and clusterware bundle patches that support this methodology, and when you want to maintain high availability of your targets, so when one node is being patched, the other nodes are available for service. For example, if you are patching a clusterware that has five nodes, then the first node is shut down, patched, and restarted, and then the process is rolled over to the next node, and so on, until all the nodes are patched successfully.

This section covers the following:

- Prerequisites
- Patching Procedure

24.5.1.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers
- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
  - Section 3.4, "Setting Up Infrastructure for Offline Patching"
- Ensure that you have the latest release of the OPatch software on the target host (where the clusterware to be patched is running).
- Ensure that operating system users such as *oracle* and *crsuser* are available on all the nodes of the cluster. These users must be a part of the relevant operating system groups such as *dba* and *oinstall*.
- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

- Ensure that you are the owner of the Oracle home of the targets being patched. If you are not the owner or if you do not have direct access to those targets, then customize the Deployment Procedure to run as another user.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

24.5.1.2 Patching Procedure

To patch an Oracle Cluterware, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Patching section, click **Patching through Deployment Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Oracle Clusterware - Rolling**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

   a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

   This default value refers to a location on the target host where the clusterware to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

   You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

   Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

   b. In the CRS Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

   On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search Software Library**, specify the details of the patches you want to search, and click **Go**.

   Enterprise Manager Grid Control searches the patches in Oracle Software Library, and displays details about them on the Search and Select Software

Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

Note: Selection of multiple patches is supported.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. *(Optional)* In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the database targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. *(Optional)* In the Black Out Associated Targets section, retain the default selection so that all targets associated with the clusterware to be patched are blacked out while the patching is in progress.

Black Out Associated Targets
☑ Black Out Associated Targets
Black out targets affected by patching. For example, black out listeners

e. *(Optional)* In the Clean Up Backup Files section, select the option if you want to clean up all patching backup files after the patches are applied. The backup files are cleaned up by running OPatch utility cleanup command. Select this option depending on the space available on the target Oracle home.

f. *(Optional)* In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. Typically, you can use this section when you have patch conflicts or when inventory location is not default for the target Oracle home being patched. For example, after running the Deployment Procedure in *Analyze Mode*, if you identified any patch conflicts, then you can pass options to rectify the patch conflicts.

Advanced OPatch Options
Specify the advanced options to be passed to the Apply Patch Step.
Options [   ]

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

> **Note:** If you want to specify multiple options, then separate the list with a white space. For example, -skip_duplicate -skip_conflict

g. Click **Next**.

5. On the Target List page, in the Clusterware Targets to be Patched section, click **Add** to add the clusterware targets you want to patch. Then click **Next**.

> **Note:** You can patch only one cluster at a time.

When you click **Add**, Enterprise Manager Grid Control displays the Select Target page. On this page, search and select the clusterware you want to patch, and click **Continue**.

> **Note:** Do not deselect any of the hosts from the **Select Hosts** column of the Target List table.

6. On the Library Step Properties page *(appears only if you had customized the Deployment Procedure and marked some steps to be prompted at runtime)*, specify values for the library steps.

7. *(Optional)* On the Oracle Configuration Manager page, specify an e-mail address to configure Oracle Configuration Manager. If you had already set up *My Oracle Support* credentials for the Enterprise Manager user account being used, then the fields are pre- filled with those details. If you have a proxy server, then provide details about the proxy server, and click **Next**.

If you do not want to configure Oracle Configuration Manager, then delete the details appearing on this screen, and click **Next**.

8. On the Credentials page, do the following:

   ■ In the **Clusterware Credentials** section, specify the Oracle home credentials required to patch the clusterware Oracle homes.

- In the RAC Credentials section, specify the Oracle home credentials required to patch the RAC databases associated withthe clusterware.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host. After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

If you are the owner of the clusteware and RAC Oracle homes or if you have the owner credentials, then for RAC Credentials section, you can choose to use the same clusterware credentials you specified in the Clusterware Credentials section. To do so, in the RAC Credentials section, select **Use Clusterware Credentials**.

9. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

10. On the Review page, review the details and do one of the following:

    - If you want to check for prerequisites, then click **Analyze**.

- If you have already checked for prerequisites or if you do not want to check for prerequisites, then click **Deploy** to patch the selected cluster.

24.5.2 Patching Oracle Clusterware Nodes Collectively

The prerequisites and the deployment instructions for patching all the nodes of Oracle Clusterware at a time, in parallel, using the *All Nodes* methodology, are the same as the ones described in Section 24.5.1, "Patching Oracle Clusterware Nodes Individually". However, in Step (3), on the Deployment Procedure Manager page, in the Procedures tab, from the table, select **Patch Oracle Clusterware - All Nodes**. Then click **Schedule Deployment**.

Although the instructions remain the same, when you submit the Deployment Procedure, Enterprise Manager Grid Control internally patches all the nodes of the clusterware collectively, that is, all the nodes are shut down and the patch is applied on all of them at the same time.

This methodology is best suited when you want to patch a shared Oracle home clusterware. For example, if you are patching a shared Oracle home clusterware that has five nodes, the patch and the scripts within it are run only once on the shared Oracle home, and the other nodes are ignored or skipped. This saves time as the patching operation is performed only once, that is, only on the shared Oracle home, not on all the nodes.

> **Note:** Patching of shared Oracle homes is supported only for Oracle Clusterware 10g Release 2 (10.2.0.1) or higher.

> **Note:** If the Deployment Procedure abruptly stops or terminates at any step while it is being run, you may rerun the same procedure from where it stopped or from the beginning, all over again.

25

Patching Oracle Automatic Storage Management (ASM)

This chapter explains how you can patch Oracle ASM. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedure
- Supported Releases
- Patching a Standalone Oracle ASM in Online Mode
- Patching a Standalone Oracle ASM in Offline Mode

25.1 Getting Started

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

Table 25–1   Getting Started with Patching Oracle ASM

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Oracle ASM. Know how the Deployment Procedure functions, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Section 25.2, "Deployment Procedure".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle ASM can be patched by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 25.3, "Supported Releases".

Table 25–1   (Cont.) Getting Started with Patching Oracle ASM

Step	Description	Reference Links
Step 3	**Selecting the Patching Mode** Identify whether the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support. If it has, then choose to patch in online mode. Otherwise, choose to patch in offline mode.	■ To learn about patching a standalone Oracle ASM in online mode, see Section 25.4, "Patching a Standalone Oracle ASM in Online Mode". ■ To learn about patching standalone Oracle ASM in offline mode, see Section 25.5, "Patching a Standalone Oracle ASM in Offline Mode".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the environment for online or offline patching, applying mandatory patches, setting up of Oracle Software Library, and so on.	■ To learn about the prerequisites for patching a standalone Oracle ASM in online mode, see Section 25.4.1, "Prerequisites". ■ To learn about the prerequisites for patching a standalone Oracle ASM in offline mode, see Section 25.5.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Oracle ASM.	■ To patch a standalone Oracle ASM in online mode, follow the steps explained in Section 25.4.2, "Patching Procedure". ■ To patch a standalone Oracle ASM in offline mode, follow the steps explained in Section 25.5.2, "Patching Procedure".

25.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedures for patching Oracle ASM:

- *Patch Standalone Oracle ASM*

25.3 Supported Releases

Using this Deployment Procedure, you can patch the following releases of Oracle ASM:

- Oracle Database 10g Release 1 (10.1.x.x)

- Oracle Database 10g Release 2 (10.2.x.x)

- Oracle Database 11g Release 1 (11.1.x.x)

25.4 Patching a Standalone Oracle ASM in Online Mode

This section describes how you can patch a standalone Oracle ASM in online mode, that is, by connecting to My Oracle Support. This option is best suited when you have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites

- Patching Procedure

25.4.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

    - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

    - Section 3.2, "Optional Infrastructure Requirements"

    - Section 3.3, "Setting Up Infrastructure for Online Patching"

- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

    If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

    For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

    For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

25.4.2 Patching Procedure

To patch a standalone Oracle ASM instance, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Patch Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Standalone Oracle ASM**. Then click **Schedule Deployment**.

    Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

    a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

    This default value refers to a location on the target host where the ASM to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

b. In the ASM Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search My Oracle Support**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control connects to My Oracle Support, searches the patches, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the ASM targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. In the Black Out Associated Targets section, retain the default selection so that all targets associated with the Oracle home of the ASM target to be patched are blacked out while the patching is in progress.

e. In the OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. For example, if you want the Deployment Procedure to ignore duplicate patches, then specify `-skip_duplicate`.

Advanced OPatch Options

Specify the advanced options to be passed to the Apply Patch Step.

Options [   ]

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

f. Click **Next**.

5. On the Target List page, in the Clusterware Targets to be Patched section, click **Add** to add the database targets that contain the ASM targets you want to patch. Then click **Next**

6. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch ASM Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the ASM targets are present. Then click **Next**.

Home Credentials

⦿ Use Oracle Home Preferred Credentials   ○ Override Oracle Home Preferred Credentials

Oracle Home	Host	Target	Username	Password
/scratch/akash/rac	stamy04.us.oracle.com	horcl_horcl1		

Note: When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

7. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

8. On the Review page, review the details you have provided for patching the selected databases that contain the ASM instances, and click **Analyze**. After running the prerequisite checks, click **Deploy** to patch the selected databases that contain the ASM instances.

25.5 Patching a Standalone Oracle ASM in Offline Mode

This section describes how you can patch a standalone Oracle ASM in offline mode, that is, without connecting to My Oracle Support. This option is best suited when you do not have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites
- Patching Procedure

25.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
  - Section 3.4, "Setting Up Infrastructure for Offline Patching"
- Ensure that you run the Deployment Procedure in the Analyze mode.

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".
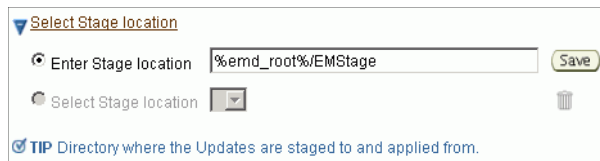
25.5.2 Patching Procedure

To patch a standalone Oracle ASM instance, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Patch Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Standalone Oracle ASM**. Then click **Schedule Deployment**.

    Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

    a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

        This default value refers to a location on the target host where the ASM to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

        You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

        Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

    b. In the ASM Database Updates section, click **Add** to search and select software updates from Oracle Software Library (Software Library). Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

        On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search Software Library**, specify the details of the patches you want to search, and click **Go**.

        Enterprise Manager Grid Control searches the patches in the Software Library, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

        On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

    c. In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target hosts are upgraded before the patches are applied on the ASM targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. In the Black Out Associated Targets section, retain the default selection so that all targets associated with the Oracle home of the ASM target to be patched are blacked out while the patching is in progress.

e. In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. For example, if you want the Deployment Procedure to ignore duplicate patches, then specify `-skip_duplicate`.

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

f. Click **Next**.

5. On the Target List page, in the Clusterware Targets to be Patched section, click **Add** to add the database targets that contain the ASM targets you want to patch. Then click **Next**

6. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch ASM Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the ASM targets are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home
> Credentials section appears. When there are host-level active steps,
> the Host Credentials section appears. When Oracle home active steps
> and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials
stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of
credentials. If you choose to do so, you can specify either a common set of
credentials to be used across Oracle homes and hosts, or a unique set of credentials
for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click
Save OH Credential if you want to store the new credentials in the Management
Repository.

7. On the Schedule page, schedule the Deployment Procedure to run either
   immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure
   instance so that it can be tracked later and reused with the same settings.

8. On the Review page, review the details you have provided for patching the
   selected databases that contain the ASM instances, and click **Analyze**. After
   running the prerequisite checks, click **Deploy** to patch the selected databases that
   contain the ASM instances.

26

Patching Primary and Standby Databases Configured with Oracle Data Guard

This chapter describes the following:

- Overview of Patching Primary and Standby Databases
- Patching Primary and Standby Databases

26.1 Overview of Patching Primary and Standby Databases

To patch the primary and standby databases that are configured with Oracle Data Guard, you can use the default database patching deployment procedures described in Part VI. The databases that act as primary or standby can be either a standalone database (single-instance) or an Oracle Real Application Cluster (Oracle RAC) database running with or without a broker.

However, the following are the limitations with patching of primary and standby databases:

- There is no support for orchestrated patching between a primary database and its standby database. Although the default database patching deployment procedures automatically identify the primary and the standby databases, they do not recognize the association between them. As a result, you must manually identify the primary database and its associated standby database, and patch them separately in such a way that the standby database switches over to production mode while the primary database is being patched.

- There is no support automatically stopping and starting log shipping. You must manually stop and start the log shipping either directly or via a broker. The deployment procedures do not handle this by default. However, if you want, you can customize the deployment procedure to include additional steps that will handle this issue.

- There is no support for Oracle Data Guard in logical standby configuration.

26.2 Patching Primary and Standby Databases

To patch the primary and the standby database, follow these steps:

1. Customize the database patching deployment procedure as described in Section 26.2.1 to include additional steps to handle stopping and starting of log shipping with or without a broker.

2. Identify the primary and standby database.

3. Patch the standby database using the deployment procedure.

4. Patch the primary database using the deployment procedure.

26.2.1 Customizing Patching Deployment Procedure

To automate the stopping and starting of log shipping, you must customize the patching deployment procedure to include two steps—one to stop the log shipping on the primary database and another to start the log shipping once the patching operation ends.

To do so, follow these steps:

1. On the Deployment Procedure Manager page, select one of the following patching deployment procedure based on the target type you want to patch, and click **Create Like.**

- *Patch Oracle Database*, for Oracle Database

- *Patch Oracle RAC Rolling*, for the Oracle RAC Databases

2. On the Create Like page, select the step **Stop Database in Normal Mode,** and insert a new step of the type *Host Command.* Ensure that you insert the new step before the step **Shutdown Database.**

3. On the Create Host Command Step, in the Script textbox, copy the following lines:

```
# Script for Stopping Log Shipping on Primary Database, Insert before Shutdown
Database step
echo ' Executing Stop Log Shipping command on Primary Database'
PRIMDB_LIST=${target.primaryDBSIDs}
DGMGRL_CMD=${target.oraHome}/bin/dgmgrl
STATE='LOG-TRANSPORT-OFF'
if [ -z $PRIMDB_LIST ]
then
 echo " Selected Databases are not Primary Databases found.  Step will be
skipped"
 exit 0
fi
PRIMDB_SPCSEP_LIST=`echo $PRIMDB_LIST | tr ',' ' '`
for EACH_PRIMDB in $PRIMDB_SPCSEP_LIST
do
        echo Running $DGMGRL_CMD on $EACH_PRIMDB
        export ORACLE_HOME=${target.oraHome}
        export ORACLE_SID=$EACH_PRIMDB
        $DGMGRL_CMD -silent / "edit database $EACH_PRIMDB SET STATE=$STATE"
$DGMGRL_CMD  -silent / "show database $EACH_PRIMDB"
done

exit 0
```

Figure 26–1 shows how the Script textbox will look after you copy the preceeding lines.

Figure 26–1   Creating a New Host Command for Starting and Stopping the Log Shipping

4. On the Create Like page, select the step **Apply Post SQL Script,** and insert a new step of the type *Host Command.* Ensure that you insert the step after the step **Apply post SQL script.**

5. On the Create Host Command Step, in the Script textbox, copy the following lines:

```
# Script for Starting Log Shipping on Primary Database, Insert after Apply Post
SQL step
echo ' Executing Start Log Shipping command on Primary Database'
PRIMDB_LIST=${target.primaryDBSIDs}
DGMGRL_CMD=${target.oraHome}/bin/dgmgrl
STATE=' ONLINE '
if [ -z $PRIMDB_LIST ]
then
  echo " Selected Databases are not Primary Databases found.  Step will be
skipped "
  exit 0
fi
PRIMDB_SPCSEP_LIST=`echo $PRIMDB_LIST | tr ',' ' '`
for EACH_PRIMDB in $PRIMDB_SPCSEP_LIST
do
        echo Running $DGMGRL_CMD on $EACH_PRIMDB
        export ORACLE_HOME=${target.oraHome}
        export ORACLE_SID=$EACH_PRIMDB
        $DGMGRL_CMD -silent / "edit database $EACH_PRIMDB SET STATE=$STATE"
        $DGMGRL_CMD -silent / "show database $EACH_PRIMDB"
done
exit 0
```

6. Save the customized deployment procedure with a unique name. For example, *Patch Oracle Database in Physical Standby Configuration.* Use this procedure for patching the primary and the standby databases.

> **Note:** You can customize the deployment procedure even further if you want. For more information on customizing deployment procedures, refer to Chapter 31.

27

Patching Oracle Application Server

This chapter explains how you can patch Oracle Application Server. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedure

- Core Components Patched

- Supported Releases

- Patching Oracle Application Server in Online Mode

- Patching Oracle Application Server in Offline Mode

27.1 Getting Started

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

Table 27–1   Getting Started with Patching Oracle Application Server

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Oracle Application Server. Know how the Deployment Procedure functions, what use cases it covers, what core components are patched, and so on.	To learn about the Deployment Procedure, see Section 27.2, "Deployment Procedure". To learn about the core components that are patched, see Section 27.3, "Core Components Patched".
Step 2	**Knowing About the Supported** Releases Know what releases of Oracle Application Server can be patched by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 27.4, "Supported Releases".

Table 27–1   (Cont.) Getting Started with Patching Oracle Application Server

Step	Description	Reference Links
Step 3	**Selecting the Patching Mode** Identify whether the host that runs Enterprise Manager Grid Control has an Internet connection to connect to My Oracle Support. If it has, then choose to patch in online mode. Otherwise, choose to patch in offline mode.	■ To learn about patching Oracle Application Server in online mode, see Section 27.5, "Patching Oracle Application Server in Online Mode". ■ To learn about patching Oracle Application Server in offline mode, see Section 27.6, "Patching Oracle Application Server in Offline Mode".
Step 4	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the environment for online or offline patching, applying mandatory patches, setting up of Oracle Software Library, and so on.	■ To learn about the prerequisites for patching Oracle Application Server in online mode, see Section 27.5.1, "Prerequisites". ■ To learn about the prerequisites for patching Oracle Application Server in offline mode, see Section 27.6.1, "Prerequisites".
Step 5	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Oracle Application Server.	■ To patch Oracle Application Server in online mode, follow the steps explained in Section 27.5.2, "Patching Procedure". ■ To patch Oracle Application Server in offline mode, follow the steps explained in Section 27.6.2, "Patching Procedure".

27.2 Deployment Procedure

Enterprise Manager Grid Control offers the following Deployment Procedure for patching Oracle Application Server:

■ *Patch Application Server*

27.3 Core Components Patched

When you patch Oracle Application Server, essentially, the Deployment Procedure patches the following core components:

■ Oracle HTTP Server

■ Oracle Web Cache

■ Oracle Containers for J2EE (OC4J)

■ Oracle Portal

27.4 Supported Releases

Using this Deployment Procedure, you can patch the following releases of Oracle Application Server:

■ Oracle Application Server 10g Release 2 (10.1.2.0.2)

■ Oracle Application Server 10g Release 3 (10.1.3.x)

27.5 Patching Oracle Application Server in Online Mode

This section describes how you can patch an Oracle Application Server in online mode, that is, by connecting to My Oracle Support. This option is best suited when you have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites
- Patching Procedure

27.5.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
  - Section 3.3, "Setting Up Infrastructure for Online Patching"

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

27.5.2 Patching Procedure

To patch an Oracle Application Server, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Patch Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Application Server**. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

This default value refers to a location on the target host where the application server to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library (Software Library), and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

b. In the Application Server Updates section, click **Add** to search and select software updates from My Oracle Support. Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search My Oracle Support**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control connects to My Oracle Support, searches the patches, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target host is upgraded before the patches are applied on the application server targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. In the Black Out Associated Targets section, retain the default selection so that all targets associated with the application server to be patched are blacked out while the patching is in progress.

    e. In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. For example, if you want the Deployment Procedure to ignore duplicate patches, then specify `-skip_duplicate`.

    The following are the options you can specify:

    `-skip_duplicate`, skips duplicate patches

    `-skip_conflict`, skips conflicting patches

    `-skip_missing_component`, skips patches whose components are missing

    `-invPtrLoc`, the location of oraInst.loc file

    Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

    f. Click **Next**.

5. On the Target List page, in the Application Server Targets to be Patched section, click **Add** to add the application server targets you want to patch. Then click **Next**.

6. On the Oracle Configuration Manager page, do the following:

    a. If the host where the application server is being patched has a *direct* connection to the Internet, then specify an e-mail address and My Oracle Support password.

    An e-mail address is required so that security updates and install updates can be sent. You can specify any e-mail address, but Oracle recommends you to specify the My Oracle Support user name. For example, `john.mathew@xyz.com`.

    If the My Oracle Support password is incorrect, you will be allowed two more attempts. However, if your password is incorrect in all three attempts or if it is left blank, then you are registered anonymously, which means, the configuration information will be collected and uploaded to My Oracle Support but the uploaded information will not be associated with your My Oracle Support account. Therefore, if you log in to My Oracle Support with your credentials, you will not see this information displayed against your account. However, if you had specified an e-mail address, then you will continue to receive security updates and other notifications from Oracle to that e-mail address.

    b. If the host where the application server is being patched has an *indirect* connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

> **Note:** You can change the proxy server settings any time after the Deployment Procedure ends. To do so, run the `configCCR` command from the `/ccr/bin/` directory within the Oracle home directory of the provisioned database.

    c. If the host where the application server is being patched does not have a *direct* or *indirect* connection to the Internet, then specify the e-mail address and leave the other fields blank.

    In this case, after you complete the installation process, manually collect the configuration information and upload it to My Oracle Support. To understand how the configuration information can be manually collected and uploaded, see the steps outlined in Section 5.6.2.1, "Manually Collecting and Uploading Configuration Information to My Oracle Support".

7. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch the application server Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the application servers are present. Then click **Next**.

> **Note:** When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

8. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

9. On the Review page, review the details you have provided for patching the selected application servers, and click **Submit**.

27.6 Patching Oracle Application Server in Offline Mode

This section describes how you can patch an Oracle Application Server in offline mode, that is, without connecting to My Oracle Support. This option is best suited when you do not have an Internet connection on the host where Enterprise Manager Grid Control is running.

This section covers the following:

- Prerequisites
- Patching Procedure

27.6.1 Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

  - Section 3.4, "Setting Up Infrastructure for Offline Patching"

Prerequisites for Operators

- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure, and that can switch to *root* user and run all commands on the target hosts. For example, commands such as mkdir, ls, and so on.

  If you do not have the privileges to do so, that is, if you are using a locked account, then request your administrator (a designer) to either customize the Deployment Procedure to run it as another user or ignore the steps that require special privileges.

  For example, user account A might have the root privileges, but you might use user account B to run the Deployment Procedure. In this case, you can switch from user account B to A by customizing the Deployment Procedure.

  For information about customization, see Chapter 31.7, "Customizing Steps to Run as Another User".

27.6.2 Patching Procedure

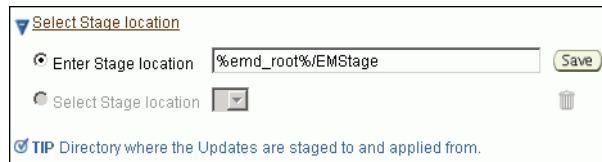To patch Oracle Application Server, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, in the Deployment Procedure Manager section, click **Patch Procedures**.

3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select **Patch Application Server**. Then click **Schedule Deployment**.

    Enterprise Manager Grid Control displays the Software Updates page of the Deployment Procedure.

4. On the Software Updates page, do the following:

a. In the Select Stage Location section, retain the default selection, that is, **Enter Stage Location** that has the value `%emd_root%/EMStage`.

This default value refers to a location on the target host where the database to be patched is running. The software updates are downloaded from My Oracle Support to Oracle Software Library, and then staged in this location on the target hosts.

You can modify the location if you want to, and click **Save**. If you save multiple custom locations, then the Select Stage Location option gets enabled and allows you to select any of the saved custom locations the next time you visit the same page.

Instead of %emd_root%, you can also specify `%oracle_home%` to indicate the current home location being patched or `%perlbin%` to indicate the location of perl binary used by Oracle Management Agent (Management Agent).

b. In the Application Server Updates section, click **Add** to search and select software updates from Oracle Software Library (Software Library). Alternatively, if you have a text file that contains details about the patches, then click **Upload From File** to select and upload that text file.

On clicking **Add**, Enterprise Manager Grid Control displays the Search and Select Software Updates page. On this page, select **Search Software Library**, specify the details of the patches you want to search, and click **Go**.

Enterprise Manager Grid Control searches the patches in the Software Library, and displays details about them on the Search and Select Software Updates page. Select the ones you want to apply using the Deployment Procedure, and click **Select**.

On clicking **Upload From File**, Enterprise Manager Grid Control displays the Upload File page. On this page, select **Upload from Local Machine** or **Upload from Software Library** depending on where the text file is available, click **Browse** to search and select that file, and then click **Upload**.

c. In the Upgrade OPatch section, retain the default selection so that the OPatch software on the target host is upgraded before the patches are applied on the application server targets.

Enterprise Manager Grid Control uses OPatch software to apply the one-off patches. It is assumed that this software is already available on the target hosts managed by Oracle Management Agents (Management Agent), but Oracle recommends you to retain the selection so that the existing software is upgraded to the latest release.

d. In the Black Out Associated Targets section, retain the default selection so that all targets associated with the application server to be patched are blacked out while the patching is in progress.

e. In the Advanced OPatch Options section, specify any opatch-related options you want to pass while running this Deployment Procedure. For example, if you want the Deployment Procedure to ignore duplicate patches, then specify `-skip_duplicate`.

The following are the options you can specify:

`-skip_duplicate`, skips duplicate patches

`-skip_conflict`, skips conflicting patches

`-skip_missing_component`, skips patches whose components are missing

`-invPtrLoc`, the location of oraInst.loc file

Note: If you want to specify multiple options, then separate the list with a white space. For example, `-skip_duplicate -skip_conflict`

---

f. Click **Next**.

5. On the Target List page, in the Application Server Targets to be Patched section, click **Add** to add the application server targets you want to patch. Then click **Next**.

6. On the Oracle Configuration Manager page, click **Next**.

7. On the Credentials page, in the **Home Credentials** section, specify the Oracle home credentials required to patch the application server Oracle homes, and in the Host Credentials section, specify the operating system credentials to log in to the hosts where the application servers are present. Then click **Next**.

Note: When there are Oracle home active steps, the Home Credentials section appears. When there are host-level active steps, the Host Credentials section appears. When Oracle home active steps and host-level active steps are present, both the sections appear.

