

Oracle® Enterprise Manager Lifecycle Management Administrator's Guide



13c Release 3
F19095-01
June 2019

ORACLE®

Copyright © 2015, 2019, Oracle and/or its affiliates. All rights reserved.

Primary Author: Oracle Corporation

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

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

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

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

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

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

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

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

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

Contents

Preface

Audience	xlii
Scope and Coverage	xlii
Documentation Accessibility	xlii
Related Documents	xlili
Conventions	xlili

Part I Overview and Setup Details

1 Introduction to Lifecycle Management

Overview of the New Lifecycle Management Solutions	1-1
Information Map for Lifecycle Management Solutions	1-5

2 Setting Up Your Infrastructure

Getting Started with Setting Up Your Infrastructure	2-1
Setting Up Oracle Software Library	2-2
Setting Up Credentials	2-4
Creating Enterprise Manager User Accounts	2-6
Overview of User Accounts	2-6
Creating Designer User Account	2-8
Creating Operator User Account	2-9
(Optional) Setting Up My Oracle Support	2-10
(Optional) Configuring Self-Update	2-10
(Optional) Setting Up E-mail Notifications	2-11
(Optional) Setting Restricted Accesses for the Root Components	2-11
Root Components for Patching	2-11
Manually Staging the Root Components	2-11
Restricting the Root User Access	2-12
Root Components for Provisioning	2-13
Manually Staging the Root Components	2-13

Part II Discovery

3 Discovering Hosts and Software Deployments

Discovering Hosts Automatically and Adding Targets Manually	3-1
Discovering Hosts Manually and Adding Targets Manually	3-1

Part III Database Provisioning

4 Overview of Database Provisioning

Introduction to Database Provisioning	4-1
Supported Use Cases and Targets Provisioned Using Database Provisioning Procedures	4-2
Setting Up Database Provisioning	4-5
Meeting Basic Infrastructure and Host Requirements	4-6
Understanding Administrator Privileges for Provisioning Database	4-6
Prerequisites for Designers	4-6
Prerequisites for Operators	4-8
Creating Database Provisioning Profiles	4-8
Viewing, Creating, and Deleting Database Provisioning Profiles Using EMCLI	4-10
Viewing Database Provisioning Profiles Using EMCLI	4-11
Creating Database Provisioning Profiles Using EMCLI	4-11
Deleting Database Provisioning Profiles Using EMCLI	4-12
Creating Installation Media	4-12
Creating Database Templates	4-13
Uploading Database Templates to Software Library	4-14
Creating Database Provisioning Entities	4-15
Creating an Entity for Oracle Database Software from a Shared NFS location	4-15
Creating an Oracle Database Clone from an External Storage	4-17
Creating an Oracle Clusterware Clone from a Reference Home	4-18
Creating an Oracle Clusterware Clone from an External Storage	4-19
Downloading Cluster Verification Utility	4-20

5 Creating and Provisioning Oracle Databases

Oracle Database Topology	5-1
Creating Databases	5-3

Getting Started with Creating Databases	5-3
Creating an Oracle Database	5-4
Prerequisites for Creating an Oracle Database	5-4
Procedure for Creating an Oracle Database	5-5
Creating Oracle Real Application Clusters Database	5-9
Prerequisites for Creating an Oracle Real Application Clusters Database	5-9
Procedure for Creating an Oracle Real Application Clusters Database	5-10
Creating Oracle Real Application Clusters One Node Database	5-13
Prerequisites for Creating an Oracle RAC One Node Database	5-13
Procedure for Creating an Oracle Real Application Clusters One Node Database	5-14
Provisioning Databases	5-16
Getting Started with Provisioning Oracle Databases	5-16
Provisioning and Creating Oracle Databases	5-17
Prerequisites for Provisioning Databases	5-18
Procedure for Provisioning Databases	5-18
Provisioning Oracle Databases with Oracle Automatic Storage Management	5-23
Prerequisites for Provisioning Oracle Databases with Oracle Automatic Storage Management	5-23
Procedure for Provisioning Databases	5-23
Provisioning Oracle Database Software Only	5-28
Prerequisites for Provisioning Oracle Database Software Only	5-28
Procedure for Provisioning Oracle Database Software Only	5-29
Provisioning Oracle Database Software from a Shared NFS location	5-31
Using Non-Root Credentials for Provisioning Oracle Databases	5-32