For both sections, choose to use the preferred credentials so that the credentials stored in the Management Repository can be used.

Of course, you can always override the preferred credentials with a new set of credentials. If you choose to do so, you can specify either a common set of credentials to be used across Oracle homes and hosts, or a unique set of credentials for each Oracle home and host.

After specifying a new set of credentials to override the preferred credentials, click **Save OH Credential** if you want to store the new credentials in the Management Repository.

8. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.

   For **Instance Name**, specify a unique name for this Deployment Procedure instance so that it can be tracked later and reused with the same settings.

9. On the Review page, review the details you have provided for patching the selected application servers, and click **Submit**.

28

Patching Linux Hosts

Linux Host Patching is a feature in Enterprise Manager Grid Control that helps in keeping the machines in an enterprise updated with security fixes and critical bug fixes, especially in a data centre or a server farm.

This feature support in Enterprise Manager Grid Control enables you to:

- Set up Linux RPM Repository based in Unbreakable Linux Network (ULN) channels.

- Download Advisories (Erratas) from ULN.

- Set up 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 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)

This chapter explains how you can patch Linux hosts. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedures
- Supported Releases
- Patching Linux Hosts
- Linux Patching Home Page
- Configuration File Management

28.1 Getting Started

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

Table 28–1 Getting Started with Patching Linux Hosts

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Linux hosts.	To learn about the Deployment Procedure, see Section 28.2, "Deployment Procedures".
Step 2	**Knowing About the Supported** Releases Know what releases of Linux are supported by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 28.3, "Supported Releases".
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the patching environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for patching Linux hosts, see Section 28.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Linux hosts.	To patch Linux hosts, follow the steps explained in Section 28.4.2, "Patching Procedure".
Step 5	**Performing Other Linux Patching Activities** There are other activities that you can perform such as registering custom channels, configuration file management, and so on.	■ To register a custom channel, follow the steps explained in Section 28.5.4, "Registering a Custom Channel". ■ To clone a channel, follow the steps explained in Section 28.5.5, "Cloning a Channel". ■ To perform configuration file management, follow the steps explained in Section 28.6, "Configuration File Management".

28.2 Deployment Procedures

Enterprise Manager Grid Control provides the following Deployment Procedures for Linux patching:

■ *Patch Linux Hosts*

This deployment procedure allows you to patch Linux hosts. The steps to run this deployment procedure is specified in the next section.

■ *Linux RPM Repository server setup*

This deployment procedure allows you to set up a Linux RPM repository server. Follow the steps mentioned in Section 3.5.2, "Setting Up the RPM Repository" to set up the Linux RPM repository server.

28.3 Supported Releases

Following are the releases supported for Linux patching:

Table 28–2   Supported Releases

Feature	Linux Distributions Supported
Compliance	Oracle Enterprise Manager Linux, RedHat
Update Job	Oracle Enterprise Manager Linux, RedHat, SuSE Linux
Emergency Patching	Oracle Enterprise Manager Linux, RedHat
Linux Patching Deployment Procedures	Oracle Enterprise Manager Linux, RedHat, SuSE Linux
Undo Patching	Oracle Enterprise Manager Linux, RedHat
Channel management	Oracle Enterprise Manager Linux, RedHat

28.4 Patching Linux Hosts

This section explains how you can patch Linux hosts. This section covers the following:

- Prerequisites
- Patching Procedure

28.4.1 Prerequisites

Before running the deployment procedure, ensure that you meet the following requirements:

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Software Library must be configured in OMS.
- All the target machines must have yum or up2date installed in it, and must also have sudo enabled for the patch user.

28.4.2 Patching Procedure

Follow these steps to patch Linux hosts using deployment procedures:

1. In the **Deployments** tab, click **Patching through Deployment Procedures**. Select **Patch Linux Hosts** and click **Schedule Deployment**.

2. In the Linux Distribution section, select the correct distribution and also select the update tool to use.

3. In the Package Repository section, click the torch icon to select the RPM Repository.

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

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

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

7. In the Pre/Post script page, enter the pre/post scripts, if any.

8. In the Schedule page, enter the schedule to be used.

9. Review the update parameters and click **Finish**.

10. A deployment procedure will be submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

11. Once the procedure is completed, you can view the updated versions of the packages by clicking Targets, selecting the target name, clicking Configuration, and then clicking Packages.

28.5 Linux Patching Home Page

In the Linux Patching Home page, you can do the following:

■ Manage and monitor out-of-date and rogue packages for each host in the enterprise

■ Configure targets that need to be automatically updated

■ View compliance reports

28.5.1 Viewing Compliance History

The Compliance History page provides compliance details for the selected Group, for a specific time period. You can select the required time period from the View Data drop-down list and click **Refresh**.

28.5.1.1 Prerequisites

- At least one Linux patching group must be defined.

- You must have View privileges on the Linux host comprising the patching group.

28.5.1.2 Procedure to View Compliance History

To view the compliance history of a Linux patching group:

1. In the **Deployments** tab, select **Linux Patching**.

2. In the Linux Patching page, under Related Links, click **Compliance History**.

3. In the Compliance History page, the Groups table in this page 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 will display 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.

28.5.2 Patching Non-Compliant Packages

Patch non-compliant packages from the Linux Patching home page.

28.5.2.1 Prerequisites

Before patching non-compliant packages, ensure that:

- A Linux Patching group is created and the Compliance Collection job has succeeded.

28.5.2.2 Procedure to patch Non-Compliant Packages

Do the following to patch non-compliant packages:

1. Log in to Enterprise Manager Grid Control.

2. In the **Deployments** tab, select **Linux Patching**.

3. In the Linux Patching page, in the Compliance Reporting section, select the Group and click **Schedule Patching**.

4. Control is transferred to the Credentials page of the "Patch Linux Hosts" deployment procedure. Click **Back** to change the list of packages to update or the targets to update.

5. Enter the credentials and click **Next**.

6. Enter the pre/post scripts, if any.

7. Enter the schedule.

8. Review the updated parameters and click **Finish**.

9. A deployment procedure is submitted to update the host. Check if all the steps finished successfully.

28.5.3 Undo Patching

You can uninstall packages from the Linux Patching home page.

28.5.3.1 Prerequisites

Before rolling back packages, ensure that:

- A Linux Patching group is created.

- The lower version of the package must be present in the RPM repository.

28.5.3.2 Procedure to Undo Patching

Do the following to uninstall or rollback packages installed:

1. Log in to Enterprise Manager Grid Control.

2. In the **Deployments** tab, select **Linux Patching**.

3. In the Linux Patching page, from the Compliance report table, click **Undo Patching**.

4. In the Undo Patching: Action page, select **Uninstall Packages**.

5. Select the Rollback Last Update Session option and click **Next**.

6. Select the Group and click **Next**.

7. In the Undo Patching: Credentials page, specify the credentials to be used while performing the undo patching job.

8. In the Undo Patching: Patching Scripts page, select the pre/post patching scripts if any.

9. Review the selected options and click **Finish**.

10. A job is submitted to rollback the updates done in the previous session.

11. Examine the job submitted to see if all the steps are successful.

28.5.4 Registering a Custom Channel

You can register a custom channel.

28.5.4.1 Prerequisites

Following are the prerequisites for registering a custom channel:

- RPM Repository must be under /var/www/html and accessible through http protocol.

- Metadata files should be created by running yum-arch and createrepo commands.

- Management Agent must be installed in the RPM repository machine and pointed to the OMS.

28.5.4.2 Procedure to Register a Custom Channel

Do the following to register a custom RPM Repository in Enterprise Manager Grid Control:

1. Log in to Enterprise Manager Grid Control.

2. Go to **Setup** and select **Patching Setup**.

3. In the **Linux Patching Setup** tab, click the **Manage RPM Repository** link.

4. In the Manage Repository Home page, click **Register Custom Channel**.

5. In the Register Custom Channel page, specify a unique channel name.

6. Click **Browse** and select the host where custom RPM repository was setup.

7. Specify the path where RPM repository resides. The directory location must start with `/var/www/html/`. Click **OK**.

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

28.5.5 Cloning a Channel

You can clone a channel.

28.5.5.1 Prerequisites

Following are the prerequisites before you can clone a channel:

1. There must at least one channel already present.

2. Patching user must have read/write access in both the source and target channel machine.

3. Check if enough space is present in the target channel machine.

4. Patch user must have write access under the agent home. Patch user must have sudo privilege.

28.5.5.2 Procedure to Clone a Channel

Follow these steps to clone a channel:

1. Log in to Enterprise Manager Grid Control.

2. Go to **Setup** and select **Patching Setup**.

3. In the **Linux Patching Setup** tab, click the **Manage RPM Repository** link.

4. Select the source channel that you want to create-like (clone) and click **Create Like**.

5. Enter the credentials to use for the source channel. The credentials must have both read and write access.

6. Enter a unique target channel name.

7. Click **Browse** to select the target host name.

8. Enter the directory location of the target channel. This directory should be under `/var/www/html`.

9. Enter the credentials to use for the target channel. This credential should have both read and write access. Click **OK**.

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

28.5.6 Copying Packages from One Channel to Another

You can copy packages from one channel to another.

28.5.6.1 Prerequisites

Following are prerequisites to copy packages from one channel to another:

1. Atleast 2 channels must be present.

2. Patching user must have read/write access on both the source and target channel machine.

3. The target channel machine must have adequate space.

4. Patch user must have write access under the agent home. Patch user must have sudo privilege.

28.5.6.2 Procedure to Copy Packages from One Channel to Another

Follow these steps to copy packages from one channel to another:

1. Log in to Enterprise Manager Grid Control.

2. Go to **Setup** and select **Patching Setup**.

3. In the **Linux Patching Setup** tab, click the **Manage RPM Repository** link.

4. Select the source channel and click **Copy Packages**.

5. Select the target channel.

6. From the source channel section, select and copy the packages to the target channel section.

7. Enter credentials for the source and target channels. These credentials should have read/write access to the machines. Click **OK**.

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

28.5.7 Adding Custom Packages to a Channel

You can add custom RPMs to a channel.

28.5.7.1 Prerequisites

Following are the prerequisites for adding custom RPMs to a channel:

1. Atleast one channel must be present.

2. Patching user must have write access on the channel machine.

3. Patch user must have write access under the agent home. Patch user must have sudo privilege.

28.5.7.2 Procedure to Add Custom Packages to a Channel

Follow these steps to add custom RPMs to a channel:

1. Log in to Enterprise Manager Grid Control.

2. Go to **Setup** and select **Patching Setup**.

3. In the **Linux Patching Setup** tab, click the **Manage RPM Repository** link.

4. Select the channel name where you want to add the RPM and click **Add**.

5. Select the source target name and the credentials to be used for the machine. The credentials you use must have read/write access.

6. In the Upload Files section, click the search icon to browse for the RPM files.

7. Enter the credentials to be used on the channel's machine. Click **OK**.

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

28.5.8 Deleting a Channel

You can delete a channel.

28.5.8.1 Prerequisites

Following are the prerequisites for deleting a channel:

1. Atleast one channel must be present.

2. Patching user must have write access to delete the RPM files from the channel machine.

3. Patch user must have write access under the agent home. Patch user must have sudo privilege.

28.5.8.2 Procedure to Delete a Channel

Follow these steps to delete a channel:

1. Log in to Enterprise Manager Grid Control.

2. Go to **Setup** and select **Patching Setup**.

3. In the **Linux Patching Setup** tab, click the **Manage RPM Repository** link.

4. Select the channel name you want to delete and click **Delete**.

5. If you want to delete the packages from the RPM Repository machine, select the check box and specify credentials for the RPM Repository machine. Click **Yes**.

6. If you have not selected to delete the packages from RPM Repository machine, you will get a confirmation message saying "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.

28.6 Configuration File Management

This section explains the configuration file management activities. In particular, this section covers the following:

- Prerequisites
- Creating Config File Channel
- Uploading Configuration Files
- Importing Files
- Deploying Files
- Deleting Channels

28.6.1 Prerequisites

Check if Software Library is already configured in OMS or not. If not configured, create one.

28.6.2 Creating Config File Channel

To create a configuration file channel, do the following:

1. In the **Deployments** tab, click **Linux Patching**, and then click **Configuration Files**. Click **Create Config File Channel**.

2. Enter a unique channel name and description for the channel and click **OK**.

3. You will see a confirmation message saying that a new configuration file is created.

28.6.3 Uploading Configuration Files

You can upload configuration files.

28.6.3.1 Prerequisites

To upload configuration files, atleast one configuration file must be present.

28.6.3.2 Procedure to Upload Configuration Files

Follow these steps to upload configuration files:

1. In the **Deployments** tab, click **Linux Patching**, and then click **Configuration Files**.

2. Select the file and click **Upload Configuration Files**.

3. You can either upload files from local machine (where the browser is running) or from a remote machine (agent should be installed on this machine and pointed to OMS). Select the mode of upload.

4. In the File Upload section, enter the file name, path where the file will be deployed in the target machine, and browse for the file on the upload machine.

5. 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.

6. After selecting the files, click **OK**.

7. You will see a confirmation message saying "Files are uploaded".

28.6.4 Importing Files

You can import files.

28.6.4.1 Prerequisites

To import files, atleast two channels present.

28.6.4.2 Procedure to Import Files

Follow these steps to import files:

1. In the **Deployments** tab, click **Linux Patching**, and then click **Configuration Files**.

2. Select the source channel and click **Import Files**.

3. Select the target channel.

4. From Source channel section, select the files and copy it to the target channel section. Click **OK**.

5. You will see a confirmation message saying, "Selected files are successfully imported".

28.6.5 Deploying Files

You can deploy files.

28.6.5.1 Prerequisites

Following are the prerequisites to deploy files:

■ Patch user must have write access under the agent home. Patch user must have sudo privilege.

■ There should be at least one channel with some files uploaded

28.6.5.2 Procedure to Deploy Files

Follow these steps to deploy files:

1. In the **Deployments** tab, click **Linux Patching**, and then click **Configuration Files**.

2. Select the source channel and click **Deploy Files**.

3. In the wizard that opens up, select the files that you want to deploy and click **Next**.

4. Click **Add** to select the targets where you want to deploy the files.

5. Enter the credentials for the selected targets.

6. Enter the Pre/Post scripts that need to be applied before or after deploying the files.

7. Review the deploy parameters and click **Finish**.

8. A Deploy job will be submitted. Follow the job's link until it completes successfully.

28.6.6 Deleting Channels

You can delete channels.

28.6.6.1 Prerequisites

To delete channels, at least one configuration file must be present.

28.6.6.2 Procedure to Delete Channels

Follow these steps to delete a channel:

1. In the **Deployments** tab, click **Linux Patching**, and then click **Configuration Files**.

2. Select the channel to delete and click **Delete**. Click **Yes**.

3. You will see a confirmation message saying, "Config File Channel (channel name) successfully deleted".

29

Patching Solaris Hosts

Solaris Host Patching is a feature in Enterprise Manager Grid Control that helps in keeping the machines in an enterprise updated with security fixes and critical bug fixes, especially in a data centre or a server farm.

This feature support in Enterprise Manager Grid Control enables you to patch Solaris hosts.

This chapter explains how you can patch Solaris hosts. In particular, this chapter covers the following:

- Getting Started

- Deployment Procedures

- Supported Releases

- Patching Solaris Hosts

29.1 Getting Started

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

Table 29–1   Getting Started with Patching Solaris Hosts

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Solaris hosts.	To learn about the Deployment Procedure, see Section 29.2, "Deployment Procedures".
Step 2	**Knowing About the Supported** Releases Know what releases of Solaris are supported by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 29.3, "Supported Releases".

Table 29–1   (Cont.) Getting Started with Patching Solaris Hosts

Step	Description	Reference Links
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the patching environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for patching Solaris hosts, see Section 29.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Solaris hosts.	To patch Solaris hosts, follow the steps explained in Section 29.4.2, "Patching Procedure".

29.2 Deployment Procedures

Enterprise Manager Grid Control provides the following Deployment Procedures for Solaris patching:

- *Patch Solaris Hosts*

  This deployment procedure allows you to patch Solaris hosts. The steps to run this deployment procedure is specified in the next section.

29.3 Supported Releases

Following are the releases supported for Solaris patching:

- Solaris 8 and higher

29.4 Patching Solaris Hosts

This section explains how you can patch Solaris hosts. This section covers the following:

- Prerequisites
- Patching Procedure

29.4.1 Prerequisites

Before running the deployment procedure, ensure that you meet the following requirements:

- Ensure that you meet the prerequisites described in the following sections:
  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"
  - Section 3.2, "Optional Infrastructure Requirements"
- Software Library must be configured in OMS.
- All the target machines must have sudo enabled for the patch user.

29.4.2 Patching Procedure

Follow these steps to patch Solaris hosts using deployment procedures:

1. In the **Deployments** tab, click **Patching through Deployment Procedures**. Select **Patch Solaris Hosts** and click **Schedule Deployment**.

2. In the Package Repository page, specify the Solaris package repository source.

    You can choose to get the updates from the following locations:

    ■   **Update from SUN Server** - Connect to the SUN server and download the Solaris updates from this Web site.

    ■   **Update from Local Server** - set up a repository on a local server. To do this, you will require a Sun Spectrum contract account. Refer to the Sun web site for more details.

    Depending on the type of application policy you select, you can either just stage the updates (and not apply them automatically), or choose to stage and apply the updates when the deployment procedure is executed.

    To simply stage the updates and not apply, select **Stage Updates**.

    To stage and apply the updates, select **Apply Updates**.

The patches when applied will have metadata (updates) that are classified as the following types:

- **Standard:** Updates of this type will take effect as soon as they are applied. The effects of this patch are visible as soon as it is applied, unless the application being patched is running while the patch is applied.

- **Reboot-After:** Updates of this type will take effect only after a reboot.

- **Reconfigure-After:** Updates of this type will take effect only after the application is reconfigured.

- **Reboot-Immediate:** If you select updates of this type, the selected hosts must be restarted immediately after the patch is applied.

- **Reconfigure-Immediate:** If you select updates of this type, the selected hosts must be reconfigured immediately after the patch is applied. If the patched hosts are not reconfigured immediately, such hosts become unstable.

3. In the Select Updates page, select the packages to be updated. Click **Next**.

4. In the Select Hosts page, select the targets to be updated. You can also select a group by changing the target type to group. Click **Next**.

5. In the Credentials page, enter the credentials to be used for the updates. Click **Next**.

6. In the Pre/Post Scripts page, specify any scripts to be run before and after patching. Click **Next**.

7. In the Schedule page, enter the schedule to be used. Click **Next**.

8. Review the update parameters and click **Finish**.

9. A deployment procedure will be submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

10. Once the procedure is completed, you can view the updated versions of the packages by clicking Targets, selecting the target name, clicking Configuration, and then clicking Packages.

30

Patching Microsoft Windows Hosts

Microsoft Windows Host Patching is a feature in Enterprise Manager Grid Control that helps in keeping the machines in an enterprise updated with security fixes and critical bug fixes, especially in a data centre or a server farm.

This feature support in Enterprise Manager Grid Control enables you to patch Microsoft Windows hosts.

This chapter explains how you can patch Microsoft Windows hosts. In particular, this chapter covers the following:

- Getting Started
- Deployment Procedures
- Supported Releases
- Patching Microsoft Windows Hosts

30.1 Getting Started

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

Table 30–1   Getting Started with Patching Microsoft Windows Hosts

Step	Description	Reference Links
Step 1	**Understanding the Deployment Procedure** Understand the Deployment Procedure that is offered by Enterprise Manager Grid Control for patching Microsoft Windows hosts.	To learn about the Deployment Procedure, see Section 30.2, "Deployment Procedures".
Step 2	**Knowing About the Supported** Releases Know what releases of Windows are supported by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Section 30.3, "Supported Releases".

Table 30–1 (Cont.) Getting Started with Patching Microsoft Windows Hosts

Step	Description	Reference Links
Step 3	**Meeting the Prerequisites** Before you run any Deployment Procedure, you must meet the prerequisites, such as setting up of the patching environment, applying mandatory patches, setting up of Oracle Software Library.	To learn about the prerequisites for patching Microsoft Windows hosts, see Section 30.4.1, "Prerequisites".
Step 4	**Running the Deployment Procedure** Run the Deployment Procedure to successfully patch Microsoft Windows hosts.	To patch Microsoft Windows hosts, follow the steps explained in Section 30.4.2, "Patching Procedure".

30.2 Deployment Procedures

Enterprise Manager Grid Control provides the following Deployment Procedures for Microsoft Windows patching:

- *Patch Windows Hosts*

  This deployment procedure allows you to patch Microsoft Windows hosts. The steps to run this deployment procedure is specified in the next section.

30.3 Supported Releases

Following are the releases supported for Microsoft Windows patching:

- Windows XP

- Windows 2000

30.4 Patching Microsoft Windows Hosts

This section explains how you can patch Microsoft Windows hosts. This section covers the following:

- Prerequisites

- Patching Procedure

30.4.1 Prerequisites

Before running the deployment procedure, ensure that you meet the following requirements:

- Ensure that you meet the prerequisites described in the following sections:

  - Section 3.1, "Mandatory Infrastructure Requirements for Provisioning and Patching"

  - Section 3.2, "Optional Infrastructure Requirements"

- Software Library must be configured in OMS.

- Microsoft Windows repository must be set up.

30.4.2 Patching Procedure

Follow these steps to patch Microsoft Windows hosts using deployment procedures:

1. In the **Deployments** tab, click **Patching through Deployment Procedures**. Select **Patch Windows Hosts** and click **Schedule Deployment**.

2. In the Package Repository page, enter the Microsoft Windows installer updates repository URL or select a repository configured earlier. You can also remove the selected URL from the database. Click **Next**.

3. In the Select Updates page, select the packages to be updated. Click **Next**.

4. In the Select Hosts page, select the targets to be updated. Click **Next**.

5. In the Credentials page, enter the credentials to be used for the updates. Click **Next**.

6. In the Pre/Post Scripts page, specify the scripts to be run before or after patching. Click **Next**.

7. In the Schedule page, enter the schedule for patching. Click **Next**.

8. Review the update parameters and click **Finish**.

9. A deployment procedure will be submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

Once the procedure is completed, you can view the updated versions of the packages by clicking Targets, selecting the target name, clicking Configuration, and then clicking Packages.

Part VII

Customization

This part explains how you can customize Deployment Procedures. If you are already familiar with Deployment Procedures and have already used the out-of-box procedures for provisioning and patching, then this part will help you understand how you can create custom copies of those default procedures with changes specific to your environment. As "*power users*" of Deployment Procedures, you will now learn to add additional custom steps, use authentication utilities to switch as another user, and so on.

In particular, this part contains the following chapters:

- Chapter 31, "Customizing Deployment Procedures"

31

Customizing Deployment Procedures

The Deployment Procedures offered by Enterprise Manager Grid Control are default procedures that have been created considering all the best practices in the industry. The steps embedded within a Deployment Procedure ensure that they meet all your provisioning and patching requirements. You can, of course, use them with the default settings to provision or patch your targets in the environment, however, you also have the choice of customizing them to include additional custom steps, disable unwanted steps, and use authentication tools to run some steps as another user.

By customizing Deployment Procedures, you can also implement different error handling methods. For example, in a patching operation where multiple hosts are patched in parallel, it may be wise to skip the host on which a failure occurs. However, failure on a device creation could render the remaining provisioning operation ineffective. Therefore, it may be necessary to abort the entire procedure for failure of such a step.

This chapter helps you understand how you can customize Deployment Procedures to make them suit your needs. In particular, this chapter covers the following:

- Understanding Customization Types
- Understanding the Limitations
- Creating Copy of Deployment Procedure
- Adding Phases or Steps
- Deleting Phases or Steps
- Enabling or Disabling Phases or Steps
- Customizing Steps to Run as Another User
- Changing Error Handling Modes
- Customizing Directives
- Setting Up E-Mail Notifications
- Customizing Deployment Procedures to RollBack Patches

31.1 Understanding Customization Types

The following describes the types of customization you can perform with Deployment Procedures:

Type 1	Type 2
Editing Default Deployment Procedures	**Creating Custom Deployment Procedures**
You can edit an existing, default Deployment Procedure that is offered by Enterprise Manager Grid Control.	You can create your own custom Deployment Procedure with new steps, phases, privilege levels, and so on. However, note that you can create custom Deployment Procedures only using Patching Deployment Procedures.
You can perform the following tasks:	You can perform the following tasks:
■ Add your own phases and steps	■ Delete all default phases and steps
■ Enable and disable phases and steps	■ Add your own phases and steps
■ Delete phases and steps	■ Enable and disable phases and steps
■ Change privilege levels	■ Change privilege levels
■ Change error handling modes	■ Change error handling modes
■ Enable e-mail notifications	■ Enable e-mail notifications

The following shows the different sections that need to be edited while customizing a Deployment Procedure to add or delete steps, enable or disable steps, enable e-mail notifications, and enable other privilege-level settings.

31.2 Understanding the Limitations

Before you start customizing the Deployment Procedures, keep in mind the following points:

■ When you customize any default patching Deployment Procedure offered by Enterprise Manager Grid Control, you **must not disable or delete the Stage step and the Apply Patch step**. You can disable or delete one of them, but not both. If you still do so, then the <Target> Updates section on the Software Updates page

of the Deployment Procedure will not appear. As a result, you will not be able to select a patch. If you want to add a custom step to replace them, them without disabling or deleting these steps, add an additional step while customizing the Deployment Procedure.

- **When you create a new custom patching Deployment Procedure** with the default steps disabled or deleted, you **must also disable or delete the Initialize step**. Otherwise, the custom Deployment Procedure will fail.

- If you do not have the privileges of a super administrator, then you will not be able to customize the Deployment Procedures (*The Create Like button will be disabled for the Deployment Procedures*). You will also not be able to update Oracle Software Library (Software Library).

- While cloning or provisioning an Oracle Database using a reference host, do the following:

  - If the Oracle home of the Management Agent is owned by the locked user, then run the Management Agent user-specific steps (for example, *Targets Discovery* step) as *sudo -S -u <locked_user>*. Otherwise, run the steps as a normal user.

  - If the Oracle home of the database you are cloning from the reference host is owned by the locked user, then run the reference host user-specific steps (for example, *Database Archive* step or *Grid Archive* step) as *sudo -S -u <locked_user>*. Otherwise, run the steps as a normal user.

31.3 Creating Copy of Deployment Procedure

The first step towards customizing a Deployment Procedure is to create a copy of the default Deployment Procedure that is offered by Enterprise Manager Grid Control. Note that only a copy can be edited and customized with your changes; the default Deployment Procedures must always and will always remain unchanged.

To create a copy of a default Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, edit the procedure to customize it according to your needs, and click **Save**.

31.4 Adding Phases or Steps

You can add additional phases or steps to a Deployment Procedure to run additional custom scripts, host commands, or jobs. For more information about phases and steps, see Section 2.2.2, "Understanding Phases and Steps".

This section explains how you can add different types of phases or steps to a Deployment Procedure. In particular, it covers the following:

- Adding Rolling or Parallel Phase

- Adding Job Step

- Adding Directive Step

- Adding Generic Component Step
- Adding Registered Component Step
- Adding Host Command Step
- Adding Manual Step

31.4.1 Adding Rolling or Parallel Phase

To add a rolling phase to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the phase before or after which you want to add a new phase, and click **Insert**.

5. In the Create wizard, do the following:

   a. On the Create page, specify general information about the phase as described in Table 31–1.

   b. On the Select Target List page, select a target list to indicate the type of targets on which the new phase should run.

   c. On the Review page, review the information you have provided for creating a new phase, and click **Finish**.

Table 31–1   Field Description - Adding Rolling Phase

Field Name	Description
Select	Select *Phase*.
Name	Specify a name for the custom phase.
Description	Provide a description for the custom phase.
Condition	Leave this field blank.
Insert Location	If you want to insert the custom phase after the phase or step you selected, then select **After <phase or step name>**. Otherwise, select **Before <phase or step>**.
Type	If you are adding a rolling phase, then select *Rolling*. If you are adding a parallel phase, then select *Parallel*.
Error Handling	Select the error handling mode you want to set for the custom phase. To understand these error handling modes, see Section 31.8, "Changing Error Handling Modes".

31.4.2 Adding Job Step

To add a job step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

   a. On the Create page, specify general information about the step as described in Table 31–2.

   b. On the Select Type page, select a job type that best describes the task that you want the step to perform. For example, if you want to job to transfer files across the network, then select **File Transfer**.

   c. On the Map Properties page, specify values for the parameters that are required by the selected job type.

   d. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

31.4.3 Adding Directive Step

To add a directive step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

   a. On the Create page, specify general information about the step as described in Table 31–2.

   b. On the Select Directive page, select a directive from the Software Library that you want to run on the targets.

   c. On the Map Properties page, specify values for the properties associated with the selected directive.

   d. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

31.4.4 Adding Generic Component Step

To add a generic component step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

a. On the Create page, specify general information about the step as described in Table 31–2.

b. On the Select Component page, select a component from the Software Library that you want to stage on the targets.

c. On the Select Directive page, select a directive from the Software Library that you want to run on the targets.

d. On the Map Properties page, specify values for the properties associated with the selected component and directive.

e. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

31.4.5 Adding Registered Component Step

To add a registered component step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

a. On the Create page, specify general information about the step as described in Table 31–2.

b. On the Select Type page, select a component type.

The list is a subset of the component types displayed on the Components page in the Provisioning tab. Each type in the list has a specific job associated with it. These jobs stage and install the components.

c. On the Select Component page, select a component from the Software Library that you want to stage on the targets.

d. On the Map Properties page, specify values for the parameters that are required by the job associated with the selected component.

e. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

31.4.6 Adding Host Command Step

To add a host command step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

   a. On the Create page, specify general information about the step as described in Table 31–2.

   b. On the Enter Command page, specify the command or script, which you want to run on the target, and the privilege to run it.

      To run the host command as a script, select **Script** from the **Command Type** list. Specify the shell that can interpret the script. The script is passed as standard input to the specified interpreter.

      To run the host command as a command line, select **Single Operation** from the **Command Type** list. Specify the text you want to execute used as a command line. No assumptions are made about the shell to interpret this command line. The first entry in the command line is assumed to be the process to spawn and the rest of the command line as passed as arguments to this process. Therefore, a command line of `ls -a /tmp` spawns a process of "ls" (from the current path; also depends on the Oracle Management Agent) and passes "-a" as the first argument and then "/tmp" as the second argument to this process.

      Note: The command line mode assumes that the first part of the command line is the process to be spawned. Therefore, shell internals and the commands that rely on the PATH environment variable for resolution are not recognized. If any such commands need to be used, then you need to prepend the shell that interprets the command line.

      For example, the command `cd /tmp && rm -rf x` expands to "cd" as a process and then "/tmp, &&, rm, -rf, x" as arguments. To fix this, change the command line to `/bin/csh -c "cd /tmp && rm -rf x"`.

      Another example, the command `export PATH=/opt:${PATH}; myopt -install` expands to "export" as a process and then "PATH=/opt:${PATH};, myopt, -install" as arguments. To fix this, use `/bin/sh -c "export PATH=/opt:${PATH}; myopt -install"`.

   c. On the Map Properties page, specify values for the parameters that are required by the job associated with the selected component.c

   d. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

31.4.7 Adding Manual Step

To add a manual step to a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step before or after which you want to add a new step, and click **Insert**.

5. In the Create wizard, do the following:

   a. On the Create page, specify general information about the step as described in Table 31–2.

b. On the Enter Instructions page, provide a message to inform the operator about a manual step. For example, if want to instruct the operator to log in to a system and update the kernel parameter, then specify the following:

```
You have been logged out of the system. Log in and update the Kernel
parameters.
```

c. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Table 31–2   Field Description - Adding Steps

Field Name	Description
Select	Select *Step*.
Name	Specify a name for the custom step.
Description	Provide a description for the custom step.
Condition	Leave this field blank.
Insert Location	If you want to insert the custom step after the step you selected, then select **After <step name>**. Otherwise, select **Before <step>**.
Type	- For a job step, select **Job**. - For a directive step, select **Library: Directive**. - For a generic component, select **Library: Generic Component**. - For a registered component, select **Library: Registered Component**. - For a manual step, select **Manual**. - For a host command step, select **Host Command**.
Error Handling	Select the error handling mode you want to set for the custom step. To understand these error handling modes, see Section 31.8, "Changing Error Handling Modes".

31.5 Deleting Phases or Steps

You can delete the phases or steps that you do not want in a Deployment Procedure. For more information about phases and steps, see Section 2.2.2, "Understanding Phases and Steps".

To delete phases or steps in a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select the step or phase you want to delete, and click **Delete**.

   Oracle recommends that you disable the steps or phases instead of deleting them because steps or phases once deleted cannot be retrieved, but steps or phases disabled can always be enabled later. For information about enabling and disabling steps or phases, see Section 31.6, "Enabling or Disabling Phases or Steps".

31.6 Enabling or Disabling Phases or Steps

If you do not want to have some phases or steps in a Deployment Procedure, you can always disable them instead of deleting them. This is a preferred option because phases or steps once deleted cannot be retrieved, but phases or steps disabled can always be enabled later. For more information about phases and steps, see Section 2.2.2, "Understanding Phases and Steps".

To enable or disable phases or steps in a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, do the following:

    a. To disable a phase or step, select the phase or step you want to disable, and click **Disable**.

    b. To enable a phase or step, select the phase or step you want to enable, and click **Enable**.

The following is an example that illustrates how you customize the *Oracle Database Provisioning* Deployment Procedure to enable the *Prepare Create Directory* step:
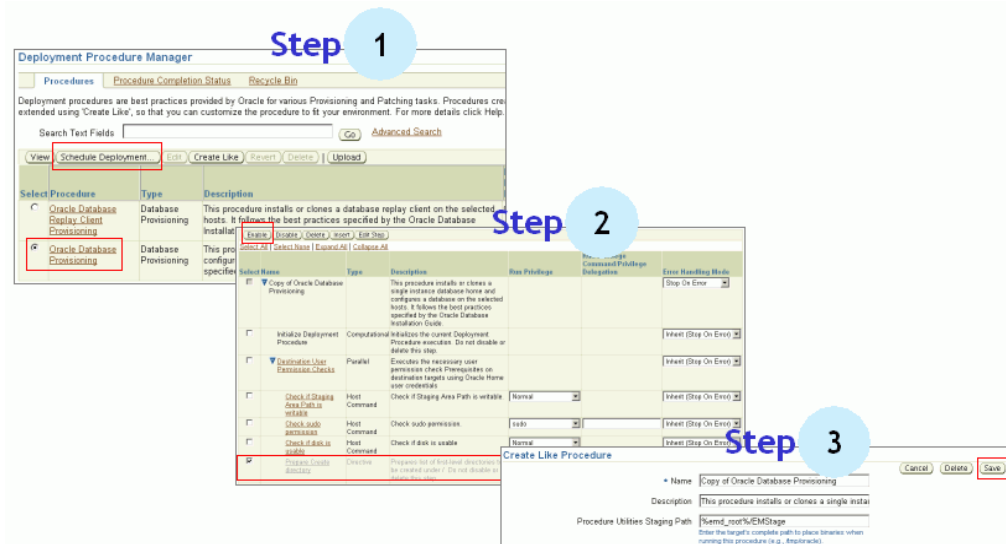
31.7 Customizing Steps to Run as Another User

All the Deployment Procedures offered by Enterprise Manager Grid Control require administrator privileges to run. While most steps within a Deployment Procedure can be run as a normal user, there are some steps that require special permissions and privileges, and unless you provide the administrator's credentials, you cannot proceed with the deployment. For more information about steps, see Section 2.2.2, "Understanding Phases and Steps".

Under such circumstances, you can do one of the following. Although the former option is recommended, you are always free to use the latter option to suit your needs.

- Customize the Deployment Procedure to disable the steps that require special privileges, run the other steps as a normal user, and have the administrator run the disabled steps later.

- Use authentication utilities to run some steps within the Deployment Procedure with the privileges of another user. The authentication utilities supported by Enterprise Manager Grid Control are SUDO, PowerBroker, and so on. This support is offered using the Privilege Delegation mechanism available in Enterprise Manager Grid Control. For more information, see Section 31.7.3, "Creating Privilege Delegation Templates".

This section how you can customize the steps within a Deployment Procedure to run them as another user. In particular, this section covers:

- Using Privilege Delegation

- Updating SUDOERS File with a List of Commands

- Creating Privilege Delegation Templates

- Using SUDO, PowerBroker, and Privilege Delegation Settings

31.7.1 Using Privilege Delegation

While SUDO and PowerBroker are third-party utilities supported in Enterprise Manager Grid Control, Privilege Delegation is proprietary to Oracle. Privilege Delegation is a framework that allows you to use either SUDO or PowerBroker to perform an activity with the privileges of another user.

Privilege Delegation can use either SUDO or PowerBroker, but not both, and the settings are only for a single host. Therefore, if a host is set up with pbrun, then it will use only pbrun.

Caution: The SUDO and PAM options are displayed in the **Run Privilege** list, which is available while customizing a Deployment Procedure. However, the options are displayed only for legacy reasons, and they will soon be deprecated. Consequently, you must use only Privilege Delegation.

Note: The certified SUDO versions are 1.6.7 to 1.6.9. The certified PBRUN versions are 4.0.8 and 5.x. Higher versions of these utilities may continue to work unless some fundamental changes have been introduced to their behavior.

Privilege Delegation offers the following advantages:

- You have the flexibility to use either SUDO or PowerBroker within the same framework.

- Using the framework, you can now run PowerBroker in a passwordless or password-protected mode.

- You can create a template with these Privilege Delegation settings and reuse it for multiple hosts. This not only allows you to standardize Privilege Delegation setting across your enterprise, but also facilitates the process of configuring Privilege Delegation Settings. It simplifies the Privilege Delegation setting management as well.

■ You can use the Privilege Delegation settings not only for deployment procedures, but also for jobs in Enterprise Manager Grid Control.

■ Privilege Delegation can read passwords from both STDIN and TTY.

31.7.2 Updating SUDOERS File with a List of Commands

If you want to use SUDO authentication utility, then before editing a Deployment Procedure, update the `/etc/sudoers` file to allow a normal user to switch to another user who has the privileges to run the Deployment Procedure. Also, if you want to restrict the normal user to have SUDO access only to certain commands, then specify a list of those commands for which SUDO access is required.

The following provides a list of commands for which SUDO access is required for different Deployment Procedures:

Table 31–3   Updating SUDOERS File with a List of Commands

Deployment Procedure Name	Commands Requiring SUDO Access

All Deployment Procedures	The following are the minimum additions required to the `/etc/sudoers` file to run any job or Deployment Procedure:
	
	`$<ORACLE_HOME>/bin/nmosudo *`
	
	`$<ORACLE_HOME>/perl/bin/perl`
	
	Here, `<ORACLE_HOME>` refers to the fully qualified path to the Oracle home of the Management Agent.

■ Oracle Clusterware / RAC Provisioning For UNIX ■ Oracle Clusterware / RAC Provisioning For Windows ■ One Click Extend Cluster Database ■ Delete/Scale down Oracle Real Application Clusters ■ Oracle Grid Infrastructure Provisioning for Standalone Servers ■ Oracle Grid Infrastructure / RAC Provisioning	`/bin/sh, $ENV{EMDROOT}/perl/bin/perl, $ENV{EMDROOT}/bin/nmupm, $ENV{EMDROOT}/bin/nmhs, /bin/pwd, /usr/bin/tee, /bin/echo, /bin/cut,/sbin/ifconfig, /usr/bin/tr, /bin/awk,/bin/rpm, /bin/grep,/bin/sed, /sbin/sysctl, /usr/bin/free, /usr/sbin/useradd, /bin/cp,/sbin/service, /usr/bin/head, /usr/bin/tail, /bin/chmod, /bin/chown, /sbin/ifconfig, /usr/sbin/usermod, /sbin/depmod, /bin/uname, /bin/ln, /bin/umount, /bin/mount, /bin/mkdir, /bin/ls, /bin/rm, /bin/basename, /bin/netstat, /bin/df, /bin/hostname, /bin/mv`  In addition, add the following line:  `env_keep +="PERL5LIB"`

31.7.3 Creating Privilege Delegation Templates

If you want to use Privilege Delegation authentication utility, then before editing a Deployment Procedure, create a Privilege Delegation template with the required settings for a host. To do so, follow these steps:

1. In Grid Control, from the top-right corner, click **Setup**.

2. On the Overview of Setup page, from the vertical menu, click **Manage Privilege Delegation Settings**.

3. On the Manage Privilege Delegation Settings page, from the Related Links section, click **Manage Privilege Delegation Settings Template**.

4. On the Manage Privilege Delegation Settings Templates page, from the **Create** list, select a privilege delegation type, either **Sudo** or **PowerBroker**, and click **Go**.

5. On the Create '<delegation type>' Setting Template page, provide a name for the template and specify the location where SUDO or PowerBroker is located (for PowerBroker, you can optionally provide the password prompt), and click **Save**.

6. On the Manage Privilege Delegation Setting Templates page, select the template you created and click **Go**.

7. On the Apply '<delegation type>' Setting: New page, click **Add Targets** to apply the privilege delegation template settings to selected hosts, and click **Apply**.

> **Note:** If you do not apply the privilege delegation template to a target, and if you configure a step in the deployment procedure to run in Privilege Delegation mode, then the deployment procedure for that target runs the step in normal mode instead.

---

8. After creating the template, to understand how it can be applied to Deployment Procedures, see Section 31.7.5, "Using SUDO, PowerBroker, and Privilege Delegation Settings".

31.7.4 Testing Privilege Delegation Settings

After creating a privilege delegation template and before applying it to a Deployment Procedure, Oracle recommends you to test the privilege delegation setting.

The following is an example that describes how you can register your credentials as preferred credentials and also choose to run as another user, and then test the settings by creating a job that checks whether a command is being as normal user or as another user using privilege delegation mechanism.

1. In Grid Control, from the top-right coner, click **Preferrences**.

2. On the General page, from the left menu bar, click **Preferred Credentials**.

3. On the Preferred Credentials page, for the Host target type, click the Set Credentials icon.

4. On the Host Preferred Credentials page, in the Target Credentials section, specify the normal user name, the normal password, and the Run as user name that you want to switch over to using the privilege delegation mechanism. Then click **Apply**.

5. After registering the credentials as preferred credentials, click the **Jobs** tab.

6. On the Job Activity page, from the Create Job list, select **OS Command**, and click **Go**.

7. On the Create OS Command Job page, in the General tab, specify a name for the job. Then, from the Target section, click **Add** to add the host on which you want to run the OS command.

8. In the Parameters tab, for **Command**, specify the command `id`.

9. Click **Submit**.

10. On the Job Activity page, click the job name you just created. Enterprise Manager Grid Control displays the status of the job. Click the status column to view its results.

Ideally, Enterprise Manager Grid Control should have switched over from normal user to another user, which you specified for **Run as** on the Host Preferred Credentials page, and run then run the OS command.

31.7.5 Using SUDO, PowerBroker, and Privilege Delegation Settings

You can use SUDO, PowerBroker, and so on globally using the Privilege Delegation mechanism, that is, for all steps, or only for some steps within a Deployment Procedure. This section explains the following:

- Making Global Settings
- Making Step-Level Settings

31.7.5.1 Making Global Settings

On the Deployment Procedure Manager page, select a Deployment Procedure and click **Create Like** to customize it. On the Create Like Procedure page, for Privilege Delegation, you can specify the user and the profile that the step must run as. The values you provide here will be used for all the steps within a Deployment Procedure.

31.7.5.2 Making Step-Level Settings

On the Deployment Procedure Manager page, select a Deployment Procedure and click **Create Like** to customize it. On the Create Like Procedure page, for a particular step, from the Run Privilege column, select **Privilege Delegation**, and specify the *Run As* value in the first text box and the *Profile* value in the second text box. For example, for SUDO, the *Run As* value can be /usr/bin/sudo -u root, and the *Profile* value can be left blank. Similarly, for PbRun, the *Run As* value can be /opt/pbrun -l -u root, and the *Profile* value can be admin.
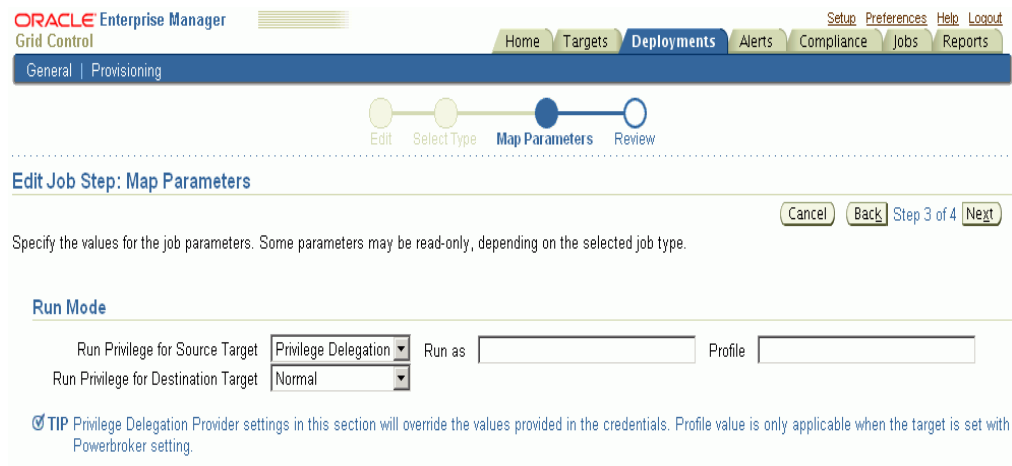
> **Note:** These values need to be overridden only if a different user is required to run that step. Otherwise, the default values are used from preferred credentials registered with Enterprise Manager Grid Control. For example, if the preferred credentials for the host have been set to Oracle as the *Run As* user, but if the current step needs to be run as root, then you can enter /usr/bin/sudo -u root in the **Run As** text box.

For file transfer-based job steps, you can apply the Privilege Delegation settings while editing that step. For example, one of the steps in the One Click Extend Cluster Database deployment procedure is *Copy Archives*. This is a file transfer-based job step for which you can apply the Privilege Delegation settings. To do so, you can click the step name and on the Map Parameters page, in the Run Mode section, from the **Run**

Privilege for Source Target list and **Run Privilege for Destination Target** list, select **Privilege Delegation**. Then provide the *Run As* value and the profile that must be used.

> **Note:** For configuration collection-based job step, you can NOT apply any Privilege Delegation settings. For example, one of the steps in the One Click Extend Cluster Database deployment procedure is *Refresh Host Configuration*. This is a configuration collection-based step for which Privilege Delegation settings cannot be applied.

31.8 Changing Error Handling Modes

Every step in a Deployment Procedure is preconfigured with an error handling mode that indicates how the Deployment Procedure will behave when the phase or step encounters an error. The error handling modes offered by Enterprise Manager Grid Control are:

- *Inherit* - Inherits the error handling mode that was set for the enclosing phase. (When set for a step that is outside a phase, it inherits the error handling mode from the Deployment Procedure).

- *Stop On Error* - Stops when an error is encountered. Deployment Procedure does not proceed to the next step until you correct the errors or override them.

- *Continue On Error* - Continues even when an error is encountered.

- *Skip Target* - Ignores the failed target on the list and continues with other targets.

For more information about steps, see Section 2.2.2, "Understanding Phases and Steps".

To change the error handling modes, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, select a phase or step, and from the list provided in the **Error Handling Mode** column against that phase or step, select the mode you want to set.

Once the mode is selected from the list, Enterprise Manager Grid Control automatically refreshes the page with the newly selected mode for that phase or step.

The following is an example that illustrates how you customize the *Oracle Database Provisioning* Deployment Procedure to change the error handling mode of the *Destination User Permission Checks* phase:
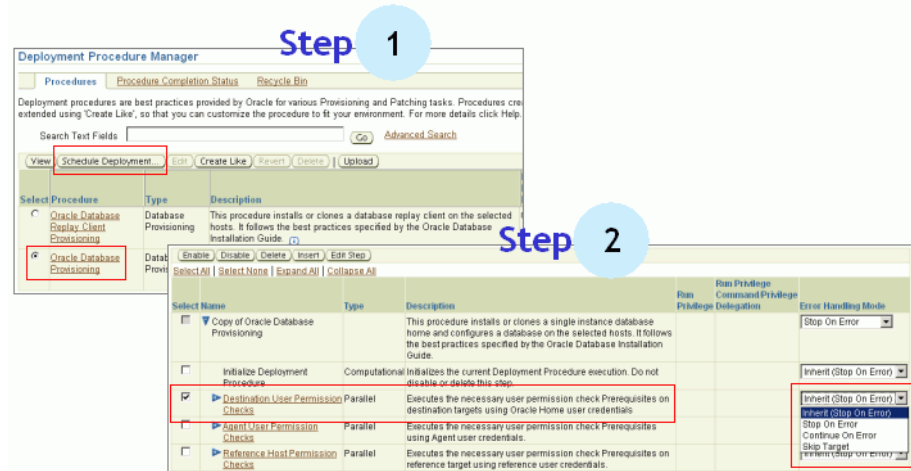
31.9 Customizing Directives

Directives are essentially scripts stored in the Software Library. They are used in Deployment Procedures within a Directive Step, which is a special type of action step. For more information about Directive Step, see Section 2.2.2, "Understanding Phases and Steps".

If you want to customize a directive offered by Enterprise Manager Grid Control, then first create a copy of the Perl script associated with that Directive and make a new directive out of that copy. Then customize the Deployment Procedure to modify a step to use this new directive, and then schedule the deployment. This section explains the following:

- Creating Copy of a Default Directive

- Storing the New Directive

- Customizing Deployment Procedure to Use the New Directive

- Running the Customized Deployment Procedure

31.9.1 Creating Copy of a Default Directive

To create a new customized directive using a default directive, follow these steps:

1. In Grid Control, click the **Deployments** tab and then the **Provisioning** secondary tab.

2. On the Provisioning page, click **Directives**.

3. On the Directives page, from the table, expand **Oracle Directives** and then expand other levels under this category to reach the directive you want to copy.

   For example, if you want to copy the Apply Patch directive of a patching operation, then expand **Oracle Directives** and then expand **Patching**. From this level, expand **Common**, and then **All**, and finally **Generic**. Under Generic, you should see the directive **Apply Patch**.

4. Select the directive you want to copy and click **Edit**. Open the Perl script, copy the contents of it to a local editor, and save it locally.

31.9.2 Storing the New Directive

To store the new directive, follow these steps:

1. In Grid Control, click the **Deployments** tab and then the **Provisioning** secondary tab.

2. On the Provisioning page, click **Directives**.

3. On the Directives page, click **Create Folder** to create a folder titled *Custom Directives*. This is to ensure that you separate the custom directives from the default directives.

4. On the Directives page, from the table, select **Custom Directives** and click **Create Directives**.

5. In the Create Directive wizard, do the following:

   a. On the Describe page, describe the directive you are creating.

   b. On the Configure page, click **Add** to specify the command line arguments to be passed to the Directive. Set the **Shell Type** to Perl because you are adding a Perl script. If the script is neither Perl nor Bash, then set it to **Defined in Script**.

      Each entry represents a single command line argument. Each argument may include a variable to be set later, a prefix and suffix. The prefix and suffix text are appended before and after the property value to produce the command line argument.

      Repeat this step to add all the command line arguments.

   c. On the Upload File page, select **Upload from Local Machine** and click **Browse** to select the script from the local system where you saved the file.

   d. On the Review page, review the information you have provided for this new directive, and click **Finish**.

31.9.3 Customizing Deployment Procedure to Use the New Directive

To customize a Deployment Procedure to use the new directive, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure for which you want to use this new directive, and click **Create Like**.

4. On the Create Like Procedure page, do the following:

a. From the table that lists all the steps within that Deployment Procedure, select the directive step with which you want to associate the new directive, and click **Edit Step**.

b. In the Edit Directive Step wizard, do the following:

   a. On the Edit page, click **Next**.

   b. On the Select Directive page, select **Select New Directive**. Then search and select the new directive you created, and click **Next**.

   c. On the Map Properties page, specify the values for the directive properties, and click **Next**.

   d. On the Review page, click **Finish**.

c. Click **Save**.

31.9.4 Running the Customized Deployment Procedure

To run the customized Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the customized Deployment Procedure and click **Schedule Deployment**.