6 Provisioning Oracle Grid Infrastructure for Oracle Databases

Getting Started with Provisioning Oracle Grid Infrastructure for Oracle Databases	6-1
Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management	6-2
Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM	6-2
Procedure for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM	6-3
Provisioning Oracle Grid Infrastructure and Oracle Database Software Only	6-8
Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only	6-8
Procedure for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only	6-9

7 Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases

Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases	7-1
Oracle Real Application Clusters Database Topology	7-3
Oracle Flex Clusters and Flex ASM	7-4
Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management	7-5
Prerequisites for Provisioning Grid Infrastructure with Oracle RAC Database	7-5
Procedure for Provisioning Grid Infrastructure with Oracle RAC Database	7-5
Requirements for Grid Infrastructure Software Location Path	7-13
Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster	7-13
Prerequisites for Provisioning Oracle RAC Database with File System on an Existing Cluster	7-13
Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster	7-13
Provisioning Oracle Real Application Clusters Database with File System on a New Cluster	7-18
Prerequisites for Provisioning Oracle RAC Database with File System on a New Cluster	7-18
Procedure for Provisioning Oracle RAC Database with File System on a New Cluster	7-18
Using No Root Credentials for Provisioning Oracle Real Application Clusters (Oracle RAC) Databases	7-26

8 Provisioning Oracle Real Application Clusters One (Oracle RAC One) Node Databases

Getting Started with Provisioning Oracle RAC One Node Databases	8-1
Deployment Procedures for Provisioning Oracle RAC One Node Databases	8-2
Provisioning Oracle RAC One Node Databases	8-2
Prerequisites for Provisioning Oracle RAC One Node Databases	8-2
Procedure for Provisioning Oracle RAC One Node Databases	8-4

9 Provisioning Oracle Real Application Clusters for 12c and 18c

Getting Started with Provisioning Oracle Real Application Clusters for 12c and 18c	9-1
Core Components Deployed When Provisioning Oracle RAC	9-2
Cloning a Running Oracle Real Application Clusters	9-3
Prerequisites for Cloning a Running Oracle Real Application Clusters	9-3
Procedure for Cloning a Running Oracle Real Application Clusters	9-4
Provisioning Oracle Real Application Clusters Using Gold Image	9-10

Prerequisites for Provisioning Oracle Real Application Clusters Using Gold Image	9-10
Procedure for Provisioning Oracle Real Application Clusters Using Gold Image	9-11
Provisioning Oracle Real Application Clusters Using Archived Software Binaries	9-17
Prerequisites for Provisioning Oracle Real Application Clusters Using Archived Software Binaries	9-17
Procedure for Provisioning Oracle Real Application Clusters Using Archived Software Binaries	9-18
Sample Cluster Configuration File	9-25
Provisioning Oracle Real Application Clusters (Oracle RAC) Databases Using No Root Credentials	9-25

10 Extending Oracle Real Application Clusters

Getting Started with Extending Oracle Real Application Clusters	10-1
Extending Oracle Real Application Clusters	10-1
Prerequisites for Extending Oracle Real Application Clusters	10-2
Procedure for Extending Oracle Real Application Clusters	10-2

11 Deleting or Scaling Down Oracle Real Application Clusters

Getting Started with Deleting or Scaling Down Oracle Real Application Clusters	11-1
Deleting the Core Components of Oracle Real Application Clusters	11-2
Deleting the Entire Oracle RAC	11-2
Prerequisites for Deleting the Entire Oracle RAC	11-3
Procedure for Deleting the Entire Oracle RAC	11-3
Scaling Down Oracle RAC by Deleting Some of Its Nodes	11-6
Prerequisites for Scaling Down Oracle RAC by Deleting Some of Its Nodes	11-6
Procedure for Scaling Down Oracle RAC by Deleting Some of Its Nodes	11-6

12 Provisioning Oracle Database Replay Client

Getting Started with Provisioning Oracle Database Replay Client	12-1
Cloning a Running Oracle Database Replay Client	12-2
Prerequisites for Cloning a Running Oracle Database Replay Client	12-2
Procedure for Cloning a Running Oracle Database Replay Client	12-3
Provisioning an Oracle Database Replay Client Using Gold Image	12-5
Prerequisites for Provisioning an Oracle Database Replay Client Using Gold Image	12-5
Procedure for Provisioning an Oracle Database Replay Client Using Gold Image	12-6
Provisioning an Oracle Database Replay Client Using Installation Binaries	12-9

Prerequisites for Provisioning an Oracle Database Replay Client Using Installation Binaries	12-9
Procedure for Provisioning an Oracle Database Replay Client Using Installation Binaries	12-11

13 Provisioning Oracle Standby Databases

Overview of Creating a Standby Database	13-1
Creating a New Physical Standby Database (single-instance only)	13-1
Step 1: Determine the backup type	13-3
Step 2: Set up the backup options	13-3
Step 3: Select the Oracle home in which to create the standby database	13-3
Step 4: Set up the location for standby database files	13-4
Step 5: Provide standby database configuration parameters	13-4
Step 6: Review the information before clicking Finish	13-4
Creating a New Logical Standby Database (single-instance only)	13-5
Step 1: Determine the backup type	13-6
Step 2: Set up the backup options	13-6
Step 3: Select the Oracle home in which to create the standby database	13-7
Step 4: Set up the location for standby database files	13-7
Step 5: Provide standby database configuration parameters	13-7
Step 6: Review the information before clicking Finish	13-8
Managing an Existing Standby Database with Data Guard Broker	13-8
Creating a Primary Database Backup Only	13-9

14 Managing Pluggable Databases Using Enterprise Manager

Getting Started With Managing Pluggable Databases Using Enterprise Manager	14-1
Overview of Managing Pluggable Databases Using Enterprise Manager	14-2
Provisioning Pluggable Databases Using Enterprise Manager	14-3
Creating a New Pluggable Database Using Enterprise Manager	14-4
Prerequisites for Creating a New Pluggable Database	14-4
Creating a New Pluggable Database	14-4
Plugging In an Unplugged Pluggable Database Using Enterprise Manager	14-10
Prerequisites for Plugging In an Unplugged Pluggable Database	14-10
Plugging In an Unplugged Pluggable Database	14-11
Cloning a Pluggable Database Using Enterprise Manager	14-18
Prerequisites for Cloning a Pluggable Database	14-18
Cloning a Pluggable Database	14-19
Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager	14-27
Prerequisites for Migrating a Non-CDB as a Pluggable Database	14-27
Migrating a Non-CDB as a Pluggable Database	14-28

Relocating Pluggable Databases Using Enterprise Manager	14-31
Prerequisites for Relocating Pluggable Databases	14-32
Relocating Pluggable Databases	14-32
Removing Pluggable Databases Using Enterprise Manager	14-33
Unplugging and Dropping a Pluggable Database Using Enterprise Manager	14-33
Prerequisites for Unplugging and Dropping a Pluggable Database	14-34
Unplugging and Dropping a Pluggable Database	14-34
Deleting Pluggable Databases Using Enterprise Manager	14-38
Prerequisites for Deleting Pluggable Databases	14-38
Deleting Pluggable Databases	14-38
Viewing Pluggable Database Job Details Using Enterprise Manager	14-42
Viewing Create Pluggable Database Job Details	14-42
Viewing Unplug Pluggable Database Job Details	14-44
Viewing Delete Pluggable Database Job Details	14-45
Administering Pluggable Databases Using Enterprise Manager	14-46
Switching Between Pluggable Databases Using Enterprise Manager	14-46
Altering Pluggable Database State Using Enterprise Manager	14-47

15 Cloning Oracle Databases and Pluggable Databases

Creating a Full Clone Database	15-1
Creating a Full Clone Database Using the Clone Wizard	15-1
Creating a Full Clone Database Using EM CLI	15-5
Refresh a Cloned Database	15-5
Creating a Full Clone Pluggable Database	15-6
Creating a Full Clone Pluggable Database Using the Clone Wizard	15-6
Creating a Full Clone Pluggable Database Using EM CLI	15-8
Cloning Databases Using the Classic Cloning Wizard	15-9
Overview of Classic Cloning Methods	15-9
Cloning an Oracle Database Using Recovery Manager (RMAN) Backup	15-10
Cloning an Oracle Database Using Staging Areas	15-12
Cloning an Oracle Database Using an Existing Backup	15-14