31.10 Setting Up E-Mail Notifications

Enterprise Manager Grid Control can send e-mail notifications to report the status of a Deployment Procedure. However, by default, Deployment Procedures do not have this feature enabled. To configure Deployment Procedures to send e-mail notifications, you must first customize them and enable the e-mail notification option, and then enable the Send E-mail option for the standard PAF Status Notification rule. This section describes how you can enable e-mail notifications for Deployment Procedures. In particular, this section covers the following:

- Enabling E-Mail Notifications
- Customizing PAF Status Notification Rule

> **IMPORTANT:** As a prerequisite, you are expected to have configured the Mail Server and set up the e-mail address in Enterprise Manager Grid Control. If you have not already done this, see the chapter on *Configuring Notifications* in the *Enterprise Manager Configuration Guide* available at:
>
> http://www.oracle.com/technology/documentation/oem.html

31.10.1 Enabling E-Mail Notifications

To enable e-mail notifications in a Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the Deployment Procedure for which you want to set up e-mail notification, and click **Create Like**.

4. On the Create Like Procedure page, do the following:

   a. Select **Enable Notification**.

   b. (Optional) For **Procedure Status Notification Job Tag**, provide a notification tag name.

   > **Note:** Ensure that you specify not more than 30 characters. The text you specify here will be prepended to the notification job name.

   c. For **Status for which Notification is to be Sent**, select the state for which you want the notifications to be sent.

   d. Click **Save**.

   Enterprise Manager Grid Control displays the Deployment Procedure Manager page.

5. On the Deployment Procedure Manager page, from the top-right corner, click **Preferences**.

6. On the General page, do the following:

   a. In the E-mail Address section, click **Add Another Row** to add a row in the table, and specify the e-mail address to which the notification should be sent.

   b. Click **Apply**.

   c. From the left menu bar, click **Rules**.

7. On the Notification Rules page, do the following:

   a. From the table, select **PAF Status Notification**, and for this notification rule, enable the Send E-mail option from the **Subscribe (Send E-mail)** column.

   b. Click **Apply**.

31.10.2 Customizing PAF Status Notification Rule

You can customize the standard PAF Status Notification rule to receive notifications in required ways for specific deployment procedures. For example, you might want to be notified by e-mail for a test system procedure, but for a production run you might want to be informed of the status through SMS Alerts.

To incorporate specific requirements and enable different methods of notification, you should use the *Create Like* function to modify the standard notification rule and edit the job with the specific notification tag name used in the Deployment Procedure and associating specific method of notification from the pre-defined notification methods.

To customize the standard PAF Status Notification rule, follow these steps:

1. In Grid Control, from the top-right corner, click **Preferences**.

2. On the General page, from the left menu bar, click **Rules**.

3. On the Notification Rules page, from the table, select **PAF Status Notification**, and click **Create Like**.

4. On the Create Notification Rule page, do the following:

a. On the General page, specify a notification rule name and a brief description for the rule.

b.

31.11 Customizing Deployment Procedures to RollBack Patches

To roll back patches, you must first download the rollback Perl script, then upload it as a directive to the Software Library, and then use it with a customized Deployment Procedure. This section explains how can perform each of these tasks. In particular, this section covers the following:

- Downloading Rollback Perl Script
- Storing the Script as a Directive
- Customizing Deployment Procedure to Use Rollback Directive
- Running the Customized Deployment Procedure

31.11.1 Downloading Rollback Perl Script

To roll back patches, you have to use the Perl script `rollback_patch.pl`. This script is available as an attachment to My Oracle Support Note: 577557.1. Access this note and download the script to your local system.

31.11.2 Storing the Script as a Directive

To store the `rollback_patch.pl` script as a directive in the Software Library, follow these steps:

1. In Grid Control, click the **Deployments** tab and then the **Provisioning** secondary tab.

2. On the Provisioning page, click **Directives**.

3. On the Directives page, click **Create Folder** to create a folder titled *Custom Directives*. This is to ensure that you separate the custom directives from the default directives.

4. On the Directives page, from the table, select **Custom Directives** and click **Create Directives**.

5. In the Create Directive wizard, do the following:

   a. On the Describe page, specify the name *Rollback Patches* and provide the description *Directive to roll back patches in rollback or nrollback mode*. Click **Next**.

   b. On the Configure page, click **Add** to specify the following command line arguments to be passed to the Directive: *oh*, *patchid*, *stageDir*, *command*, *patchsetRelease*.

   Set the **Shell Type** to Perl because you are adding a Perl script.

   c. On the Upload File page, select **Upload from Local Machine** and click **Browse** to select the script from the local system where you saved the `rollback_patch.pl` file.

   d. On the Review page, review the information you have provided for this new directive, and click **Finish**.

31.11.3 Customizing Deployment Procedure to Use Rollback Directive

To customize a Deployment Procedure to use the new *Rollback Patches* directive, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select a patching Deployment Procedure that you want to customize for using this new directive, and click **Create Like**.

4. On the Create Like Procedure page, do the following:

   a. From the table that lists all the steps within that Deployment Procedure, select the Apply Patch directive step with which you want to associate the new directive, and click **Insert**.

   b. In the Insert Directive Step wizard, do the following:

      a. On the Create page, provide a name and description. From the **Insert Location** list, select **Before "Apply Patches"**. From the **Type** list, select **Library: Directive**. Click **Next**.

      b. On the Select Directive page, select **Select New Directive**. Then search and select the new *Rollback Patches* directive you created. Select **Always Use Latest Revision**. Click **Next**.

      c. On the Map Properties page, specify the following values for the directive properties, and click **Next**.

Table 31–4   Map Properties

Directive Property	Display Name to be Specified	Description to be Specified	Variable to be Specified
Oh	ORACLE_HOME	Oracle home of the target	Select **Choose Variable** and then select **Target Variable: oraHome** from the associated list.
patchid	Patch ID	Comma-separated list of patch numbers	Select **Ask User at Run-Time**
stageDir	Stage Directory	Directory where the patch can be staged	Select **Choose Variable** and then select **replacedStageDir** from the associated list. It can be also be set to a value where the patch is staged. However, this is required only for patches of release 10.1.0.x
command	Rollback Command	Rollback command to run	Select **Ask User at Run-Time**
patchsetRelease	Patch Set Release Version	Patch set release number	Select **Ask User at Run-Time**

   d. On the Review page, click **Finish**.

   c. Click **Save**.

31.11.4 Running the Customized Deployment Procedure

To run the customized Deployment Procedure, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. On the Deployment Procedure Manager page, select the customized Deployment Procedure and click **Schedule Deployment**.

4. In the Deployment Procedure Wizard, provide the required information and submit it for deployment.

   On the Library Step Properties page, you will be prompted to specify the patch ID, the rollback command, and the patch set release version. Specify the following vales on this page:

Table 31–5   Values for Library Step Properties Page for Rolling Back Patches

Library Step	Values to be Specified
Patch ID	Specify the patch number you want to roll back. If you want to roll back multiple patches, then specify the patch number separated by a comma.
Rollback Command	Specify *rollback* for normal patches and *nrollback* for Napply CPU patches.
Patchset Release Version	Specify the release for which the patch is being applied.

Part VIII

Enterprise Manager Command Line Interface (EMCLI)

This part describes how you can use Enterprise Manager Command Line Interface (EMCLI) commands to run Deployment Procedures.

This part contains the following chapters:

- Chapter 32, "Using EMCLI to Run Deployment Procedures"

32

Using EMCLI to Run Deployment Procedures

This chapter introduces the usage of EMCLI commands to run Deployment Procedures. In particular, this chapter covers the following:

- Before You Begin

- Prerequisites for Using EMCLI to Run Deployment Procedures

- Using EMCLI to Run Deployment Procedures

- Use Cases for EMCLI-based Provisioning and Patching

- Setting Up Preferred Credentials for Targets

- Queries to Acquire Data for Patching Runtime

- EMCLI Verbs for Deployment Procedures

- EMCLI Verbs for Deployment Instances

32.1 Before You Begin

Before you begin, familiarize yourself with the following:

- **RuntimeData xml**

  Runtime data response file (known as RuntimeData xml) is required to run any out-of-the-box or customized procedures. This file provides input for the configuration parameters consumed by a given procedure during execution.

  Each time you use the Enterprise Manager Console to run an out-of-box or customized Deployment Procedure, a RuntimeData xml file is automatically created based on the user input for the various parameters required by the procedure.

- **RuntimeData Template**

  Oracle provides out of the box templates for creating runtime data response files for the Deployment Procedures used in the most common use cases. These are known as "RuntimeData templates". These templates are available under the emcli/samples directory in OMS oracle home. The user needs to modify the configuration properties in these templates in order to generate RuntimeData xml file for a executing a procedure.

  For example, in order to provision RAC/AS you need to provide inputs such as install base location, shared device paths for the OCR, Voting disk and data files (in case of RAC). Similarly, for patching procedures inputs such as targets to be patched and patch number would be needed.

- **Procedure GUID**

Both out-of-box and customized procedures are associated with a global unique identifier (GUID). This GUID is required while executing the procedures using EMCLI.

- **Procedure Instance GUID**

  Each execution of a given out-of-box or customized procedures is associated with an instance global unique identifier (GUID). This instance GUID is generated at runtime and can be used to monitor the execution of the procedure.

- **Properties File**

  For every execution of Deployment Procedure you must modify the values of the required configuration parameters for a RuntimeData xml or RuntimeData template. Instead of manually editing the xml files, you can populate a simple properties with name-value pairs for listing values of configuration parameters such as hosts, platform for deployment, and so on.

- **Procedure Execution Scripts**

  Oracle provides out of box scripts for the execution of procedures for Provisioning and Patching. These Perl scripts are available under the emcli/scripts directory in the OMS Oracle home. The user needs to copy the scripts to the working directory location where the EMCLI client is setup for execution.

  The properties file, the associated RuntimeData xml or RuntimeData template, and GUID of the relevant procedure are then submitted as input to an out-of-box script which creates a new runtime data response file and runs the procedure.

  > **Note:** For more information on EMCLI, see *Oracle Enterprise Manager Command Line Interface Guide* available at:
  >
  > http://www.oracle.com/technology/documentation/oem.html

32.2 Prerequisites for Using EMCLI to Run Deployment Procedures

You must ensure that the following requirements are met prior to using EMCLI to run Deployment Procedures:

- EMCLI client must be set up. Refer to the Installation and Configuration section of the Oracle Enterprise Manager Command Line Interface 10g Release 5 (10.2.0.5.0) for configuring the EMCLI client. The document is available in the following location:

  http://download.oracle.com/docs/cd/B16240_01/doc/nav/portal_booklist.htm

- Targets that will be supplied to the Deployment Procedures are managed by Management Agents of version 10.2.0.2 or higher.

- Download the procedure execution scripts, out-of-box templates and properties files on the machine where the EMCLI client is setup. The out-of-box templates and properties files for patching and provisioning are available in the respective directories under OMS HOME/emcli/samples/.

- EMCLI-based patching and provisioning uses the Oracle Home credentials set in Oracle Enterprise Manager. The preferred credentials can be set for the targets during the execution of the Deployment Procedures in Oracle Enterprise Manager. It can also be set explicitly from the Oracle Enterprise Manager user interface or by using EMCLI. Refer to Section 32.5, "Setting Up Preferred Credentials for Targets" to set credentials for the Oracle Homes.

32.3 Using EMCLI to Run Deployment Procedures

Oracle provides out-of-the-box templates for creating run time data for Deployment Procedures used in the most common use cases. These are known as *runtimedata templates*. You can access them under the `emcli/samples` directory in the OMS Oracle home and you can then modify the configuration properties in these templates.

The following depicts the process to use EMCLI to run Deployment Procedures:

Figure 32–1   EMCLI Process to Run Deployment Procedures

There are four required actions to run Deployment Procedures. Those steps are described below:

1. Step 1: Find the GUID of the procedure to be run using EMCLI. This is a one-time activity.

2. Step 2: Obtain the RuntimeData xml or RuntimeData template for the procedure that needs to be run. This is also a one-time activity.

3. Step 3: Create the properties file for the RuntimeData xml or RuntimeData template. This step is required for each execution of the procedure.

4. Step 4: Submit the RuntimeData template or RuntimeData, properties file for the given execution and procedure GUID as input to an out-of-box script that will generate a new runtime data response file and then use this response file to run the procedure using EMCLI.

Each of these steps is discussed in detail in the subsequent sections.

32.3.1 Step 1: Finding Procedure GUID

EMCLI is case-sensitive so be sure to use the correct EMCLI verb and pass correct input. GUID out-of-box and customized procedures can be found using the following EMCLI verbs:

get_procedures

Usage:
emcli get_procedures -type="procedure type"

Description:
Get a list of Deployment Procedures.

Option:

```
-type="procedure type"
```
Display all the Deployment Procedure of type {procedure type}.

Output Columns:
GUID, Procedure Type, Name, Version, Created By

RAC procedures are of type: **RACPROV**
AS procedures are of type: **AS Provisioning**
The Standalone Database, RAC rolling, and CRS patching procedures are of type: **PatchOracleSoftware**

Alternatively the type associated with procedures can be found using the get_ procedure_type EMCLI command.

get_procedure_types

Usage:
emcli get_procedure_types

Description:
Get the list of all Deployment Procedure types

Output Columns:
Procedure Type

32.3.2 Step 2: Obtaining RuntimeData Template And RuntimeData XML

For out-of-box procedures, RuntimeData Templates are located in the emcli/samples directory in the Oracle Management Service (OMS) Oracle home.

For customized procedures you has to download the RuntimeData xml generated from an earlier execution of this procedure. To obtain the RuntimeData xml you first need to find the Instance GUID associated with the earlier execution of the procedure and then use it to download the RuntimeData xml generated for it. The following EMCLI verbs can be used to download a RuntimeData xml:

emcli get_instances -type="procedure type"

Usage:
emcli get_instances -type="procedure type"

Description:
Display list of procedure instances. EMCLI verb to obtain Instance GUID associated with an earlier execution of the customized procedure.

Option:
```
-type="procedure type"
```
Display all the Procedure Instances of a given type.

Output Columns:
Instance GUID, Procedure Type, Instance Name, Status

To find the type associated with procedures, use the get_procedures_type verb.

emcli get_instance_data_xml -instance="instance_guid"

Usage:
emcli get_instance_data_xml -instance="instance_guid"

Description:
Download Instance Data XML. EMCLI verb to download a RuntimeData xml using Instance GUID.

Option:
-instance is used to specify the instance GUID.

Example:
```
emcli get_instance_data_xml
-instance="16B15CB29C3F9E6CE040578C96093F61"
```

Output:
The Instance Data XML.

32.3.3 Step 3: Creating Properties File

The following sections describe the properties files for out-of-box procedures, customized procedures, and extending procedure execution.

32.3.3.1 Properties File for Out-Of-Box Procedures

If you are using an out-of-box RuntimeData template, a user identifies the variables in the RuntimeData template file that need to be replaced with values for the configuration properties for a given execution of the procedure. Of all the variables present in the Runtime Data templates, only some might be mandatory for running a given procedure. Once this is done you can create the properties file, which would contain name value pairs mentioning the variables and the values with which they would be replaced for generating the RuntimeData xml file.

For out-of-box procedures, a sample properties file can be found in Appendix B, "Sample Property Files for the Out-of-Box RuntimeData Templates". The corresponding RuntimeData Templates can be obtained from the zip file where this document is present.

Note that each sample properties files in the appendix mentioned above contains a section for mandatory variables which should be present in the properties file with relevant values to be substituted at run time. You can optionally provide values for other variables present in the templates.

32.3.3.2 Properties File for Customized Procedures

In case of customized procedures, the RuntimeData xml actually have values instead of placeholder variables for the configuration properties. You need to replace the old runtime values in the RuntimeData xml of the previous run with the new runtime values, which are relevant to the new run.

For this, you can have a properties file of the form:

<old_value>=<new_value>

For example, consider this snippet from the RuntimeData xml used to patch an Oracle Database:

```
…

<scalar value="dbtarget1" classname="java.lang.String" name="targetsToPatch"/>
<scalar value=" HostPrefNormal" classname="java.lang.String"
name="hostCredentialSet"/>
   <list classname="java.util.Vector" name="patchIDs">
   <scalar value="=%oracle_home%/EMStagedPatches " classname="java.lang.String"
name="stageDir"/>
<scalar value="false" classname="java.lang.String"
    name="isPatchset"/>
<scalar value="true" classname="java.lang.String"
    name="isNotPatchset"/>
<scalar value="defaultSqlScript" classname="java.lang.String" name="sqlScript"/>
...
```

To patch another database you need to create a properties file with
oldvalue=newvalue type of entries for at least the mandatory parameters (in case of
patching on the mandatory property is targetsToPatch). Hence the new properties file
would look something as below:

dbtarget1=dbtarget2

Since this approach would simply replace an old-string with a new-string, you might
run into issues if the old-string is substring in multiple strings in the DP runtime xml.
In that case the resulting runtime xml might be erroneous. To circumvent this issue, it
is strongly advised to format the properties file a proper fashion. The thumb-rule here
is: put the specifics before the generics. A fragment of a properties file in the form of
old-value=new-value pairs shown below, illustrates this point.

node1.test.com=node2example.com

node1=node2

node1,node2=node3,node4

Also, for specifying the passwords in the properties file please make sure you include
the following line in the properties file before mentioning any passwords.

```
oracle.sysman.pp.paf.impl.EncryptedString
=oracle.sysman.pp.paf.impl.UnencryptedString
```

After this you can mention the password as shown in the examples below:

1. For a password value to replace the placeholder variable in the template file

   ```
   oracle.sysman.pp.paf.impl.EncryptedString=oracle.sysman.pp.paf.impl.
   UnencryptedStringcrsasmrac_provisioning_USER_PASSWORD=mypassword
   ```

2. For a new password value to replace an older one in a RuntimeData xml

   ```
   oracle.sysman.pp.paf.impl.EncryptedString=oracle.sysman.pp.paf.impl.
   UnencryptedStringmyOLDpass=myNEWpassword
   ```

Note that `mypassword` and `myNewpassword` used in the above examples are clear
text passwords.

Note: The elements var_runOpatchUpgrade and var_
isUpgradeStepEnabled have been added to support Opatch upgrade.
The first element should be set to "true" to run the opatch upgrade
step. var_isUpgradeStepEnabled should be set to "true" if opatch is to
be upgraded, otherwise, it should be "false".

32.3.3.3 Properties File For Extending Procedure Execution

Properties file allows you to use the same RuntimeData xml for extending the use case. For example, you might have performed a successful procedure execution (partial cluster scale- up or scale down or patching) for a target and you want to extend it to a set of new targets.

You can do this by using a Properties file and replacing the parameter values in it (Refer to the Mandatory parameter section in Appendix B, "Sample Property Files for the Out-of-Box RuntimeData Templates". For example:

node1 = node2,node 3

Wherein *node1* is the target for which you had run the procedure previously and *node2* and *node3* are the targets to which you want to extend the procedure execution.

An exception to this rule is extending the patching procedures to a different set of targets, which would require you to have a properties file with the following mandatory parameter (Instead of the approach of <old-value>=<new-value>):

PA_VAR_targetsToPatch=Tgt2, Tgt3, Tgt4

Wherein *Tgt2*, *Tgt3* and *Tgt4* are the new targets to which you want to extend the procedure execution.

Refer to Queries to Acquire Data for Patching Runtime for the list of queries which can be used to acquire data for creating properties file.

32.3.3.4 Properties File For Applying Multiple Patches At Once

The new elements added in RuntimeData xml to support multiple patches are:

- patchesToBeApplied - Enter comma-separated list of patch IDs for this element.

  Example: <scalar value="patchesToBeApplied" classname="java.lang.String" name="patchListToApply"/>

- patchSourceForPatches - Enter patch source, which is either SWLIB or METALINK. Default source is SWLIB

  Example: <scalar value="patchSourceForPatches" classname="java.lang.String" name="patchListSource"/>

- patchPlatformForPatches - This is optional. Enter supported platforms for patchOptional. You must provide a valid platform ID. To get platform IDs, run the displayPlatformInfo.pl script in <EMCLI working directory>/scripts.

  Example: <scalar value="patchPlatformForPatches" classname="java.lang.String" name="patchPlatform"/>

- patchReleaseForPatches - Enter the release of the patchset. This is optional, but required for patchsets.

  Example: <scalar value="patchReleaseForPatches" classname="java.lang.String" name="patchRelease"/>

32.3.4 Step 4: Procedure Execution

First download and apply the Metalink patch - 5890474 on OMS 10.2.0.3 which will update the EMCLI procedure execution scripts, out-of-box templates, and properties files. This updates the *emcli/samples* and *emcli/scripts* directory in OMS Oracle home. After applying the patch on the OMS, download the procedure execution scripts, out-of-box templates, and properties files on the machine where the EMCLI client is

set up. The out-of-box templates and properties files for patching and provisioning are available in the respective directories under *OMSHOME/emcli/samples/*.

Once the RuntimeData template or RuntimeData xml and properties file are ready, then the procedure can be run using the following script.

Usage:

*perl executeDP.p*l
   -t <template>
   -p <properties file name>
   -g <DP GUID>
   [*-s <schedule>* in the format yyyy/MM/dd HH:mm]
   [*-z <time zone ID>*]
    [*-d <emcli directory path>*, mandatory if emcli executable is not in the current directory]

Template -- The name of the RuntimeData template for out-of-box procedures or location of the RuntimeData xml file downloaded after the execution of a customized procedure.

Properties file name -- The location of the properties file created for executing the procedure.

DP GUID -- The GUID of the procedure that needs to be run.

emcli Directory -- The directory which contains the EMCLI executable. If the current working directory contains the EMCLI executable, this parameter is optional.

Schedule -- The time when the Deployment Procedure would be scheduled to run. If not specified it defaults to running the Deployment Procedure immediately. The HH:MM is based on 24 hrs clock, for example, 22:30.

Time Zone ID -- The time zone to which the Deployment Procedure run is scheduled. If not specified it defaults to the time zone of the OMS.

Below is a sample code string executing RAC provisioning procedure for UNIX using out-of-box procedure:

```
perl executeDP.pl -t crsasmrac_gold_prov_template.xml -p
Properties.txt -g 31ABCFF2199BB77990B057AC4A442DAC -t 2007/02/03
10:00 -z Americas/New_York -d /oracle/prod/orahome/
```

The following parameter descriptions apply to the script:

crsasmrac_gold_prov_template.xml is the name of the out-of-box template.

Properties.txt is the properties file

31ABCFF2199BB77990B057AC4A442DAC is the GUID for the RAC provisioning procedure for UNIX

2007/02/03 10:00 is the date and time during which the Deployment Procedure is scheduled to run.

Americas/New_York is the Time Zone ID for which the time schedule is set.

/oracle/prod/orahome/ is the directory location for the EMCLI executables.

The properties file and the out-of-box template are located in the same directory as the script that was run.

32.3.4.1 Executing SIDB Patching for UNIX Using An Out-of-Box Procedure

Below is a sample code string executing single instance database patching for UNIX using an out-of-box procedure:

```
perl executeDP.pl patch_standalone_DB.xml Properties.txt
2EECED3592A0175FE040578CE808291F
```

The following parameter descriptions apply to the script:

patch_standalone_DB.xml is the name of the out-of-box template.

Properties.txt is the properties file.

2EECED3592A0175FE040578CE808291F is the GUID for the Single Instance Database patching procedure for UNIX.

The templates, properties file, and the EMCLI executables are located in the same directory as the script. Also, the Deployment Procedure is scheduled to run immediately in the time zone of the OMS.

32.4 Use Cases for EMCLI-based Provisioning and Patching

The following sections describe various use cases for EMCLI-based provisioning and patching procedures. You must first download and apply the Metalink patch - 5890474 on OMS 10.2.0.3 which will update the EMCLI procedure execution scripts, out-of-box templates, and properties files. This updates the *emcli/samples* and *emcli/scripts* directory in OMS oracle home. After applying the patch on OMS, download the procedure execution scripts, out-of-box templates, and properties files on the machine where the EMCLI client is setup. The out-of-box templates and properties files for patching and provisioning are available in the respective directories under *OMS HOME/emcli/samples/*.

Before using any Real Application Cluster (RAC) related procedure, ensure that the Management Agents on the nodes are cluster Management Agents.

32.4.1 Use Cases for CRS/ASM/RAC Provisioning Procedure

Use Case 1: User wants to use the EMCLI to provision a 2-node RAC using a Gold Image from the software library. He uses the out-of-box templates and properties file to perform this operation.

- User creates a Properties file with the mandatory configuration parameters required for RAC provisioning procedure and assigns appropriate values for the variables. Refer Sample Properties file with Mandatory parameters for out-of-box RAC provisioning procedure using Gold Image

- User finds the appropriate GUID for the RAC provisioning procedure. Refer to section "Out-of-box RuntimeData Templates For RAC Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box RAC provisioning procedure using Gold Image from Software Library.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID.

Use Case 2: User wants to use EMCLI to provision another 4 node RAC using the same out-of-box templates and properties file as used in Use case 1 to perform this operation.

■ User takes the properties file from the previous use case and makes the necessary changes for the mandatory parameters to provision a 4-node RAC.

■ User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID. Sample usage of the script is shown below.

Use Case 3: User wants to use the EMCLI to provision a 2-node RAC using a Reference Installation. He uses the out-of-box templates and properties file to perform this operation.

■ User creates a Properties file with the mandatory configuration parameters required for RAC provisioning procedure and assigns appropriate values for the variables. Refer Sample Properties file with Mandatory parameters for out-of-box RAC provisioning procedure using Reference Installation.

■ User finds the appropriate GUID for the RAC provisioning procedure. Refer to section "Out-of-box RuntimeData Templates For RAC Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box RAC provisioning procedure using Reference Installation.

■ User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID.

Use Case 4: User customizes and tests the out-of-box RAC provisioning procedure using reference host. He wants to use the EMCLI to provision a 2-node RAC. He uses the runtime data xml of the trial runs of his customized procedure and properties file to perform this operation on a similar 2-node RAC.

■ Locates the instance GUID of the previous trial run.

■ User downloads the Runtime data xml for the previous execution of the procedure.

■ User identifies the parameters in the Runtime data xml that need to be substituted with new values. He then creates a properties file with name-value pairs like *<old-value>=<new-value>* for carrying out the necessary runtime substitutions. This properties file should at least contain the substitution rule for the values corresponding to the mandatory parameters mentioned in Sample Properties file with Mandatory parameters for out-of-box RAC provisioning procedure using Reference Installation in addition to the other values that he might want to substitute.

■ User finds the GUID for the customized RAC provisioning procedure. Refer to section Finding Procedure GUID.

■ User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the downloaded Runtime data xml, and procedure GUID.

Use Case 5: User customizes and tests the out-of-box RAC provisioning procedure. He wants to use EMCLI to provision a N-node RAC. He uses the runtime data xml of a trial run of his customized procedure and properties file to perform this operation on a M-node cluster (where M>N).

■ Locates the instance GUID of the previous trial run of the customized procedure.

■ User downloads the Runtime data xml for the previous execution of the procedure.

- User identifies the parameters in the Runtime data xml that need to be substituted with new values. He then creates a properties file with name-value pairs like *<old-value>=<new-value>* for carrying out the necessary runtime substitutions. This properties file should at least contain the substitution rule for the values corresponding to the mandatory parameters mentioned in Sample Properties file with Mandatory parameters for out-of-box Oracle RAC provisioning procedure using Gold Image, in addition to the other values that he might want to substitute.

- User finds the GUID for the customized Oracle RAC provisioning procedure. Refer to section Finding Procedure GUID.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the downloaded Runtime data xml and procedure GUID.

32.4.2 Use Cases for Extend Cluster Procedure

Use Case 1: User wants to use the EMCLI to extend a 2-node Oracle RAC to 4-node cluster. He uses the out-of-box templates and properties file to perform this operation.

- User creates a Properties file with at least the mandatory parameters required for cluster Extension procedure and assigns appropriate values for the variables. Refer to Sample Properties file with Mandatory parameters for out-of-box Cluster Extend procedure.

- User finds the appropriate GUID for the RAC provisioning procedure. Refer to section "Out-of-box RuntimeData Templates For RAC Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box Cluster Extend procedure.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, name of the out-of-box template and procedure GUID.

32.4.3 Use Cases For RAC Delete/Descale Procedure

Use Case 1: User wants to use EMCLI to delete the 2-node RAC cluster. He uses the out-of-box templates and properties file to perform this operation.

- User creates a Properties file with at least the mandatory parameters required for RAC Delete procedure and assigns appropriate values for the variables. Refer Sample Properties file with Mandatory parameters for out-of-box Cluster Delete procedure.

- User finds the appropriate GUID for the RAC Delete procedure. Refer to section "Out-of-box RuntimeData Templates For RAC Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box Scale Down/Delete RAC procedure. Note that template used for Cluster Scale down and Cluster Delete use cases differ.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, name of the out-of-box template and procedure GUID.

Use Case 2: User wants to use the EMCLI to descale a 2-node RAC cluster. He uses the out-of-box templates and properties file to perform this operation.

- User creates a Properties file with at least the mandatory parameters required for RAC Scale Down procedure and assigns appropriate values for the variables.

Refer Sample Properties file with Mandatory parameters for out-of-box Cluster Descale procedure.

- User finds the appropriate GUID for the RAC Descale procedure. Refer to section "Out-of-box RuntimeData Templates For RAC Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box Scale Down/Delete RAC procedure. Note that template used for Cluster Scale down and Cluster Delete use cases are different.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, name of the out-of-box template and procedure GUID.

32.4.4 Use Cases for Patching

Use Case 1: User wants to use the out-of-box Database Patching procedure to apply a one-off patch to a database. He uses the out-of-box templates and properties file to perform this operation.

- User creates a Properties file with the mandatory configuration parameters required for patching the database with a particular one-off and assigns appropriate values for these variables. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

- User finds the appropriate GUID for the Patch provisioning procedure. Refer to section "Out-of-box RuntimeData Templates For Patching Procedures" of Appendix A, "Out-Of-Box RuntimeData Templates" for GUID, procedure type, and template name information for the out-of-box Patch Oracle Database procedure.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID.

Use Case 2: User wants to use the Database Patching procedure to apply multiple one-off patches to multiple databases. He uses the out-of-box templates and properties file created in use User Case 1 above to perform this operation.

- User creates a Properties file with the mandatory configuration parameters required for patching the database with a particular one-off and assigns appropriate values for these variables. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID.

Use Case 3: User wants to use the Database Patching procedure to apply a patchset to a set of databases. He uses the out-of-box templates and properties file created in use Use Case 1 above to perform this operation.

- User creates a Properties file with the mandatory configuration parameters required for patching the database with a particular one-off and assigns appropriate values for these variables. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

- User find the appropriate GUID for the Patch provisioning procedure.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the out-of-box template and procedure GUID.

Use Case 4: User customizes and tests the out-of-box Oracle Clusterware Patching procedure. He wants to use the EMCLI to patch a 2-node cluster. He uses the runtime data xml of the trial runs of his customized procedure and properties file to perform this operation on a similar 2-node cluster.

- User creates a Properties file with the mandatory configuration parameters required for patching the database with a particular one-off and assigns appropriate values for these variables. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

- User finds the GUID for the customized patching procedure.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the run time data xml and procedure GUID.

Use Case 5: User customizes and tests the out-of-box Oracle Clusterware Patching procedure. He wants to use the EMCLI to patch a 2-node cluster. He uses the runtime data xml of the trial runs of his customized procedure and properties file to perform this operation on a similar N-node cluster.

- User creates a Properties file with the mandatory configuration parameters required for patching the database with a particular one-off and assigns appropriate values for these variables. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

  For example, use the mandatory variable of the Properties file for scaling up to multiple Targets as seen here:

  PA_VAR_targetsToPatch=NewTarget1, NewTarget2, NewTarget3…

- User finds the GUID for the customized patching procedure.

- User submits the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the runtime data xml and procedure GUID.

32.4.5 Limitations

There is a limitation to consider when using patching Deployment Procedures through EMCLI.

Out-of -box templates are not available for patching Deployment Procedures like Application Server patching, Real Application Cluster -All nodes and pre-requisite checkers for Database, Real Application Cluster (RAC), Automatic Storage Management (ASM) and Clusterware. To use EMCLI for these procedures:

- Run the procedure once through the UI.

- Export the run time data xml

  ```
  emcli get_instance_data_xml -instance="instance_guid"
  ```

  Where, instance_guid is of the Deployment Procedure. (Refer to step 2-Obtaining Runtime Data Template and Data xml.)

- Create a properties file with the mandatory configuration parameter required for patching or running pre-requisite checker and assign appropriate values for the other changes. List of the mandatory values can be found from the Sample Properties file with Mandatory parameters for all the patching procedures.

  For example, use the mandatory variable of the Properties file for scaling up to multiple Targets

  PA_VAR_targetsT=NewTarget1, NewTarget2, NewTarget3…

- Submit the procedure for execution by invoking the executeDP.pl script and providing the location details of the properties file, location of the runtime data xml and procedure GUID.

32.5 Setting Up Preferred Credentials for Targets

The EMCLI execution looks for the credentials for the Targets under the Enterprise Manager with the OMS user executing the procedures. The preset credentials is looked up for the Targets under patching or for the ones used as a reference during provisioning procedures.

The credentials can be stored while doing any patching or provisioning operations through the Enterprise Manager user interface in the 'Credentials' section of the procedure run. If not, you can set up the credentials either through the Enterprise Manager OMS explicitly or through the use of EMCLI commands.

32.5.1 Setting Credentials from the Oracle Enterprise Manager User Interface

You can set the credentials for targets through the Oracle Enterprise Manager user interface by following these steps:

1. Log in to Oracle Enterprise Manger.

2. Access the link "Preferences" on the top right corner of the page.

3. Click on "Preferred Credentials" link in the Options section of the page.

4. Setup 'Normal' or 'Preferred Credentials' from this page for the Target type. (Example: Database Instance, Cluster Database or Cluster).

32.5.2 Setting Credentials through EMCLI

You can set the credentials for targets through EMCLI command line interface using the following code sequence:

```
set_credential
            -target_type="ttype"
           [-target_name="tname"]
            -credential_set="cred_set"
           [-user="user"]
            -columns="col1:newval1;col2:newval2;..."
           [-input_file="tag1:file_path1;tag2:file_path2;..."]
           [-oracle_homes="home1;home2"]
```

The following list describes the options used in the EMCLI code:

- **target_type** - Type of target. Must be "host" in case "-oracle_homes" parameter is specified.

- **target_name** - Name of target. Omit this argument to set enterprise preferred credentials. Must be hostname in case "-oracle_homes" parameter is specified.

- **user** - Enterprise Manager user whose credentials are affected. If omitted, the current user's credentials are affected.

- **columns** - The name and new value of the column(s) to set. Every column of the credential set must be specified. Alternatively, a tag from the -input_file argument may be used so that the credential values are not seen on the command line. This argument may be specified more than once.

- **input_file** - Path of file that has -columns argument(s). This option is used to hide passwords. Each path must be accompanied by a tag, which is referenced in the -columns argument. This argument may be specified more than once.

- **oracle_homes** - Name of oracle homes on the target host. Credentials will be added/updated for all specified homes.

The list of columns and the credential sets they belong to is included in the metadata file for each target type. This and other credential information is in the <CredentialInfo> section of the metadata.

The following is an example of the sequence:

```
emcli set_credential
              -target_type=host
              -target_name=host.us.oracle.com
              -credential_set=OHCreds
              -user=admin1
              -column="OHUsername:joe;OHPassword:newPass"
              -oracle_homes="database1;mydb"
```

For more details on EMCLI, refer to the verb reference section of Oracle® Enterprise Manager Command Line Interface 10g Release 5 (10.2.0.5.0). This document is available at:

http://download.oracle.com/docs/cd/B16240_01/doc/nav/portal_booklist.htm

32.5.3 Clearing Credentials through EMCLI

You can clear preferred or monitoring credentials for given users through the EMCLI command line interface using the following code sequence:

```
emcli clear_credential
        -target_type="ttype"
        [-target_name="tname"]
        -credential_set="cred_set"
        [-user="user"]
        [-oracle_homes="home1;home2"]
```

The following list describes the options used in the EMCLI code:

- **target_type** - Type of target. Must be "host" in case "-oracle_homes" parameter is specified.

- **target_name** - Name of target. Omit this argument to set enterprise preferred credentials. Must be hostname in case "-oracle_homes" parameter is specified.

- credential_set - Credential set affected. This value is ignored for monitoring credentials.

- **user** - Enterprise Manager user whose credentials are affected. If omitted, the current user's credentials are affected.

- **oracle_homes** - Name of oracle homes on the target host. Credentials will be cleared for all specified homes.

Example 1:

```
 emcli clear_credential
        -target_type=oracle_database
        -target_name=myDB
        -credential_set=DBCredsNormal
        -user=admin1
```

Example 2:

```
emcli clear_credential
        -target_type=oracle_database
        -credential_set=DBCredsNormal
        -user=admin1
```

32.6 Queries to Acquire Data for Patching Runtime

Use the following queries to acquire data for patching runtime:

- Use the following query to acquire an Instance name from a host:

  select target_name, target_type, oracle_home from em$ECM_TARGETS_VIEW where host = '<host name>';

- Get the instance name for a given host:

  select target_name, target_type, oracle_home from em$ECM_TARGETS_VIEW where host = '<host name>';

- Get the instances of a CRS given the name of the CRS:

  select assoc_target_name, crs_instance from sysman.mgmt$target_associations where assoc_def_name='contains' and source_target_name='<crs_name>' and source_target_ type='cluster'

- Get all CRS and its instances:

  select source_target_name crs_name, assoc_target_name, crs_instance from sysman.mgmt$target_associations   where assoc_def_name='contains' and source_target_ type='cluster' order by source_target_name

- Get instances of a RAC cluster given the name of the RAC cluster:

  select assoc_target_name rac_instance from sysman.mgmt$target_associations where assoc_def_name='contains' and source_target_name='<rac_name>' and source_target_ type='rac_database'

- Get all RAC clusters and their instances:

  select source_target_name rac_name, assoc_target_name rac_instance from sysman.mgmt$target_associations where assoc_def_name='contains' and source_target_ type='rac_database' order by source_target_name

32.7 EMCLI Verbs for Deployment Procedures

Following are the EMCLI verbs related to deployment procedures:

32.7.1 get_instance_status

Description

Display Procedure Instance status identified by the GUID on the command line.

Usage

```
emcli get_instance_status  -guid={guid number} [-xml [-details] [-showJobOutput
[-tailLength={Last N Characters}]]]
```

Option

  -guid={guid number}

Mandotory option to display the details of a Procedure Instance identified by the guid number.

    {guid number} can be found using the emcli get_instances.

  -xml

    Show the complete status of each of the steps in XML format

  -details

    Optional option to display more details to the output of the command. Can only be used in conjunction with the -xml option.

  -showJobOutput

    Show output or error of the Job Execution steps.

  -tailLength={Last N Characters}

    Optional argument to limit the number of characters in the job step output or error. Can only be used with the -showJobOutput option.

    {Last N Characters} is a positive non-zero number till which the characters would be picked from the end of the job step output

    Max permissible characters to dump as part of the above is set by system.

    If this option is not provided, then max permissible characters would be dumped.

Example

```
emcli get_instance_status -guid=16B15CB29C3F9E6CE040578C96093F61 -xml
-showJobOutput -tailLength=1024
    emcli get_instance_status -guid=16B15CB29C3F9E6CE040578C96093F61 -xml -details
-showJobOutput
    emcli get_instance_status -guid=16B15CB29C3F9E6CE040578C96093F61 -xml -details
    emcli get_instance_status -guid=16B15CB29C3F9E6CE040578C96093F61
```

32.7.2 get_procedure_types

Description

Get the list of all Deployment Procedure types.

Usage

```
emcli get_procedure_types
```

Example

```
emcli delete_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

Output Column

Procedure Type

32.7.3 get_procedure_xml

Description

Get the Deployment Procedure XML file. XML will be printed on standard output.

Usage

```
emcli get_procedure_xml -procedure={procedure guid}
```

Option

-procedure

    The procedure GUID.

Example

```
emcli get_procedure_xml -procedure=16B15CB29C3F9E6CE040578C96093F61 > proc.xml
```

Output

The Deployment Procedure XML

32.7.4 get_procedures

Description

Get a list of Deployment Procedures.

Usage

```
emcli get_procedures [-type={procedure type}]
```

Option

-type={procedure type}

    display all the Deployment Procedure of type {procedure type}.

Output Columns

GUID, Procedure Type, Name, Version, Created By.

32.8 EMCLI Verbs for Deployment Instances

Following are the emcli verbs for deployment instances.

32.8.1 confirm_instance

Description

Confirm manual step.

Usage

```
emcli confirm_instance -instance={instance_guid} -stateguid={state_guids}
```

Option

-instance

    The instance GUID.

-stateguid

The separated list of State GUIDs.

Example

```
emcli confirm_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid=51F762417C4943DEE040578C4E087168
    emcli confirm_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid='51F762417C4943DEE040578C4E087168,51F762417C4944DEE040578C4E087168'
```

32.8.2 delete_instance

Description

Delete a stopped or completed Deployment Instance.

Usage

```
emcli confirm_instance -instance={instance_guid} -stateguid={state_guids}
```

Option

-instance

The instance GUID.

Example

```
emcli delete_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

32.8.3 get_instance_data_xml

Description

Download Instance Data XML.

Usage

```
emcli get_instance_data_xml -instance={instance_guid}
```

Option

-instance

The instance GUID.

Example

```
emcli get_instance_data_xml -instance=16B15CB29C3F9E6CE040578C96093F61 > data.xml
```

Output

The Instance Data XML.

32.8.4 get_instances

Description

Display list of Procedure Instances.

Usage

```
emcli get_instances -type={procedure type}
```

Option

-type={procedure type}

display all the Procedure Instances of type {procedure type}.

Example

```
emcli delete_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

Output Columns

GUID, Procedure Type, Instance Name, Status

32.8.5 get_retry_arguments

Description

Get arguments of failed steps which can be retried.

Usage

```
emcli get_retry_arguments -instance={instance_guid} [-stateguid={state_guid}]
```

Option

-instance

The instance GUID.

-stateguid

The State GUID.

Example

```
emcli get_retry_arguments -instance=16B15CB29C3F9E6CE040578C96093F61
    emcli get_retry_arguments -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid=51F762417C4943DEE040578C4E087168
```

32.8.6 ignore_instance

Description

Ignore failed step.

Usage

```
emcli ignore_instance -instance={instance_guid} -stateguid={state_guid}
```

Option

-instance

The instance GUID.

-stateguid

The State GUID.

Example

```
emcli ignore_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid=51F762417C4943DEE040578C4E087168
    emcli ignore_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid='51F762417C4943DEE040578C4E087168,51F762417C4944DEE040578C4E087168'
```

32.8.7 reschedule_instance

Description

Reschedule a submitted Procedure Instance

Note that only scheduled instances can be rescheduled

Usage

```
emcli reshedule_instance
        -instance={instance guid}
        -schedule=
            start_time:yyyy/MM/dd HH:mm;
            tz:{java timezone ID};
            grace_period:xxx
```

Option

-instance

The GUID of the instance to execute.

-schedule

The schedule for the Procedure Instance.

start_time - when the procedure should start

tz   - the timezone ID (optional)

grace_period - grace period in minutes(optional)

Example

```
emcli reschedule_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-schedule="start_time:2006/6/21 21:23;tz:America/New_York;grace_period:60
```

32.8.8 resume_instance

Description

Resume a suspended Deployment Instance.

Usage

```
emcli resume_instance -instance={instance_guid}
```

Option

-instance

The instance GUID

Example

```
emcli resume_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

32.8.9 retry_instance

Description
Retry failed instance or failed step.

Usage
```
emcli retry_instance -instance={instance_guid} [-stateguid={state_guid}]
```

Option
-instance

> The instance GUID

 -stateguid

> The State GUID.

Example
```
emcli retry_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid=51F762417C4943DEE040578C4E087168
    emcli retry_instance -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid='51F762417C4943DEE040578C4E087168,51F762417C4944DEE040578C4E087168'
```

32.8.10 set_instance_jobgrants

Description
Set the user access privilege for jobs of the Deployment Instance.

Usage
```
emcli set_instance_jobgrants -instance_guid={instance guid}
-grants={user:privilege}
```

Option
-instance

> The instance GUID

-grants

> String of user:privilege pairs each separated by ; .
>
> > where user = em user name,
> >
> > > privilege = VIEW_JOB or FULL_JOB.

Example
```
emcli set_instance_jobgrants -instance_guid=16B15CB29C3F9E6CE040578C96093F61
-grants="user1:VIEW_JOB;user2:FULL_JOB"
```

32.8.11 start_paf_daemon

Description
Start Deployment Procedure Manager Daemon.

Usage

```
emcli start_paf_daemon -interval={number in minutes}
```

Option

-interval

Number in minutes that Deployment Procedure Manager Daemon should wait between each run.

32.8.12 status_paf_daemon

Description

Get Deployment Procedure Manager Daemon status.

Usage

```
emcli status_paf_daemon
```

32.8.13 stop_instance

Description

Stop a scheduled, failed, or running Deployment Instance.

Usage

```
emcli stop_instance -instance={instance_guid}
```

Option

-instance

The instance GUID

Example

```
emcli stop_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

32.8.14 stop_paf_daemon

Description

Stop Deployment Procedure Manager Daemon.

Usage

```
emcli stop_paf_daemon
```

32.8.15 submit_procedure

Description

Submit a Deployment Procedure.

Usage

```
emcli submit_procedure
        -procedure="guid of the procedure"
        -input_file="data:{file_path}"
```

```
                    [-instance_name="name for the procedure instance"]
                    [-grants="users and their correspoding access levels"]
                    [-schedule=
                        start_time:yyyy/MM/dd HH:mm;
                        tz:{java timezone ID};
                        grace_period:xxx];
                    ]
```

Option

-procedure

    The GUID of the procedure to execute.

  -input_file=data:file_path

    The input data for the Deployment Procedure.

    The file_path should point to a file containing the data XML file.

  -instance_name

    The name of the procedure instance (optional)

  -grants

    Users and their corresponding access levels (optional)

    String of user:privilege pairs each separated by ; .

      where user = em user name,

          privilege = VIEW_JOB or FULL_JOB.

  -schedule

    The schedule for the Deployment Procedure (optional).

    If not specified, the procedure will be executed immediately.

      start_time   - when the procedure should start

      tz   - the timezone ID (optional)

      grace_period - grace period in minutes(optional)

Example

```
emcli submit_procedure -input_file=data:data.xml
-procedure=16B15CB29C3F9E6CE040578C96093F61 -grants="user1:VIEW_JOB;user2:FULL_
JOB" -schedule="start_time:2006/6/21 21:23;tz:America/New_York;grace_period:60"
-instance_name="MyProcedureInstance_0001"
```

Output

Instance GUID

32.8.16 suspend_instance

Description

Suspend a running Deployment Instance.

Usage

```
emcli suspend_instance -instance={instance_guid}
```

Option

-instance

The instance GUID

Example

```
emcli suspend_instance -instance=16B15CB29C3F9E6CE040578C96093F61
```

32.8.17 update_and_retry_step

Description

Update arguments of the failed step and retry it.

Usage

```
emcli update_and_retry_step -instance={instance_guid} -stateguid={state_guid}
[-args="command1:value1;command2:value2;..."]
```

Option

-instance

The instance GUID

-stateguid

The State GUID.

-args

Arguments, which will be updated.

Example

```
emcli update_and_retry_step -instance=16B15CB29C3F9E6CE040578C96093F61
-stateguid=51F762417C4943DEE040578C4E087168 -args="command:ls"
```

Part IX

Additional Information

This part contains additional information related to provisioning and patching features offered by Enterprise Manager Grid Control.

This part contains the following chapters:

- Appendix A, "Troubleshooting Issues"

- Appendix B, "Using emctl partool Utility"

- Appendix C, "Using Oracle Software Library"

- Appendix D, "Creating Templates for Oracle Database"

- Appendix E, "Out-Of-Box RuntimeData Templates"

- Appendix F, "Sample Property Files for the Out-of-Box RuntimeData Templates"

- Appendix G, "PXE Booting and Kickstart Technology"

- Appendix H, "About Patching Automation Reports"

- Appendix I, "Converting Standalone Oracle Management Agent to Cluster Agent"

A

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 Provisioning and Patching Issues
- Frequently Asked Questions

A.1 Troubleshooting Provisioning and Patching Issues

This section provides troubleshooting tips for some common provisioning and patching issues. In particular, this section covers the following:

- Troubleshooting Provisioning Issues
- Troubleshooting Virtualization Management Issues
- Troubleshooting Linux Provisioning Issues
- Troubleshooting Linux Patching Issues

A.1.1 Troubleshooting Provisioning Issues

To troubleshoot provisioning issues, see My Oracle Support note 816656.1. You can access My Oracle Support at:

https://support.oracle.com/

A.1.2 Troubleshooting Patching Issues

This section describes how you can troubleshoot patching issues. In particular, this section covers the following:

- Using Troubleshooting Options Offered by Deployment Procedure
- Collecting Information from Procedure Logs
- Instance XML of the Procedure Executed
- Raising a Service Request

A.1.2.1 Using Troubleshooting Options Offered by Deployment Procedure

When a patching Deployment Procedure fails at a step, you can choose to do one of the following:

- **Ignore:** Ignore the failure and proceed with the step.

- **Retry:** Retry the step. For example, locked processes issue.

- **Update and Retry:** Update and then retry the step in case of any errors in the runtime values.

Figure A–1 shows the troubleshooting option available as part of the procedure steps.

Figure A–1   Troubleshooting options as part of the procedure steps

Additionally, during the **Update and Retry** mode, you can enable the debug option to get more information on the failures. See Figure A–2.

Figure A–2   Getting More Debug Information
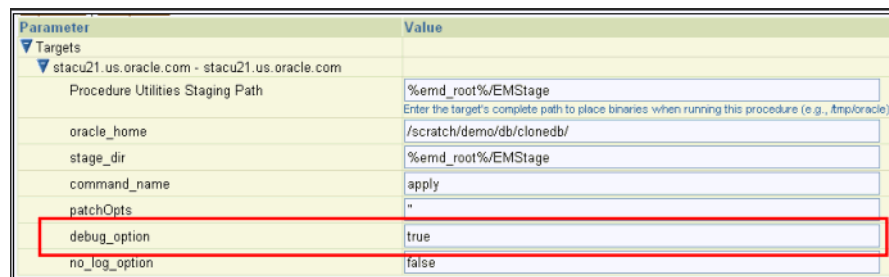
> **Note:**   This option is a part of the patching Deployment Procedures only.

A.1.2.2 Collecting Information from Procedure Logs

You can collect information about the failure from the procedure logs available in the Oracle home of the OMS and the Oracle home of the Management Agent.

The following are the OMS-related procedure logs:

- Generic Enterprise Manager trace files:

  `$<ORACLE_HOME>/sysman/log/emoms.trc`

- Provision Advisory Framework logs:

  `$<ORACLE_HOME>/ sysman/log/pafLogs/`

For specific Deployment Procedure instance, the logs are in the following file:

`$<ORACLE_HOME>/sysman/log/pafLogs/<name>_<instance_guid>.log`

The following are the Management Agent-related procedure logs:

`$<ORACLE_HOME>/sysma/logs/emagent.nohup`

`$<ORACLE_HOME>/sysma/logs/emagent.trc`

Optionally, to capture more details, you can make the logging to finer levels. Follow these steps to reset the log level and capture the logs mentioned above. (Note: It is advised to archive the old logs and have a fresh run after resetting the log level to capture the fresh logs.)

- On the OMS side:

  1. Open the following file available in the Oracle home of the OMS:

     `$<ORACLE_HOME>/sysman/config/emomslogging.properties`

  2. Set the `@ log4j.rootCategory=` parameter to `DEBUG`.

  3. Restart the OMS by running the following commands:

     `$<ORACLE_HOME>/bin/emctl stop oms`

     `$<ORACLE_HOME>/bin/emctl start oms`

- On the Management Agent side:

  1. Open the following file available in the Oracle home of the Management Agent:

     `$<ORACLE_HOME>/sysman/config/emd.properties`

  2. Make the following settings:

     tracelevel.Dispatcherr=DEBUG

     tracelevel.command=DEBUG

  3. Reload the Management Agent by running the following command:

     `$<ORACLE_HOME>/bin/emctl reload agent`

The above settings are to be set only when you want additional details and when the logs do not have sufficient information to debug the issue. Make sure to set the debug level back to the original levels after reproducing the issue.

A.1.2.3 Instance XML of the Procedure Executed

Instance XML provides with insight into all the values embedded with the procedures from the selected targets and user inputs. To generate an instance XML, you need to set up EMCLI either on the OMS or on the targets as SYSMAN or as an Enterprise Manager user who executed the procedure.

For more information, see *Oracle Enterprise Manager Command Line Interface* available at:

http://www.oracle.com/technology/documentation/oem.html

Run the following commands, where `instance_guid` is derived from the URL of the failed procedure (see Figure A–3):

`emcli get_instance_data_xml -instance={instance_guid}`

```
emcli get_instance_data_xml
-instance=16B15CB29C3F9E6CE040578C96093F61 > data.xml
```

Figure A–3 Instance GUID for the Procedure Instance Run
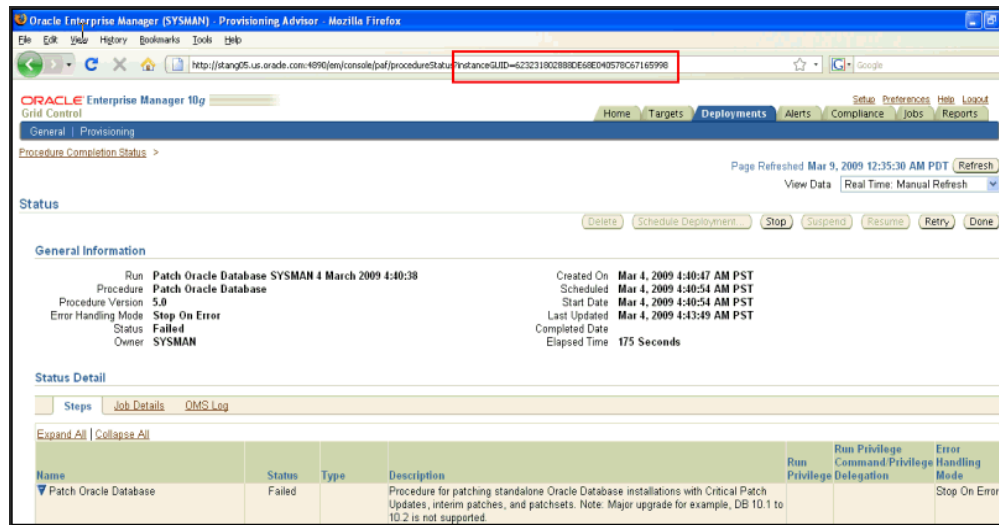
A.1.2.4 Raising a Service Request

Associate the following with a service request (SR) to Oracle Support to enable speedy resolution:

- Screenshots of the failed procedure and step

- Log from the failed step UI

- Tar/zip of the logs

- Optionally, instance XML of the run

A.1.3 Troubleshooting Virtualization Management Issues

How do I view Enterprise Management agent errors?

- Check the Alerts from the Virtual Server Home Page where the error occurs.

- Check the following EM Agent logs for errors. This will be the EM Agent monitoring the Virtual Server on which we have observed a collection error.

```
<AGENT_STATE>/sysman/log/sysman/log/emagent.log
<AGENT_STATE>/sysman/log/sysman/log/emagent.nohup
<AGENT_STATE>/sysman/log/sysman/log/emagent.trc
<AGENT_STATE>/sysman/log/sysman/log/emagentfetchlet.log
<AGENT_STATE>/sysman/log/sysman/log/emagentfetchlet.trc
```

- Enhance the logging level at the Agent to Debug, reproduce the issue, and take a look at the above-mentioned files for errors. Do remember to restore the logging level to what it was after getting the required data.

  Update the following properties in <AGENT_STATE>/sysman/config/emagentlogging.properties file to the below-mentioned values:

  ```
  log4j.rootCategory=DEBUG, emagentlogAppender, emagenttrcAppender
  ```

  Also update the following properties in the <AGENT_STATE>/sysman/config/emd.properties file to the below-mentioned values:

```
tracelevel.command=DEBUG
tracelevel.Dispatcher=DEBUG
```

After making these changes, the EM Agent will need to be bounced as follows:

```
<AGENT_HOME>/bin/emctl reload agent
```

How do I view OMS errors?

To enable enhanced logging at the OMS update the following properties in OMS_ ORACLE_HOME/sysman/config/emomslogging.properties on each OMS with the following and bounce each OMS after the update:

```
log4j.category.em.ovm=DEBUG, emovmlogAppender
log4j.appender.emovmlogAppender=org.apache.log4j.RollingFileAppender
log4j.appender.emovmlogAppender.File=/private/rahgupta/app/oracle/product/10.2.0/E
MHomes/oms10g/sysman/log/emovm.log
log4j.appender.emovmlogAppender.Append=true
log4j.appender.emovmlogAppender.MaxFileSize =5000000
log4j.appender.emovmlogAppender.MaxBackupIndex=100
log4j.appender.emovmlogAppender.layout=org.apache.log4j.PatternLayout
log4j.appender.emovmlogAppender.layout.ConversionPattern=%d [%t] %-5p %c{2} %M.%L
- %m\n
```

From each OMS the following logs need to be looked into:

```
$ORACLE_HOME/sysman/log/emoms.log
$ORACLE_HOME/sysman/log/emoms.trc
$ORACLE_HOME/sysman/log/emovm.log*
$ORACLE_HOME/opmn/logs/OC4J~OC4J_EM~default_island~1
```

It is possible that the log files get rolled over, particularly if reproducing the issue takes a while. In that case increase the following parameter in the above-mentioned file to a larger number:

```
LogFileMaxRolls=50
```

How do I view Virtual Server errors?

View the following logs for virtual server errors:

```
/var/log/ovs-agent/ovs_operation.log
/var/log/ovs-agent/ovs_query.log
/var/log/ovs-agent/ovs_root.log
```

A.1.4 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/Boot Server machine or its not uploading data to the OMS. Refer to the Agent Deployment Best Practices document available at http://www.oracle.com/ocom/technology/products/oem/pdf/10gr2_agent_ deploy_bp.pdf for troubleshooting information and known issues.

"Cannot create under the software library, please contact your administrator" error comes up while creating default image.

This may happen because the Software Library is not configured. Refer to Section 3.1.2.1, "Configuring Oracle Software Library". If it is configured and the error still occurs, then check if the library directory is accessible from the OMS server. The

software library location should be accessible from the OMS and should have enough space.

Default Image or OS component creation fails with a '404' status while copying the header.info.

This may happen because of the RPM Repository URL being incorrect or because the RPM Repository has not been configured properly. Please refer to Section 3.7.5, "Setting Up RPM Repository" and Section 3.7.11, "Configuring RPM Repository".

Default Image or OS component creation fails with an error message "Following RPMs are not found in Repository".

Reference machine has additional RPMs that are not available in the RPM repository. One of the following alternatives can be selected:

- Choose a different reference machine.

- Add extra RPMs in RPM repository (Make sure the RPMs are added in the repository as per Redhat specifications. Header files should be updated with the new RPM).

- If the RPMs are not required on the provisioned machine, add them to the Deny RPM list in the OS Component creation user interface.

Default Image or OS component creation job is suspended.

This may happen if the Preferred Credentials are not set for Reference machine. Please set the Preferred Credentials.

Default Image or OS component creation fails "Sudo error".

This may happen if:

- Sudo is not installed on the reference machine

- The user mentioned for the Reference Machine is not in the sudoers list. Edit /etc/sudoers file and add the user in the list.

Default Image staging fails with "Cannot stage default image".

This happens if there is already a default image staged for the same IP address range as the current one. Remove the staged image by clicking on the Remove button in the Default Image section on the Administration tab.

Default Image or OS component staging/provision fails with "Sudo error".

This may happen if:

- Sudo is not installed on the reference machine.

- The user mentioned for the Reference Machine is not in the sudoers list. Edit /etc/sudoers file and add the user in the list.

Default Image or OS component staging/provision job is suspended.

This may happen if Preferred Credentials are not set for Reference machine. Please set the Preferred Credentials.

Default Image or OS component staging/provision job fails "with directory permission error".

This error happens because of insufficient user privileges on the stage server machine. STAGE_TOP_LEVEL_DIRECTORY should have the write permission for the Stage

user. In case of NAS, the NAS directory should be mounted on the staging server. If the error is while writing to the '/tftpboot directory then it has to have the write permission for the Stage user.

Bare metal machine is not coming up since it cannot locate the Boot file.

Verify the dhcp settings (/etc/dhcpd.conf) & tftp settings for the target machine. Check whether the services (dhcpd, xinted, portmap) are running. Make the necessary setting changes if required or start the required services if they are down.

Even though the environment is correctly setup, 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.

Provisioning Default image on the bare metal box fails with "reverse name lookup failed" error.

Verify that the DNS has the entry for the IP address and the host name.

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 OS?

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 'no' (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 but following steps below would fix the problem:

■ On the Reference Machine modify 'ethtool -s eth0 duplex half' entry to 'ethtool -s eth0 duplex full' in the kickstart file.

■ Once done recreate the Default Image.

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/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 DHCP is configured to deliver correct DNS information to the target machine. If the problem persists, specify the IP address of the stage/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 EM 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.

Component creation fails with "Cannot create under the software library, please contact your administrator".

This may happen if the Software Library is not configured with the Enterprise Manager. Refer Section 3.1.2.1, "Configuring Oracle Software Library". Create the components once the software library is configured.

"Cannot create under the Software Library, please contact your administrator" error comes up while creating default image.

This may happen because the Software Library is not configured. Refer to Section 3.1.2.1, "Configuring Oracle Software Library". If it is configured and the error still occurs, then check if the library directory is accessible from the OMS server. The software library location should be accessible from the OMS and should have enough space.

Target does not appear under the hardware tab or while selecting hardware during assignment creation.

When you have multiple Oracle Management Services managed by a load balancer, you must configure a common directory that can be used by the Management Agents to upload data. To configure a common directory for this purpose, run the following command:

```
emctl config oms loader -shared yes -dir <directory>
```

With 64-bit OS provisioning, agent is not installed.

During OS provisioning, specify the full path of the agent RPM in the Advanced Properties page.

A.1.5 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?

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.

I see a UI error message saying "Package list is too long". How do I fix it?

Unselect some of the selected packages. The UI error message tells you from which package to unselect.

A.2 Frequently Asked Questions

This section provides answers to some frequently asked questions. In particular, this section covers the following:

- Frequently Asked Questions on Setting Up the Provisioning Environment
- Frequently Asked Questions on Directives and Components

■ Frequently Asked Questions on Network Profiles and Images

A.2.1 Frequently Asked Questions on Setting Up the Provisioning Environment

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. Refer to Appendix G, "PXE Booting and Kickstart Technology" for more information.

Can my Boot server and Stage server be located on different machine?

It is recommended that the Boot server and Stage server are co-located on the same physical machine. If this is not done, then one has to ensure that the network boot directory (say /tftpboot/pxelinux0.cfg) is exposed to Stage server via NFS and can be mounted by it.

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. Refer to section **Using agent rpm for Oracle Management Agent Installation** on the following page for more information:

http://www.oracle.com/technology/software/products/oem/htdocs/provisioning_agent.html

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.

A.2.2 Frequently Asked Questions on Directives and Components

Are directives associated with Components or Images?

Directives can be associated with both components and images. As explained later during the creation of components and images, directives are actually associated with one of Staging, Pre-Install, Install or Post-Installation provisioning steps of a component. In case of an Image, directives can be associated with Stage, Cleanup, Post-Install, and Diagnostics provisioning steps of an Image.

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 A–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	his 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?

Look at the following table to know the possible Status values and what they signify:

Table A–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 images?

Yes. Components are reusable and a given component can be a part of multiple images at the same time.

Do I need to modify the images associated with a component, if the component is edited?

No. If the component is changed for some reason then the change is propagated to all the images with contain that component.

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 if the working directory does not have enough space while creating a clone component?

The component creation job that is kicked off in the end will indicate that the component creations failed because of insufficient space. Create the component again by specifying a temporary directory with sufficient space and the creation will succeed.

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 component? How can one change it?.

What is the purpose of associating Properties with Generic Components?

Properties are used to provide flexibility for customizing the components depending on need. The property values are fed to the directives, which then customize the installation depending upon the values.

A.2.3 Frequently Asked Questions on Network Profiles and Images

Can I use 10.1.x IP addresses for Private Network Configuration?

No. Because of a known limitation, this type of addressing cannot be done for private IPs of the cluster nodes.

What is Reset Timeout?

After the OS installation on a target machine, management agent is installed on it, which reports to OMS and makes the machine a managed target of the Enterprise Manager. The time duration between the start of OS provisioning to agent reporting back to OMS server is known is the Reset Timeout. This is used by the provisioning application for deciding it provisioning process has timed out. If the machine or network is slow then it is advisable to set a high reset value.

B

Using emctl partool Utility

This appendix introduces you to the emctl partool utility and explains how you can use it to perform critical tasks such as exporting Deployment Procedures, importing PAR files, and so on. In particular, this appendix covers the following:

- Overview of Provisioning Archive (PAR) Files
- Overview of emctl partool Utility
- Checking Oracle Software Library
- Exporting Deployment Procedures
- Importing PAR Files

B.1 Overview of Provisioning Archive (PAR) Files

Provisioning Archive (PAR) files are archive files that contain one of the following:

- Deployment Procedures and the associated Oracle Software Library (Software Library) entities, such as components, directives, images, and so on.
- Only the Software Library entities, such as components, directives, images, and so on.

When you install Enterprise Manager Grid Control, by default, all out-of-box Deployment Procedures are installed in the form of Provisioning Archive (PAR) files at the following location:

```
$ORACLE_HOME/sysman/prov/paf
```

For example, all application server-related Deployment Procedures are stored in `asprov.par` file.

> **Note:** The PAR file for Oracle VM Management-specific Deployment Procedures is available in the following location. Here, <ORACLE_ HOME> is the Oracle home of the OMS:
>
> ```
> $<ORACLE_HOME>/sysman/admin/deploymentprocedures
> ```

To use the Deployment Procedures in Enterprise Manager Grid Control, you need not import or export these PAR files. The only prerequisite is to set up Oracle Software Library (Software Library). Once you set up the Software Library, Enterprise Manager Grid Control will automatically extract the contents of these default PAR files and configure them so that you can start using the Deployment Procedures from the Grid Control console. When importing or exporting components and/or directives that

contain properties with secret values, password will need to be provided to create Oracle Wallet to securely store/retrieve these properties.

B.2 Overview of emctl partool Utility

Over a period of time, you might have customized some Deployment Procedures, and you might want to reuse them in another instance of Enterprise Manager Grid Control. Under such circumstances, you might want to export the customized Deployment Procedures from one instance of Enterprise Manager Grid Control, and deploy them to another instance of Enterprise Manager Grid Control.

The emctl partool utility is a tool offered by Enterprise Manager Grid Control that helps you perform these functions using the command line interface. Essentially, the emctl partool utility helps you:

- Export Deployment Procedures and its associated components and directives as PAR files

- Import PAR files to the same instance or any other instance of Enterprise Manager Grid Control

The emctl partool utility is located in the `$ORACLE_HOME/bin` directory, and the out-of-box PAR files are located at `$ORACLE_HOME/sysman/prov/paf`.

> **Note:**   The PAR file for Oracle VM Management-specific Deployment Procedures is available in the following location. Here, <ORACLE_HOME> is the Oracle home of the OMS:
>
> `$<ORACLE_HOME>/sysman/admin/deploymentprocedures`

The following is the usage information displayed when you run `$ORACLE_HOME/bin/emctl partool`:

```
emctl partool <deploy|view> -parFile <file> -force(optional)
emctl partool <deploy|view> -parFile <file> -force(optional) -ssPasswd <password>
emctl partool <deploy|view> -parDir <dir> -force(optional)
emctl partool export -guid <procedure guid> -file <file> -displayName <name>
-description <desc> -metadataOnly(optional)
emctl partool -check
emctl partool -help
```

> **Note:**   After you successfully deploy the PAR files, restart the OMS.

Table B–1 describes the additional options that can be used with the emctl partool utility.

Table B–1   emctl partool Options

Option	Description
-repPasswd <repPassword>	Indicates the repository password. User will be prompted for the repository password if -repPasswd is not specified on the command line.
	Note: Providing a password on the command line is insecure and should be avoided in a production environment.

Table B–1   (Cont.) emctl partool Options

Option	Description
-force	Forces the Software Library entities to be created or uploaded again. If already present, it creates a new revision.
-check	Checks whether the Software Library is configured.
-file <file>	Represents the PAR file.
-action <deploy	view	export>	Deploys, views, or exports PAR files.
-verbose	Indicates verbose mode.
-help	Displays Help information.
-displayName <displayName>	Indicates PAR file name.
-parDir <dir>	Directory where the PAR files are located.
-metadataOnly	Flag for metadata-only exports.
-guid <guid>	Procedure GUID to export. To export multiple procedures, provide the GUIDs separated by ","
-parFile <file>	Path to the PAR file.
-description <description>	PAR file description.
-ssPasswd <secretStorePassword>	This is optional. If used with -action export; if any of the exported Software Library entity contains a secret property, an Oracle Wallet is created to store the value of the secret property. Oracle Wallet is created using the specified password. You are prompted to enter a password if -ssPasswd switch is used and if password is not supplied as a command line argument. You must use the same password while importing the PAR file in a new repository. If used with -action <deploy	view>; if the PAR file contains any password protected Oracle Wallet (that stores an entity's secret property values), then this parameter is required to open the store. You are prompted to enter a password if -ssPasswd switch is used and password is not specified as a command line argument.

B.3 Checking Oracle Software Library

Before running the emctl partool utility to export or import PAR files, ensure that the $ORACLE_HOME environment variable is set to the Oracle home directory of Oracle Management Service (OMS) and a Software Library path is configured.

To check the Software Library, run the following command:

```
$ORACLE_HOME/bin/emctl partool check
```

B.4 Exporting Deployment Procedures

To export Deployment Procedures, you must first obtain the GUID of those Deployment Procedures, and then run the emctl partool utility to create a PAR file. This section explains the following:

- Obtaining Deployment Procedure's GUID

- Creating PAR File

> **Note:** If you want to export custom Deployment Procedures (PAR files) from Enterprise Manager 10g Grid Control Release 5 (10.2.0.5) and import them to Enterprise Manager 11g Grid Control Release 1 (11.1.0.0), then refer to the *My Oracle Support* note 1110052.1.

B.4.1 Obtaining Deployment Procedure's GUID

To obtain the GUID of a Deployment Procedure using Enterprise Manager Grid Control, follow these steps:

1. In Grid Control, click the **Deployments** tab.

2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.

3. Deployment Procedure Manager page, in the Procedures table, click the deployment procedure you want to export.

4. On the View Procedure page, note the URL of the page from the address bar of the browser.

   The format of the URL should be similar to this:

   ```
   http://<OMS
   host>:<port>/em/console/paf/procedureView?guid=<value of
   GUID>
   ```

To obtain the GUID of a Deployment Procedure using Enterprise Manager Command Line Interface (EM CLI), run the `get_procedures` verb. This command displays information such as GUID, Procedure Type, Name, Version, and Created By.

```
emcli get_procedures [-type={procedure type}]
```

[ ] denotes that the parameter is optional. If you do not provide this paramaeter, then the verb displays information about all the Deployment Procedures. If you provide this parameter, then the verb displays information only about the procedure type you have specified.

B.4.2 Creating PAR File

To create a PAR file that contains one or more Deployment Procedures, run the emctl partool utility with the *export* option as the *action*, and quote the GUIDs of the Deployment Procedures you want to export. Ensure that you separate the GUIDs by a comma.

```
$ORACLE_HOME/bin/emctl partool export -guid <GUID> -file
exportedDP.par -displayName "User exported DP" -description
"<description>"
```

For example, if the GUID of the Deployment Procedure that you want to export is FAC05DD31E3791C3E030579D23106C67, then run the following command:

```
$ORACLE_HOME/bin/emctl partool export -guid
FAC05DD31E3791C3E030579D23106C67 -file exportedDP.par
-displayName "User exported DP" -description "Deployment
Procedure to be copied to other OMS"
```

After you run this command, a new PAR file named exportedDP.par is created in the directory where you ran the command. You can then import this PAR file to the same instance of Enterprise Manager Grid Control or another instance, multiple times.

To export multiple deployment procedures, separate the GUIDs with commas as follows:

```
$ORACLE_HOME/bin/emctl partool export -guid
"06B62B6ED5DA20BCE040578C850862A7,0C96E96D9818BC5FE040578C850862
0F,09AEFF331025AAD0EE40578C85FB5772" -file $ENV{T_WORK}/tvmgf_
partool_multi_dp.par -displayName "partool multi dp test"
-description "partool multi dp test description" -repPasswd
sysman
```

> **Note:** When a procedure is exported using emctl partool, any directives or components referred by the procedure are also exported. However, only the latest revision of these directives or components will be exported. If you do not want to export components or directives, you can specify the -metadataOnly flag when running emctl partool.

> **Note:** If you want to export custom Deployment Procedures (PAR files) from Enterprise Manager 10g Grid Control Release 5 (10.2.0.5) and import them to Enterprise Manager 11g Grid Control Release 1 (11.1.0.0), then refer to the *My Oracle Support* note 1110052.1.

B.5 Importing PAR Files

You can import PAR files using the command line interface or the graphical user interface offered by Enterprise Manager Grid Control, that is, the console. This section explains the following:

- Importing Using Command Line Interface
- Importing Using Grid Control Console

> **Note:**
> - If you want to export custom Deployment Procedures (PAR files) from Enterprise Manager 10g Grid Control Release 5 (10.2.0.5) and import them to Enterprise Manager 11g Grid Control Release 1 (11.1.0.0), then refer to the *My Oracle Support* note 1110052.1.
> - After you successfully import the PAR files, restart the OMS.

B.5.1 Importing Using Command Line Interface

This section covers the following:

- Importing Specific PAR File
- Importing All PAR Files

Note:

- If you want to export custom Deployment Procedures (PAR files) from Enterprise Manager 10g Grid Control Release 5 (10.2.0.5) and import them to Enterprise Manager 11g Grid Control Release 1 (11.1.0.0), then refer to the *My Oracle Support* note 1110052.1.

- After you successfully import the PAR files, restart the OMS.

B.5.1.1 Importing Specific PAR File

To import or deploy a specific PAR file, run the following command:

```
$ORACLE_HOME/bin/emctl partool deploy -parFile $ORACLE_
HOME/sysman/prov/paf/<par_file_name> -force
```

For example:

```
$ORACLE_HOME/bin/emctl partool deploy -parFile $ORACLE_
HOME/sysman/prov/paf/asprov.par -force
```

Note: If you have multiple OMSes in your environment, then you need run the emctl partool utility only once to deploy any PAR files or to perform other related operations.

B.5.1.2 Importing All PAR Files

To import or deploy all the PAR files in a directory, run the following command:

```
$ORACLE_HOME/bin/emctl partool deploy -parDir $ORACLE_
HOME/sysman/prov/paf/ -force
```

B.5.2 Importing Using Grid Control Console

To import PAR files or deploy them to an OMS, you can use the emctl partool utility. Alternatively, you can also log in to Enterprise Manager Grid Control, navigate to the Deployment Procedure Manager page, and click **Upload** to upload a PAR file.

If a PAR file is being deployed using Grid Control, the SQL files in the PAR file will not be executed.

Note: When importing or exporting components and/or directives that contain properties with secret values, you must use the -ssPasswd command and provide the secret store password to create Oracle Wallet. This ensures that the properties with secret values are securely stored using an Oracle Wallet, and can be accessed while importing with only the Oracle Wallet password.

For more information about the -ssPasswd command, see Table B–1, " emctl partool Options". For more information about importing and exporting Software Library entities, see Section C.1.1, "Exporting and Importing Entities Across Oracle Enterprise Manager Deployments".

Note:

- If you want to export custom Deployment Procedures (PAR files) from Enterprise Manager 10g Grid Control Release 5 (10.2.0.5) and import them to Enterprise Manager 11g Grid Control Release 1 (11.1.0.0), then refer to the *My Oracle Support* note 1110052.1.

- After you successfully import the PAR files, restart the OMS.

C

Using Oracle Software Library

This appendix provides an overview of Oracle Software Library (Software Library) and explains how you can use it to export and import entities, delete and purge entities, and so on. In particular, this appendix covers the following:

- Using the Software Library
- De-Configuring a Software Library Location
- Software Library Maintenance Tasks
- Software Library Issues

C.1 Using the Software Library

Oracle Software Library (Software Library) is a repository that stores certified software images (for example, Oracle Database, operating system, Oracle Real Application Clusters, third party software, and so on) and other related entities. These can then be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manger Grid Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

At any point, only one Software Library can exist in your environment. You can, however, configure multiple files ystem locations for the Software Library for storing files associated with its entities.

The graphical user interface of the Provisioning application has various tabs for creating Components, Directives, Images, Network and Hardware Templates, which are created and stored in the Software Library. Various subdirectories for storing the entities can be created in the Software Library. These tabs also allow a user to delete and edit an entity stored in the Software Library. They also allow a user to view metadata information for an entity. Multiple Software Library locations can be configured and when a file/binary is being associated with an entity, the location with the most space is picked.

On the Administration tab of the Provisioning page, the Software Library Configuration section lists the directory location, free space, used space, and space used by deleted entities. In addition, the following functionalities are also available in the Software Library Configuration section:

Purge

When an entity is deleted, the binary file associated with the entity continues to exists on the disk till it is completely removed. The space usage by such binary files can be monitored using "Used Space By Deleted Entities" column against every location. Use

the Purge functionality in the Software Library section to permanently remove deleted entities and their respective binary files from the Software Library.

Refresh

The space usage for all Software Library locations can be computed using the Refresh functionality. On clicking **Refresh**, the Free Space, Used Space, and Used Space by Deleted Entities display the latest space usage details for each location. The Last Computed column gives the date when free space was computed last.

Check Accessibility

When you add a new Software Library directory location or a new OMS, then use the Check Accessibility option to verify that the Software Library location is accessible from all Oracle Management Services (OMS), in case of multiple OMS in your Enterprise Manager deployment. Ensure that normal preferred credentials are set for all hosts running Oracle Management Service to use the Check Accessibility functionality. You can also use the out-of-box Enterprise Manager policy "Non-Shared Software Library Existence" to check if any non-shared Software Library location exists among all OMS.

Export and Import

You can export Software Library entities and import Provisioning Archive (PAR) files in the Software Library. For more details, see Section C.1.1, "Exporting and Importing Entities Across Oracle Enterprise Manager Deployments".

> **Caution:**   Do *not* attempt to export Software Library entities from a higher version of Enterprise Manager Grid Control and import them to a lower version of Enterprise Manager Grid Control. Attempting such an unsupported operation will produce undesired results. However, you can always export the entities from a lower version to a higher version of Enterprise Manager Grid Control.

Add, Remove, and Edit

You can add a new Software Library directory location and edit or remove an existing Software Library directory location.

> **Note:** The following are the limitations and known issues related to the Software Library:
>
> - Editing or removing an existing Software Library directory location is possible only if there are no entities using this location for storing its files or binaries.
>
> - Relocation of the Software Library is not supported from the Enterprise Manager Grid Control console.
>
> - The directory /tmp on Linux and Unix, or C:\TEMP on Microsoft Windows must not be used as the Software Library location.
>
> - (For Microsoft Windows Only) In order to have the Software Library configured across multiple OMSes, you must have a disk partition with the same name for all the OMS hosts and have a cluster file system running on the partition. This can be accomplished by installing OCFS via an Oracle CRS version 10.2.0.1 installation.
>
>   The OMS service must be configured to login as a domain user account, and the domain user account must have the required permissions for accessing the shared storage locations. All the OMS services must login using the same domain user account.

C.1.1 Exporting and Importing Entities Across Oracle Enterprise Manager Deployments

Software Library entities can be exported and imported across different Software Libraries used by different Enterprise Manager deployments.

> **Caution:** Do *not* attempt to export Software Library entities from a higher version of Enterprise Manager Grid Control and import them to a lower version of Enterprise Manager Grid Control. Attempting such an unsupported operation will produce undesired results. However, you can always export the entities from a lower version to a higher version of Enterprise Manager Grid Control.

Oracle provides PAR files that contain deployment procedures and Software Library entities such as components and directives, which are required to run them for Provisioning and Patching operations.

A PAR file is generated only when you export some entities from the Software Library. Therefore, to generate a PAR file, in Grid Control, click **Deployments** and then **Provisioning**. On the Provisioning page, click the **Administration** tab, and on the Administration page, from the Software Library Configuration section, select a directory location and click **Export**.

The following shows the Export Software Library Entities page.

> **Note:** Although a PAR file can accommodate both Deployment Procedures and Software Library entities, when you export from the Export Software Library Entities page, only the Software Library entities are exported.

On the Export Software Library Entities page, specify the following details:

Element	Description
Directory Location	Location where the PAR file will be created. This is a required parameter.
PAR Filename	The file name must end with a .par extension. If the given filename already exists in the specified location, then filename_1.par is created. If that also exists, then filename_2.par is created and so on. The name of file can be viewed in the Job output page. This is a required parameter.
Oracle Wallet Password	This is used to securely store all secret property values in an encrypted Oracle Wallet. This is a required parameter. However, the password is used only if any exported entity has a secret property.
Exclude File	Select this option to exclude file, binary, or script associated with an entity to be exported.
Export	You can choose to either export all entities or export a selected entity.
	If you choose to export a selected entity, click **Add** to select entity. In popup to select the entity, search for the entities you want to export, select and add them. You need not select all referenced entities. For example, if an image is referring to a component and a directive, you can select only the image. During the export process, the component and the directive are also exported.

Click **Import** to import a PAR (Provisioning Archive) file in the Software Library. The following shows the Import Software Library Entities Page.

On the Import Software Library Entities page, specify the following details:

Element	Description
PAR File Location	Location of the PAR file. This location must be accessible by the Oracle Management Server.
Oracle Wallet Password	This is an optional field. The password is required if the PAR file contains an Oracle Wallet that stores secret property values. PAR files provided by Oracle will not have any associated Oracle Wallet. Hence, if you are importing an Oracle provided PAR file, then password need not be provided. However, if you are importing a PAR file that was created by the Export process (using Enterprise Manager Console or PARDeploy script) then you will need to provide this password.

The same password that was used during export should be used during import. You can ignore this password if there is no Oracle Wallet contained in the PAR file. If the PAR file has an associated Oracle Wallet, the import process will fail if no password is provided. |
| Force New Revision | If an entity being imported is already present in the Software Library, then this option allows the user to force a newer revision of the entity to be created during import. If the option is not selected and the PAR file contains an entity in a directory that already exists in the repository, then the import process will fail. |

You can also use the `deploymentLibraryExport.pl` and `deploymentLibraryImport.pl` scripts for importing and exporting the Software Library entities respectively. These scripts are present in `/bin` directory of the OMS Oracle Home and provide the flexibility to transfer entities from a test to production or between different production environments.

The scripts support the following import/export use cases:

- Exporting all directives and importing the same.

- Exporting all components and importing the same.

- Exporting all entities in the Software Library (including Network profiles, Images etc.) and importing the same.

The export script is available at the <OMS HOME>/bin/deploymentLibraryExport.pl location on Linux and Windows.

The import script is available at the <OMS HOME>/bin/deploymentLibraryImport.pl location on Linux and Windows.

For information on using the import and export scripts, type `--help` in the command line.

C.1.2 Deleting and Purging Software Library Entities

Software Library entities can be deleted from the relevant tabs provided by the provisioning application. By deleting an entity, the entity is no longer available for selection, view, or edit, and will not show up in the user interface. But it is to be noted that deleting the entity does not purge the file associated with it from the Software Library file system. If you delete an entity, it will not appear in the user interface, but will continue to exist on disk and take up disk space. To clean up and delete entities completely from the file system of the Software Library, you will need to purge the deleted entities.

To purge deleted entities, click **Purge** in the Software Library section.

You can also purge deleted entities by running the purgeDeploymentLibrary script. The script is located at the following location:

For OMS on Linux: `<OMS HOME>/bin/purgeDeploymentLibrary`

C.2 De-Configuring a Software Library Location

Go to the Administration tab, Software Library Configuration section, choose the Software Library entry, and then click **Remove** to de-configure a Software Library location. Once a Software Library entry is deleted, it becomes inaccessible.

> **Note:** De-configuring an existing Software Library directory location is possible only if there are no entities using this location for storing its files or binaries.

Even though the deleted location is not used further for storing binaries/files/scripts of newly-created Software Library entities, it may still store files of already created entities. If the location is not accessible from all Oracle Management Servers, user can experience problems during deployment of such entities later.

If multiple locations for the Software Library are configured, when you deconfigure the Software Library locations, you can remove all except one Software Library location.

C.3 Software Library Maintenance Tasks

It is recommended that the system administrator periodically performs the following administrative tasks to ensure that the Software Library is functioning properly:

- Refresh the Software Library regularly to compute the free and used disk space.
- Purge deleted entities to conserve disk space.
- Check accessibility of the configured Software Library locations to ensure that they are accessible to all Oracle Management Servers. If you add another Oracle Management Server, use the Check Accessibility feature to ensure that the Software Library location is accessible to the new Oracle Management Server.

For information about these features, see Section C.1, "Using the Software Library".

C.4 Software Library Issues

Component/Directive/Image creation sometimes fail with "There was an error saving '<Component Name>' below '<Folder Name>' in the Software Library.

Contact your administrator. You can find a detailed report in the log file at ORACLE_HOME/sysman/log/emoms.log."

This may happen if the Software Library is not configured with Enterprise Manager. Configure the Software Library as explained in Section 3.1.2, "Setting Up Oracle Software Library". Create the components once the Software Library is configured.

"Unable to access binary of a Software Library Entity (Components/Directives) during deployment time even though all Software Library locations are accessible"

If the contents of Software Library location (that stores binaries of Software Library entities such as components or directives) are deleted, then the user will encounter problems during the deployment of such entities. Since the metadata of the entities is not removed from the Management Repository, the entities will continue to show up on the Console. Upload the binary of all such entities again.

D

Creating Templates for Oracle Database

This appendix explains how you can use Database Configuration Assistant (DBCA) to create templates for structuring newly created databases.

Note that DBCA can also be used to create scripts for creating databases using the templates. The necessary options for creating templates and scripts in the DBCA help have been highlighted in this appendix.

This appendix covers the following:

- Overview
- Creating Database and Templates Using DBCA
- DBCA Command Line Help Options

D.1 Overview

For detailed information on using DBCA to create and configure a database, see the *Oracle Database 2 Day DBA Guide* available at:

http://www.oracle.com/technology/documentation/database.html

D.2 Creating Database and Templates Using DBCA

You can launch DBCA in interactive mode, provide details for creating a database, and in the end, instead of creating a database, you can save the specified information as a template. You can then use this template to create databases.

This section explains how you can create a database and other templates using DBCA.

Creating a Database using DBCA

Example: Create db glob - use template glob1.dbc

```
$ORACLE_HOME/bin/dbca -silent -createDatabase -templateName
glob1.dbc -gdbname glob glob-sid -characterset AL32UTF8
```

Example: Create db glob with the location of datafiles - creating a database with the location where datafiles have to be placed in a given location

```
$ORACLE_HOME/bin/dbca -silent -createDatabase -templateName
glob1.dbc -gdbname glob glob-sid-characterset AL32UTF8
-datafiledestination D: \ oracle\product\10.2.0\oradata\glob
```

> **Note:** If you want these databases to be monitored by Enterprise Manager Grid Control, then unlock the DBSNMP password of the databases.

Creating .dbt Template File from Existing Databases

```
template dbt dbca -silent -createTemplateFromDB -sourceDB GLOB
-sysDBAUserName sys -sysDBAPassword change_on_install
-templateName glob_template
```

Creating .dbc Template File from Existing Databases

```
template dbc dbca -silent -createCloneTemplate -sourceDB GLOB
-sysDBAUserName sys -sysDBAPassword change_on_install
-templateName glob_clone
```

Creating a Template File and Deleting Databases Using a Template

- Create a clone template glob_clone from an existing database by specifying the following parameters:

  ```
  dbca -silent -createCloneTemplate -sourceDB GLOB
  -sysDBAUserName sys -sysDBAPassword change_on_install
  -templateName glob_clone
  ```

- Delete database glob silently or copy glob_clone.dbc and GLOB_CLONE.DFB to other host:

  ```
  dbca -silent -deleteDatabase -sourceDB GLOB -sysDBAUserName
  sys -sysDBAPassword change_on_install
  ```

- Create database using database template clone:

  ```
  dbca -silent -createDatabase -templateName glob_clone.dbc
  -gdbname glob glob-sid-characterset AL32UTF8
  ```

Creating Sample Script (template) and Databases

- Create db scripts in the directory D: \ oracle \ product \ 10.2.0 \ db_1 \ assistants \ dbca \ templates \ scr

  ```
  dbca -silent -generateScripts -templateName glob_clone.dbc
  -gdbName glob -scriptDest D: \ oracle \ product \ 10.2.0 \
  db_1 \ assistants \ dbca \ templates \ scr
  ```

- Create database using scripts

  ```
  set ORACLE_SID = GLOB
  ```

  ```
  glob.bat
  ```

D.3 DBCA Command Line Help Options

```
dbca [-silent |-progressOnly |-customCreate] (<command>
<options>) | ([<command> [options]]-responseFile <response
file>) [-continueOnNonFatalErrors <true | false>]
```

Creating a Database

- - createDatabase
- - templateName <name of an existing template>

- [-cloneTemplate]

- - gdbName <global database name>

- [- Sid <database system identifier>]

- [-sysPassword <SYS user password>]

- [-systemPassword <SYSTEM user password>]

- [-emConfiguration
  <CENTRAL|LOCAL|ALL|NOBACKUP|NOEMAIL|NONE>

  [-emConfiguration
  <CENTRAL|LOCAL|ALL|NOBACKUP|NOEMAIL|NONE>

- -dbsnmpPassword <DBSNMP user password> - dbsnmpPassword <DBSNMP
  user password>

- -sysmanPassword <SYSMAN user password> - sysmanPassword <SYSMAN user
  password>

- [-hostUserName <Host user name for EM backup job> [-hostUserName <Host
  user name for EM backup job>

- -hostUserPassword <Host user password for EM backup job> - hostUserPassword
  <Host user password for backup EM job>

- -backupSchedule <Daily backup schedule in the form of hh:mm>] -
  backupSchedule <Daily backup schedule in the form of hh:mm>]

- [-smtpServer <Outgoing mail (SMTP) server for email notifications> [smtpServer
  <Outgoing-mail server for (SMTP) email notifications>

- -emailAddress <Email address for email notifications>] - emailAddress <Email
  email address for notifications>]

- [-centralAgent <Enterprise Manager central agent home>]] [-centralAgent
  <Enterprise Manager central agent home>]]

- [-datafileDestination <destination directory for all database files> |
  -datafileNames <a text file containing database objects such as controlfiles,
  tablespaces, redo log files and spfile to their corresponding raw device file names
  mappings in name=value format.>] [-datafileDestination <destination directory
  database for all files> |-datafileNames <a text file containing database objects such
  as controlfiles, tablespaces, redo log files and spfile to their corresponding raw
  device file names mappings in name=value format.>]

- [-recoveryAreaDestination <destination directory for all recovery files>]
  [-recoveryAreaDestination <destination directory for all recovery files>]

- [- datafileJarLocation < location of the data file jar, used only for clone database
  creation>] [- DatafileJarLocation <location of the data file jar, used only for clone
  database creation>]

- [-storageType < FS | ASM | RAW> [-storageType <FS | ASM | RAW>

- [-asmSysPassword <SYS password for ASM instance>] [-asmSysPassword <SYS
  password for ASM instance>]

- [-diskString <disk discovery path to be used by ASM>] [-diskString <disk
  discovery path to be used by ASM>]

- [-diskList <comma seperated list of disks for the database area disk group>
  [diskList <comma-seperated list of disks for the database area disc group>

- -diskGroupName <database area disk group name> diskGroupName <database-area disk group name>

- -redundancy <HIGH|NORMAL|EXTERNAL>] - redundancy <HIGH|NORMAL|EXTERNAL>]

- [-recoveryDiskList <comma seperated list of disks for the recovery area disk group> [recoveryDiskList <comma-seperated list of disks for the recovery area disc group>

- -recoveryGroupName <recovery area disk group name> recoveryGroupName <recovery-area disk group name>

- -recoveryGroupRedundancy <HIGH|NORMAL|EXTERNAL>]] - recoveryGroupRedundancy <HIGH|NORMAL|EXTERNAL>]]

- [-characterSet <character set for the database>] [-characterSet <character set for the database>]

- [- nationalCharacterSet < national character set for the database>] [-NationalCharacterSet <national character set for the database more]

- [-registerWithDirService <true | false> [-registerWithDirService <true | false>

- -dirServiceUserName <user name for directory service> - dirServiceUserName <user name for directory service>

- -dirServicePassword <password for directory service > - dirServicePassword <password for directory service>

- -walletPassword <password for database wallet >] - walletPassword <password wallet database for more]

- [- listeners < list of listeners to configure the database with>] [- Listeners <list of listeners to configure the database with more]

- [-variablesFile <file name for the variable-value pair for variables in the template>]] [-variablesFile <file name for the pair variable-value for variables in the template>]]

- [- variables < comma seperated list of name=value pairs>] [- Variables <comma seperated list of name = value pairs>]

- [-initParams <comma seperated list of name=value pairs>] [initParams <comma-seperated list of name=value pairs>]

- [-memoryPercentage <percentage of physical memory for Oracle>] [-memoryPercentage <percentage of physical memory for Oracle>]

- [-databaseType <MULTIPURPOSE|DATA_WAREHOUSING|OLTP>]] [-databaseType <MULTIPURPOSE|DATA_WAREHOUSING|OLTP>]]

Configuring a Database

- - configureDatabase

- -sourceDB <source database sid > - sourceDB <source database sid>

- [-sysDBAUserName <user name with SYSDBA privileges> [-sysDBAUserName <user name with SYSDBA privileges>

- -sysDBAPassword <password for sysDBAUserName user name>] - sysDBAPassword <password for sysDBAUserName user name>]

- [-registerWithDirService|-unregisterWithDirService|-regenerateDBPassword <true | false> [registerWithDirService-|-unregisterWithDirService |-regenerateDBPassword <true | false>

- -dirServiceUserName <user name for directory service> - dirServiceUserName <user name for directory service>

- -dirServicePassword <password for directory service > - dirServicePassword <password for directory service>

- -walletPassword <password for database wallet >] - walletPassword <password wallet database for more]

- [-emConfiguration <CENTRAL|LOCAL|ALL|NOBACKUP|NOEMAIL|NONE> [-emConfiguration <CENTRAL|LOCAL|ALL|NOBACKUP|NOEMAIL|NONE>

- -dbsnmpPassword <DBSNMP user password> - dbsnmpPassword <DBSNMP user password>

- -symanPassword <SYSMAN user password> - symanPassword <SYSMAN user password>

- [-hostUserName <Host user name for EM backup job> [-hostUserName <Host user name for EM backup job>

- -hostUserPassword <Host user password for EM backup job> - hostUserPassword <Host user password for backup EM job>

- -backupSchedule <Daily backup schedule in the form of hh:mm>] - backupSchedule <Daily backup schedule in the form of hh:mm>]

- [-smtpServer <Outgoing mail (SMTP) server for email notifications> [smtpServer <Outgoing-mail server for (SMTP) email notifications>

- -emailAddress <Email address for email notifications>] - emailAddress <Email email address for notifications>]

- [-centralAgent <Enterprise Manager central agent home>]] [-centralAgent <Enterprise Manager central agent home>]]

Creating a Template from an Existing Database

- - createTemplateFromDB

- -sourceDB <service in the form of <host >: <port>:<sid>> - sourceDB <service in the form of <host>: <port>: <sid>>

- -templateName <new template name> - templateName <new template name>

- -sysDBAUserName <user name with SYSDBA privileges> - sysDBAUserName <user name with SYSDBA privileges>

- -sysDBAPassword <password for sysDBAUserName user name> - sysDBAPassword <password for sysDBAUserName user name>

- [-maintainFileLocations <true | false>] [-maintainFileLocations <true | false>]

Creating a Clone Template from an Existing Database

- - createCloneTemplate

- -sourceSID <source database sid > - sourceSID <source database sid>

- -templateName <new template name> - templateName <new template name>

- [-sysDBAUserName <user name with SYSDBA privileges> [-sysDBAUserName <user name with SYSDBA privileges>

- -sysDBAPassword <password for sysDBAUserName user name>] -sysDBAPassword <password for sysDBAUserName user name>]

- [-maintainFileLocations <true | false>] [-maintainFileLocations <true | false>]

- [-datafileJarLocation <directory to place the datafiles in a compressed format>] [-datafileJarLocation <directory to place the datafiles in a compressed format>]

Generate Scripts to Create Database

- - generateScripts

- -templateName <name of an existing template > - templateName <name of an existing template>

- -gdbName <global database name> - gdbName <global database name>

- [-scriptDest <destination for all the scriptfiles>] [-scriptDest <destination for all the scriptfiles>]

Deleting a Database

- - deleteDatabase

- -sourceDB <source database sid > - sourceDB <source database sid>

- [-sysDBAUserName <user name with SYSDBA privileges> [-sysDBAUserName <user name with SYSDBA privileges>

- -sysDBAPassword <password for sysDBAUserName user name>] -sysDBAPassword <password for sysDBAUserName user name>]

Configure ASM Disk Groups

- - configureASM

- [-asmSysPassword <SYS password for ASM instance>] [-asmSysPassword <SYS password for ASM instance>]

- [-diskString <disk discovery path to be used by ASM>] [-diskString <disk discovery path to be used by ASM>]

- [-diskList <comma seperated list of disks for the database area disk group> [diskList <comma-seperated list of disks for the database area disc group>

- -diskGroupName <database area disk group name> diskGroupName <database-area disk group name>

- -redundancy <HIGH | NORMAL | EXTERNAL>]] - redundancy <HIGH | NORMAL | EXTERNAL>]]

- [-recoveryDiskList <comma seperated list of disks for the database area disk group> [recoveryDiskList <comma-seperated list of disks for the database area disc group>

- -recoveryGroupName <database area disk group name> recoveryGroupName <database-area disk group name>

- -recoveryGroupRedundancy <HIGH | NORMAL | EXTERNAL>] -recoveryGroupRedundancy <HIGH | NORMAL | EXTERNAL>]

- [-emConfiguration <CENTRAL | NONE> [-emConfiguration <CENTRAL | NONE>

- -centralAgent <Enterprise Manager central agent home>]] - centralAgent <Enterprise Manager central agent home>]]

- Query for help by specifying the following options: -h ｜ -help Query for help by specifying the following options:-h ｜-help

E

Out-Of-Box RuntimeData Templates

This apendix lists the out-of-box runtime data templates for different Deployment Procedures. In particular, this appendix covers the following:

- Out-Of-Box RuntimeData Templates for Database Deployment Procedures
- Out-Of-Box RuntimeData Templates for Oracle Application Server Deployment Procedures
- Out-Of-Box RuntimeData Templates for Patching Procedures

E.1 Out-Of-Box RuntimeData Templates for Database Deployment Procedures

The following tables lists the out-of-box runtimedata templates for database deployment procedures.

Table E–1 Out-of-Box RuntimeData Templates for Database Deployment Procedures

Procedure Name	Provisioning Mode	Procedure Type	Template Name
Oracle Clusterware/Oracle Real Applications Clusters (RAC) Provisioning for UNIX	Using a Gold Image from the Software Library	RACPROV	crsasmrac_gold_prov_template.xml
Oracle Clusterware/Oracle Real Applications Clusters (RAC) Provisioning for UNIX	Using an existing installation, that is, a reference host	RACPROV	crsasmrac_inshome_prov_template.xml
Extend Cluster Database	NA	RACPROV	crsasmrac_extend_cluster_template.xml
Scale Down/Delete Oracle Real Applications Clusters (RAC)	Scale down	RACPROV	rac_partial_delete_node_template.xml
Scale Down/Delete Oracle Real Applications Clusters (RAC)	Complete deletion	RACPROV	rac_complete_delete_node_template.xml
Oracle Database Provisioning	Database provisioning	SIDB	single_instance_database_template.xml

E.2 Out-Of-Box RuntimeData Templates for Oracle Application Server Deployment Procedures

The following tables lists the out-of-box runtimedata templates for Oracle Application Server deployment procedures.

Table E–2   Out-of-Box RuntimeData Templates for Oracle Application Server Deployment Procedures

Procedure Name	Provisioning Mode	Procedure Type	Template Name
10.1.3 AS Provisioning	Using a gold image from the Software Library	ASPROV	as_gold_imgage_cloning_ template.xml

E.3 Out-Of-Box RuntimeData Templates for Patching Procedures

The following tables lists the out-of-box runtimedata templates for patching procedures.

Table E–3   Out-of-Box RuntimeData Templates for Patching Procedures

Procedure Name	Template Name
Patch Oracle Database	StandAlone_template.xml

F

Sample Property Files for the Out-of-Box RuntimeData Templates

This appendix lists the sample property files for out-of-box RuntimeData templates. These property files are packaged with the product. To see the contents of a file, access the file location that is mentioned against the property file.

Table F–1   Sample Property Files for the Out-of-Box RuntimeData Templates

Procedure Name	Property File Name	Directory Location
RAC Provisioning using gold image	CRS-ASM-RAC_provisioning_using_GOLD-Image.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/crsasmracgoldprov
RAC Provisioning using an existing, running, reference node	CRS-ASM-RAC_provisioning_using_Reference-Installation.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/crsasmracinshomeprov
RAC Provisioning using software binaries archived from the installation medium	CRS-ASM-RAC_provisioning_using_SHIPHOME-Image.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/crsasmracshiphomeprov
Extend cluster	Cluster_Extend.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/crsasmracextendcluster
Delete cluster	Cluster_Complete_Delete.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/raccompletedeletenode
Delete node	Cluster_Partial_Delete.properties	%ORACLE_HOME%/emcli/samples/provisioning/racprovisioning/10.2.0.1.0/racpartialdeletenode
Database provisioning by cloning on Linux	SidbCloneLinux.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning by cloninng on Linux (software only)	SidbCloneLinuxSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning by cloning on Microsoft Windows	SidbCloneWin.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2

Table F–1 (Cont.) Sample Property Files for the Out-of-Box RuntimeData Templates

Procedure Name	Property File Name	Directory Location
Database provisioning by cloning on Microsoft Windows (software only)	SidbCloneWinSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning using gold image on Linux	SidbGoldImgLinux.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning using gold image on Linux (software only)	SidbGoldImgLinuxSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning using gold image on Microsoft Windows	SidbGoldImgWin.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning using gold image on Microsoft Windows (software only)	SidbGoldImgWinSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning on Linux using software binaries archived from the installation medium	SidbInstallLinux.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning on Linux using software binaries archived from the installation medium (software only)	SidbInstallLinuxSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning on Microsoft Windows using software binaries archived from the installation medium	SidbInstallWin.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2
Database provisioning on Microsoft Windows using software binaries archived from the installation medium (software only)	SidbInstallWinSwOnly.txt	%ORACLE_HOME%/emcli/samples/provisioning/dbprovisioning/10.2

G

PXE Booting and Kickstart Technology

This appendix explains the PXE booting and kickstart technology.

G.1 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.

G.2 Subnet Provisioning Usecases

Following are examples of subnet provisioning usecases:

Subnet of size 256

IP Prefix: 192.168.1.0

Hex value: C0 A8 01 00

Trimmed Hex Value: C0 A8 01

Covers IPs from 192.168.1.0 - 192.168.1.255

Subnet of size 16

IP Prefix: 192.168.1.16

Hex value: C0 A8 01 10

Trimmed Hex Value: C0 A8 01 1

Covers IPs from 192.168.1.16 - 192.168.1.31

> **Tip:** If you want to stage an image for a different size of subnet say 64, you can stage four images for four continuous 16-size subnets.

H

About Patching Automation Reports

This appendix describes the patching automation reports offered by Enterprise Manager Grid Control. In particular, this appendix covers the following:

- About Patchability Report
- About Patching Deployment Procedure Execution Summary Report

H.1 About Patchability Report

The *Patchability Report* helps you analyze the environment and identify the Oracle Databases, or Application Servers, or Oracle Clusterware targets that are suitable and not suitable for a patching operation. Before running a patching Deployment Procedure for Oracle Database and Oracle RAC, Oracle strongly recommends you to run this report so that you know whether the database targets you want to patch can be patched successfully. You can diagnose whether the database target is not suitable for patching due to missing properties or unsupported configuration.

To access the *Patchability Report*, in Grid Control, click **Reports**. On the Report Definitions page, in the table, expand **Deployment and Configuration**, and then expand **Patching Automation Reports** and click **EM Target Patchability Report**.
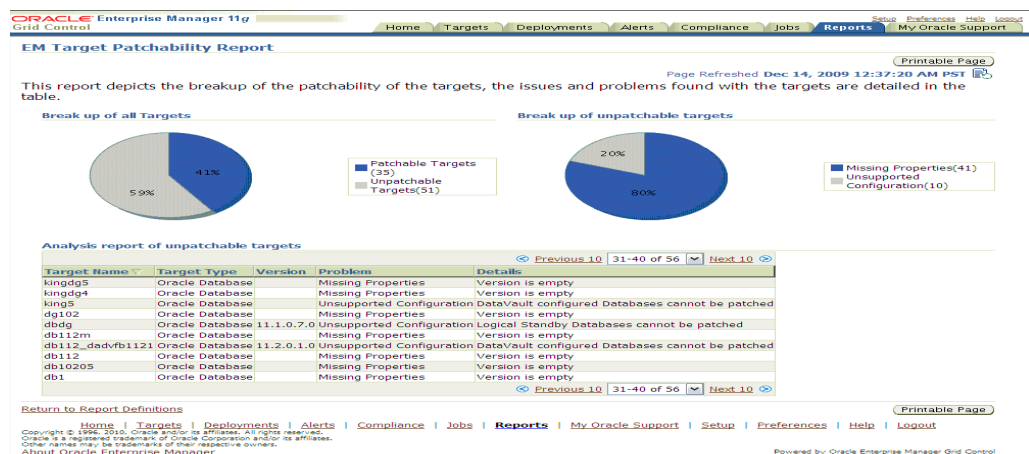
> **Note:** For information about missing property errors, see Section H.1.1, "Workarounds for Missing Property Errors". For information about unsupported configuration errors, see Section H.1.2, "Workarounds for Unsupported Configuration Errors".

H.1.1 Workarounds for Missing Property Errors

Table H–1 describes the possible missing property errors and the workarounds you can try to resolve them.

Table H–1   Missing Properties Error and Workarounds

Problem	Workaround
Empty target propertyversion	The target is not properly configured or maybe the target is unavailable. Reconfigure the target or check for metric collection errors .

Inadequate or incomplete target information collected by Oracle Management Agent.	To resolve this issue, recompute the dynamic properties and refresh the host configuration so that Oracle Management Repository is updated with the latest configuration of the host. To do so, follow these steps:
	
	1. To recompute the dynamic properties, do one of the following:
	
	Option A: Stop and restart Oracle Management Agent:
	
	`$ emctl stop agent`
	
	`$ emctl start agent`
	
	For a cluster, restart Oracle Management Agent on all the nodes of the cluster.
	
	Option B: Reload the dynamic properties. This option is simpler because you do not have to restart Oracle Management Agent.
	
	`$ emctl reload agent dynamicproperties`
	
	For a cluster, run the following command:
	
	`$ emctl reload agent dynamicproperties [<Target_name>:<Target_Type>]`
	
	For example:
	
	`$ emctl reload agent dynamicproperties crs:cluster`
	
	`$ emctl reload agent dynamicproperties oradb:oracle_database`
	
	`$ emctl reload agent dynamicproperties server1.xyz.com:host`
	
	`$ emctl reload agent dynamic properties racdb_1:rac_database`
	
	2. To update Oracle Management Repository with the latest configuration information, click **Deployments**. On the Deployments page, from the Configuration section, click **Refresh Host Configuration**. On the Refresh Host Configuration page, select the hosts whose configuration needs to be updated in Oracle Management Repository.

H.1.2 Workarounds for Unsupported Configuration Errors

Table H–2 describes the possible unsupported configuration errors and the workarounds you can try to resolve them.

Table H–2   Workarounds for Unsupported Configuration Errors

Problem	Workarounds
RAC Instance does not have an associated RAC Database	Rediscover the RAC target and add the RAC instance to the RAC database.
The database is not mediated by the OMS	The target discovery is not appropriate. Remove the target from Grid Control, and rediscover on all the Management Agents in the cluster.

H.2 About Patching Deployment Procedure Execution Summary Report

The Patching Deployment Procedure Execution Summary Report helps you analyze the results of the patching Deployment Procedures that were run over a period of time in the past. The report helps you understand the number of patching failures, prerequisite failures, successful executions, and unsupported configurations encountered by the Deployment Procedures.

To access the Patching Deployment Procedure Execution Summary Report, in Grid Control, click **Reports**. On the Report Definitions page, in the table, expand **Deployment and Configuration**, and then expand **Patching Automation Reports** and click **Patching Deployment Procedures Execution Summary**.

I
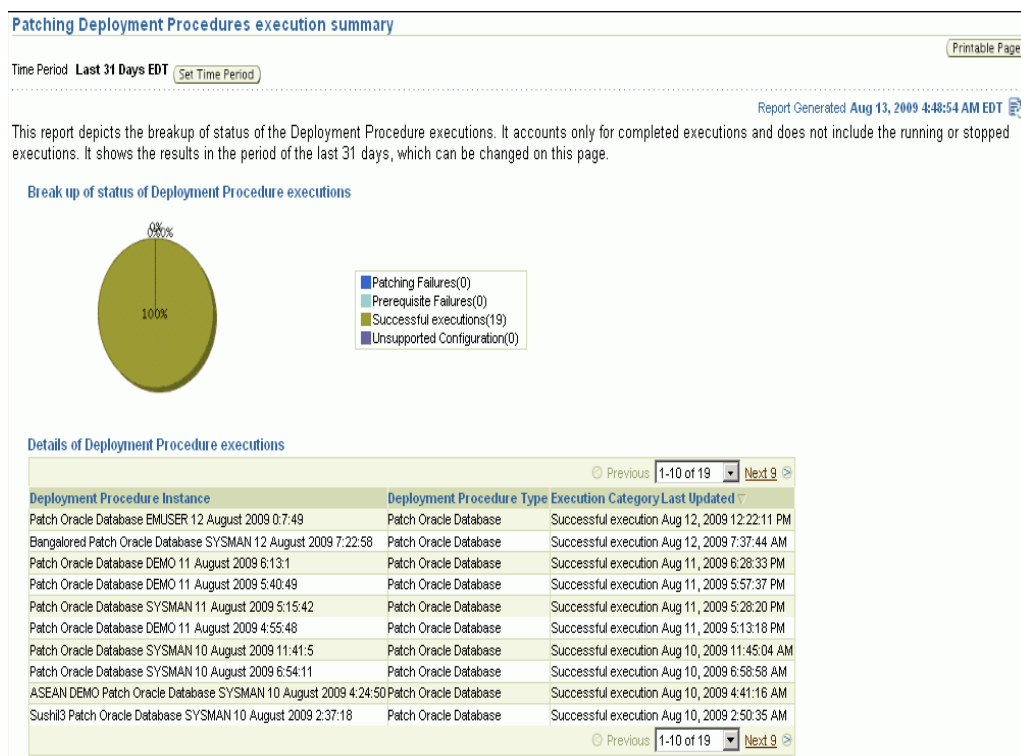
Converting Standalone Oracle Management Agent to Cluster Agent

For using the RAC-related procedures, you must have cluster-compatible Oracle Management Agents (Management Agents) on the cluster nodes, and not standalone Management Agents. However, if the Management Agents on the cluster nodes are indeed standalone Management Agents, then you must convert them to cluster-compatible Management Agents.

This appendix explains how you can convert a standalone Management Agent to a cluster-compatible Management Agent. In particular, this chapter covers the following:

- Using Graphical Mode to Convert
- Using Silent Mode to Convert

I.1 Using Graphical Mode to Convert

To convert a standalone Management Agent to a cluster-compatible Management Agent in a graphical mode, follow these steps:

1. Meet the prerequisites for installing a Management Agent as described in the *Enterprise Manager Grid Control Advanced Installation and Configuration Guide* available at:

   http://www.oracle.com/technology/documentation/oem.html

2. In Grid Control, click **Deployments**.

3. On the Deployments page, from the Agent Installation section, click **Install Agent**.

4. On the following page, click **Fresh Install**.

5. On the Installation Details page, do the following:

   – In the Source Software section, retain the default selection, that is, **Default, from Management Server location**.

   – In the Version section, select the appropriate version of the already-installed standalone Management Agents. Note that to use Deployment Procedures, the Management Agents must be at least 10g Release 2 (10.2.0.2).

   – In the Host section, select the required platform, provide a list of hosts that are part of the cluster, check **Cluster Install**, click **Populate Defaults** to automatically specify a list of cluster nodes, and then specify a cluster name.

   – In the OS Credentials section, provide the host credentials.

- In the Destination section, specify the installation base directory on the nodes that form the cluster.

- In the Additional Parameters section, supply any other optional parameters you want to run.

- Click **Continue**.

Note: If you see the following error messages on the Prerequisite Details page, ignore them and click **Continue** to proceed with the installation:

```
Some or all Agent home locations(Install base dir/agent10g or
Install Base Dir/agent11g) provided are not empty directories

The Agent Oracle Home ( InstallBaseDir/agent10g or
InstallBaseDir/agent11g), is already registered with the inventory

An agent home exists in the hosts
```

I.2 Using Silent Mode to Convert

To convert a standalone Management Agent to a cluster-compatible Management Agent in a silent mode, invoke the agentca utility using the -f and -c option from the <Agent Oracle home>/bin directory of the standalone Management Agent on each cluster node. Also use the -n option to specify the name of the cluster:

For example:

<Agent Oracle Home>/bin/agentca -f -n <cluster name> -c "{<comma separated cluster node list like node1, node 2…>}"

J

Manually Downloading Prerequisites from My Oracle Support

This appendix explains how you can manually download the provisioning Deployment Procedure-specific prerequisites from *My Oracle Support* when Enterprise Manager Grid Control does not have connectivity to *My Oracle Support* or has firewall settings that prevent it from connecting to *My Oracle Support*.

In particular, this appendix covers the following:

- Downloading Prerequisites from My Oracle Support
- Uploading Prerequisites to Oracle Software Library

J.1 Downloading Prerequisites from My Oracle Support

Table J–1 shows the URL from where you can download the prerequisite-specific ZIP file for each of the platforms, and how you must rename the downloaded ZIP file. Note that downloading the ZIP file from *My Oracle Support* requires *My Oracle Support* credentials. After downloading and renaming the ZIP file, upload the ZIP file to Oracle Software Library (Software Library) as described in Section J.2.

Table J–1   Downloading Prerequisites from My Oracle Support

Platform	URL to Download the ZIP File	Rename the ZIP File to
All Platforms *(Generic)*	`https://updates.oracle.com/Orion/Services/download/p8505528_111000_GENERIC.zip?aru=12173379&patch_file=p8505528_111000_GENERIC.zip`	`cvuprereq_2000.zip`
Linux X32	`https://updates.oracle.com/Orion/Services/download/p9288873_111000_LINUX.zip?aru=12113360&patch_file=p9288873_111000_LINUX.zip`	`cvuprereq_46.zip`
Linux.X64	`https://updates.oracle.com/Orion/Services/download/p9288873_111000_Linux-x86-64.zip?aru=12113437&patch_file=p9288873_111000_Linux-x86-64.zip`	`cvuprereq_226.zip`
Sparc64	`https://updates.oracle.com/Orion/Services/download/p9288873_111000_SOLARIS64.zip?aru=12113438&patch_file=p9288873_111000_SOLARIS64.zip`	`cvuprereq_23.zip`

Table J–1 (Cont.) Downloading Prerequisites from My Oracle Support

Platform	URL to Download the ZIP File	Rename the ZIP File to
AIXPPC64	`https://updates.oracle.com/Orion/Services /download/p9288873_111000_ AIX64-5L.zip?aru=12113597&patch_ file=p9288873_111000_AIX64-5L.zip`	`cvuprereq_ 212.zip`
HPIA64	`https://updates.oracle.com/Orion/Services /download/p9288873_111000_ HPUX-IA64.zip?aru=12113526&patch_ file=p9288873_111000_HPUX-IA64.zip`	`cvuprereq_ 197.zip`
HPUX PARISC64	`https://updates.oracle.com/Orion/Services /download/p9288873_111000_ HP64.zip?aru=12283071&patch_ file=p9288873_111000_HP64.zip`	`cvuprereq_ 59.zip`
WINDOWS32	`https://updates.oracle.com/Orion/Services /download/p9288873_111000_ WINNT.zip?aru=12301250&patch_ file=p9288873_111000_WINNT.zip`	`cvuprereq_ 912.zip`
WINDOWS64	`https://updates.oracle.com/Orion/Services /download/p9288873_111000_ MSWIN-x86-64.zip?aru=12301249&patch_ file=p9288873_111000_MSWIN-x86-64.zip`	`cvuprereq_ 233.zip`

J.2 Uploading Prerequisites to Oracle Software Library

After downloading and renaming the prerequisite-specific ZIP file, upload the ZIP file to the Software Library. To do so, follow these steps:

1. In Grid Control, click **Deployments**, then **Provisioning**.

2. On the Provisioning page, in the **Components** tab, from the table, expand **Oracle Components**.

3. Expand **CVU Prerequisite-fixup components**.

4. Expand the platform for which you want to upload the prerequisite-specific ZIP file. For example, `linux_x64` for Linux x64 platform-specific ZIP file, or `all_ platforms` for generic ZIP file.

5. Under the platform you expanded, select **CVU OSD Prerequisite-Fixup component**. If you expanded all_platforms, then select **CVU Generic Prerequisite-Fixup component**.

6. Click **Edit**.

7. On the Edit Component page, click the **Upload File** tab and upload the prerequisite-specific ZIP file.

8. After uploading the ZIP file, click **Finish**.

9. On the Provisioning page, in the **Components** tab, from the table, locate the row corresponding to your upload, and ensure that is still selected. For example, if you uploaded a ZIP file to CVU Generic Prerequisite-Fixup component, which is under all_platforms, then ensure that CVU Generic Prerequisite-Fixup component is still selected.

   You will see that the **Status** column for this component shows **Incomplete**, and the **Maturity** column shows **Untested**.

10. Click **Promote**. The **Maturity** column shows **Beta**.

11. Click **Promote** again. The Maturity column now shows **Production**.

K

Staging Gold Image to Software Library or Staging Location

This chapter describes how you can stage a gold image to either Oracle Software Library (Software Library) or a shared location. In particular, this chapter covers the following:

- Staging Gold Image to Software Library
- Staging Gold Image to Shared Location

K.1 Staging Gold Image to Software Library

To create a gold image and stage it to the Software Library, follow these steps:

> **Note:** You can use this procedure for creating a gold image of an existing Oracle Database (standalone), Oracle RAC Database, Oracle Grid Infrastructure, or Oracle Database Replay Client.

1. In Grid Control, click **Deployments**, and then click the **Provisioning** subtab.

2. On the Components page, from the table, select **Components** and click **Create Component**. Enterprise Manager Grid Control displays the Create Component Wizard.

3. On the Create Component: Describe page, from the **Type** list, select a type from Table K–1. Specify a unique name, a description (for example, *This is Database Gold Image with CPU Patches Applied*), the product/patch number, the version (for example, 11g or 10.1.0.5), and the vendor details.

> **Note:** If you are creating a gold image of Oracle Grid Infrastructure, then from the **Type** list, select **Oracle Clusterware Clone**.

Table K–1   Selecting Component Name

Component Name	Description
Oracle Database Software Clone	Creates gold image for the Oracle Database 10g Release 1 (10.1.x.x), Oracle Database 10g Release 2 (10.2.x.x), Oracle Database 11g Release 1 (11.1.x.x) ,Oracle Database 11g Release 2 (11.2.x.x).
	Note: This database gold image is used even during RAC provisioning.

Table K–1   (Cont.) Selecting Component Name

Component Name	Description
Oracle Clusterware Clone	Creates gold image for Oracle Clusterware 10g Release 2 (10.2.x.x), Oracle Clusterware 11g Release 1 (11.1.x.x), Oracle Grid Infrastructure 11g Release 2 (11.2.x.x).
Oracle Application Server	Creates gold image for Oracle Fusion Middleware 11g, Oracle SOA Suite 10.1.3.x , Oracle Web Tier 10.1.3, and Oracle Application Tier (OHS and OC4J) 10.1.3, J2EE and Oracle Web Cache 10.1.2, Oracle Forms and Reports 9.0.4.

4. Click **Next**.

5. On the Create Component: Configure page, specify the configuration details such as the host name, the credentials of the host, and the privilege type you want to set. The privilege type is for the destination host credentials. For example, if you have SUDO configured , then select the SUDO option. For more information, click **Help** on the page.

    By default, the page displays the working directory (on the destination host where you have the database home) and a list of files you can exclude during this provisioning process. If you want to set another location for the working directory and modify the list of files to exclude, then edit the fields. Ensure that the working directory has adequate space. For example, if the Oracle home if 1 GB, then the working directory must have more than 1 GB.

6. Click **Next**.

7. On the Review page, review all the information you have provided and click **Finish** to submit the component creation job.

    > **Note:**   After you submit the job, Enterprise Manager Grid Control creates the gold image and returns to the Components page. By default, the job has an ID in the format CREATEDBCLONECOMPONENT <day month date time>. For example, `CREATEDBCLONECOMPONENT MON APR 12 09:22:03 PDT 2010`. Enterprise Manager Grid Control displays this job ID on the Components page. You can click this job ID to view its status.

8. On the Components page, in the table, ensure that the gold image you just created is appearing under **Oracle Components** and its status is *Ready*.

    > **Note:**   If the component creation job fails, then click the **Jobs** tab. On the Job Activity page, select the job that failed and click **View Results**. On the Job Run page, click the step, which caused the job to fail, and view the output log that describes the cause of the failure. Correct the error and rerun the job.

K.2 Staging Gold Image to Shared Location

This section describes how you can stage a gold image to a shared location. In particular, this section covers the following:

- Manually Staging Gold Image to Shared Location

- Running Deployment Procedure to Stage Gold Image to Shared Location

K.2.1 Manually Staging Gold Image to Shared Location

To manually stage a gold image to a shared location, simply use any file transfer utility to manually copy the gold image to a staging location.

K.2.2 Running Deployment Procedure to Stage Gold Image to Shared Location

To run the Deployment Procedure only to stage a gold image in a staging location, follow these steps:

1. In Grid Control, click **Deployments**.

2. On the Deployment page, from the Deployment Procedure Manager section, click **Database Provisioning Procedures**.

3. On the Deployment Procedure Manager page, from the procedures table, select one of the following Deployment Procedures, and click **Create Like**.

   - *Oracle Database Provisioning*, for standalone Oracle Database.

   - *Oracle Grid Infrastructure Provisioning for Standalone Servers*, for Oracle Grid Infrastructure for standalone servers.

   - *Oracle Grid Infrastructure/RAC Provisioning*, for Oracle Grid Infrastructure for clustered environments.

   - *Oracle Clusterware/RAC Provisioning*, for Oracle Clusterware.

4. On the Create Like Procedure page, do the following:

   a. For Name, provide a name for the custom Deployment Procedure you are creating.

   b. From the table that lists all the phases and steps of the Deployment Procedure, select all the rows and click **Disable**.

   c. Enable the phases and steps as described in Table K–2. To enable a phase or step, select the row of a phase or step and click **Enable**.

5. Click **Save**.

   Enterprise Manager Grid Control saves the custom Deployment Procedure you created and returns to the Deployment Procedure Manager page.

6. On the Deployment Procedure Manager page, select the custom Deployment Procedure you created and click **Schedule Deployment**.

   Based on the settings you made in Step 4 (c), select the appropriate option - either using an existing installation or using the Software Library.

7. In the custom Deployment Procedure wizard, provide all the details that are required. Ensure that you expand the Advanced portion of the Specify Destination Host Settings section, select **Stage to Shared Location**, and provide the full path to the staging location.

8. After you provide all the required information, click **Finish**.

   Enterprise Manager Grid Control runs the Deployment Procedure to stage the gold image to the staging location. The Deployment Procedure does not actually provision the gold image to the destination hosts because the phases and steps to do so are disabled.

Table K–2   Enabling Phases/Stes in Deployment Procedure to Stage Gold Image

Deployment Procedure You Are Customizing	Phases/Steps to Enable or Disable
■ Oracle Database Provisioning ■ *Oracle Grid Infrastructure Provisioning for Standalone Servers*	■ Initialize Deployment Procedure ■ Create New Stage Location When you enable the phase **Create new stage location**, you will see two more steps, mainly *Create new stage location from Installed home* and *Create new stage location from Software Library*. After customizing the Deployment Procedure, when you run it, if you want to select the existing installation option to provision the database and also to create a gold image of it and store it in a shared location, then enable the step *Create new stage location from Installed home* and disable the step *Create new stage location from Software Library*. However, if you want to select the Software Library option to provision the database using a gold image present in the Software Library, and also to create a gold image of it and store it in a sharedlocation, then then disable the step *Create new stage location from Installed home* and enable the step *Create new stage location from Software Library*. ■ Destination User Permission Checks ■ Agent User Permission Checks ■ Reference Host Permission Checks
Oracle Grid Infrastructure/RAC Provisioning	■ Initialize Deployment Procedure ■ Clusterware User Permission Checks ■ RAC User Permission Checks ■ Agent User Permission Checks ■ Reference User Permission Checks ■ Prerequisites and Fix-ups ■ Clusterware Archive ■ Prepare Clusterware Stage Area ■ Stage Clusterware Archive ■ RAC Archive ■ Prepare Database Stage Area ■ Stage RAC Archive

Table K–2   (Cont.) Enabling Phases/Stes in Deployment Procedure to Stage Gold Image

Deployment Procedure You Are Customizing	Phases/Steps to Enable or Disable
Oracle Clusterware/RAC Provisioning For UNIX	■ Initialize Deployment Procedure ■ Clusterware User Permission Checks ■ ASM User Permission Checks ■ RAC User Permission Checks ■ Agent User Permission Checks ■ Reference User Permission Checks ■ Prerequisites and Fix-ups ■ Clusterware Archive ■ Prepare Clusterware Stage Area ■ ASM Archive ■ Prepare Database Stage Area (this for the ASM archives) ■ RAC Archive ■ Prepare Database Stage Area (this is for the RAC archives)

L

Downloading Installation Binaries and Uploading to Software Library or Staging Location

This chapter describes how you can download the installation binaries of the following, and upload them to Oracle Software Library (Software Library) or a shared location:

- Standalone Oracle Database
- Oracle Database Replay Client
- Oracle Grid Infrastructure for standalone servers
- Oracle Grid Infrastructure and Oracle RAC for clustered environments

In particular, this chapter covers the following:

- Downloading Installation Binaries
- Uploading Installation Binaries to Software Library or Shared Location

L.1 Downloading Installation Binaries

For information about downloading the installation binaries of Oracle Database and Oracle Database Replay Client, see *My Oracle Support* note 815567.1.

For information about downloading the installation binaries of Oracle Grid Infrastructure for standalone servers, and Oracle Grid Infrastructure and Oracle RAC for clustered environments, see *My Oracle Support* note 1083691.1.

L.2 Uploading Installation Binaries to Software Library or Shared Location

This section describes how you can upload the installation binaries to either the Software Library or a shared location. In particular, this section covers the following:

- Uploading Installation Binaries to Software Library
- Uploading Installation Binaries to Shared Location

L.2.1 Uploading Installation Binaries to Software Library

To upload the installation binaries to the Software Library, follow these steps:

1. In Grid Control, click **Deployments**, then **Provisioning**.

2. On the Components page, from the table, expand **Oracle Components**, and then **RAC Provisioning**.

3. Under RAC Provisioning, expand the release of the database. For example, 10.2.0.1.0.

4. Under the release of the database, expand the operating system.

5. Under the operating system, select one of the following and click **Edit**.

- For 11.1.0.6.0 or lower:

  – For Oracle Clusterware software binaries, select **Oracle Clusterware shiphome**.

  – For Oracle Database software binaries, select **Oracle Database shiphome**.'

  – For Oracle Database Replay Client software binaries, select **Oracle Database Replay Client Install Media**.

- For 11.2.0.1.0 or higher:

  – For Oracle Grid Infrastructure software binaries, select **Oracle Grid Infrastructure shiphome**.

  – For Oracle Database software binaries, select **Oracle Database shiphome**.

  – For Oracle Database Replay Client software binaries, select **Oracle Database Replay Client Install Media**.

6. On the Edit Component page, click the **Upload File** tab.

7. On the Upload File page, select one of the following options and upload the software binaries:

   a. **Upload from Agent Machine**, if the software binaries are available on the host where the Management Agent is running.

   b. **Upload from Local Machine**, if the software binaries are available on the local host, that is, the host from where you are accessing the Grid Control console.

8. Click **Finish**.

L.2.2 Uploading Installation Binaries to Shared Location

To upload the installation binaries to a shared location, simply use any file transfer utility to manually copy the installation binaries to a staging location.

M

Reviewing Log Files

This appendix lists the log files you must review to resolve issues encountered while running a Deployment Procedure.

M.1 OMS-Related Log Files

The following are OMS-related log files.

- **Generic EM Trace File**

  `$<ORACLE_HOME>/sysman/log/emoms.trc`

- **Provision Advisory Framework (PAF) Llogs**

  `$<ORACLE_HOME>/ sysman/log/pafLogs/`

- **For Specific Deployment Procedure Instance**

  `$<ORACLE_HOME>/sysman/log/pafLogs/<name>_<instance_guid>.log`

> **Note:** The `$<ORACLE_HOME>` refers to the Oracle home of the OMS.

M.2 Management Agent-Related Log Files

The following are Management Agent-related log files:

`$<ORACLE_HOME>/sysma/logs/emagent.nohup`

`$<ORACLE_HOME>/sysma/logs/emagent.trc`

> **Note:** The `$<ORACLE_HOME>` refers to the Oracle home of the Management Agent.

M.3 Advanced Options

Optionally, to capture more details, follow these steps to reset the log level and capture the logs mentioned in the previous sections.

> **Note:** Oracle recommends you to archive the old logs and have a fresh run after resetting the log level to capture the fresh logs.

- On the OMS side:

  1. Open the following file available in the Oracle home of the OMS:

     `$<ORACLE_HOME>/sysman/config/emomslogging.properties`

  2. Set the `@ log4j.rootCategory=` parameter to `DEBUG`.

  3. Restart the OMS by running the following commands:

     `$<ORACLE_HOME>/bin/emctl stop oms`

     `$<ORACLE_HOME>/bin/emctl start oms`

- On the Management Agent side:

  1. Open the following file available in the Oracle home of the Management Agent:

     `$<ORACLE_HOME>/sysman/config/emd.properties`

  2. Make the following settings:

     `tracelevel.Dispatcherr=DEBUG`

If Installation or Cloning Steps Fail

The OUI/Clone steps generate installation and error log files, and since they are generated at runtime, their locations are printed in the job logs.

If Database Configuration Failed

Database Configuration Assistant generates a log file, and its location is printed in the job log.

As part of target discovery, we generate the `newtargets.xml` file where the newly provisioned targets are added. The following is the location. Here, `$<ORACLE_HOME>` refers to the Oracle home of the Management Agent.

`$<ORACLE_HOME>/sysman/emd/newtargets.xml`

For Oracle Grid Infrastructure for standalone servers, the location is:

`$<ORACLE_HOME>/sysman/emd/hastargets.xml`

Index