16 Database High Availability REST API

Data Guard	16-1
Backup and Recovery	16-5

Part IV Hybrid Cloud Management

17 Hybrid Solution: Oracle Cloud Infrastructure Management

Hybrid Cloud Cloning Overview	17-1
Cloning in Hybrid Cloud Use Cases	17-1
Prerequisites for Cloning in Oracle Cloud	17-2
Cloning to Oracle Cloud	17-3
Cloning a DB to Oracle Cloud (Oracle Compute Service)	17-3
Cloning a DB to Oracle Cloud Using the Clone Wizard	17-3
Cloning a DB to Oracle Cloud Using EMCLI	17-6
Cloning a PDB to Oracle Cloud	17-8
Cloning a PDB to Oracle Cloud Using the Clone Wizard	17-8
Cloning a PDB to Oracle Cloud Using EM CLI	17-9
Cloning Schema(s) to a DB or PDB on Oracle Cloud	17-11
Cloning a DB to a PDB on Oracle Cloud	17-14
Cloning to Oracle Cloud Using Oracle Storage Service	17-17
Cloning a Database to Oracle Cloud Using the Clone Wizard (OSS)	17-17
Cloning from Oracle Cloud	17-20
Cloning a PDB from Oracle Cloud	17-20
Cloning a PDB from Oracle Cloud Using the Clone Wizard	17-21
Cloning a PDB from Oracle Cloud Using EM CLI	17-22
Cloning Schema(s) from Oracle Cloud to a DB or PDB	17-24
Cloning a DB from Oracle Cloud to an On-Premise PDB	17-26
Cloning a DB from Oracle Cloud (Oracle Compute Service)	17-29
Cloning a DB from Oracle Cloud Using the Clone Wizard	17-29
Cloning a DB from Oracle Cloud Using EMCLI	17-31
Cloning Within Oracle Cloud	17-33
Cloning a PDB Within Oracle Cloud	17-33
Creating a Full Clone Pluggable Database Using the Clone Wizard	17-33
Creating a Full Clone Pluggable Database Using EM CLI	17-35
Cloning a DB Within Oracle Cloud (Oracle Compute Service)	17-36
Cloning a DB Within Oracle Cloud Using the Clone Wizard	17-36
Cloning a DB Within Oracle Cloud Using EMCLI	17-36

18 Hybrid Solution: Oracle Cloud Platform Management

About Hybrid Cloud Management	18-1
Viewing and Monitoring Database Cloud Services	18-2
Typical Hybrid Cloud Management Workflow	18-2
Setting Up Hybrid Cloud	18-3
Setting Up Hybrid Cloud Using the Hybrid Setup Console	18-3
Registering the Identity Domain and the Cloud Account	18-4
Configuring the Hybrid Cloud Gateway	18-4

Configuring the Proxy Settings	18-5
Configuring the SSH Keys	18-5
Setting Up Hybrid Cloud Using REST API	18-6
Setting Up Hybrid Cloud Using EMCLI	18-9
Defining the Oracle Database Cloud Service Settings	18-13
Discovering the Database Cloud Service	18-15
Managing Data Sources	18-16
Creating a Data Source	18-16
Adding a Data Source	18-19
Creating Backups On the Oracle Storage Server	18-20
Defining the Database Backup Cloud Service Settings	18-20
Monitoring the Discovered Database Cloud Service	18-20
Refreshing the Cloud Database Service	18-21
Clone Management	18-21
Using Snapshots to Create Snap Clones	18-22
Using Database Backups to Create Full Clones	18-22
Defining the Oracle Compute Service Settings	18-22
Discovering the Oracle Compute Service	18-25

Part V Database Upgrade

19 Upgrading Databases

Getting Started	19-1
Supported Releases	19-2
Database Upgradability Report	19-3
Upgrading Databases Using Deployment Procedure	19-4
About Deployment Procedures	19-5
Meeting the Prerequisites	19-6
Upgrading Oracle Cluster Database Using Deployment Procedure	19-8
Upgrading Oracle Clusterware Using Deployment Procedure	19-13
Upgrading Oracle Database Instance Using Deployment Procedure	19-17
Upgrading an Oracle Database or Oracle RAC Database Instance Using the Database Upgrade Wizard	19-21
Meeting the Prerequisites	19-22
Performing the Upgrade Procedure	19-22
Data Guard Rolling Upgrade	19-26
About Rolling Upgrades	19-27
Prerequisites to Rolling Upgrades	19-27
Submitting a Rolling Upgrade Procedure for a Primary Database With One Physical Standby Database	19-28

Part VI Database Security

20 Managing Oracle Audit Vault and Database Firewall

21 Using Oracle Data Redaction

22 Managing Oracle Database Vault and Privilege Analysis

23 Managing Oracle Transparent Data Encryption

Prerequisites for Transparent Data Encryption Operations	23-2
Encrypting a Tablespace in Online Mode	23-3
Decrypting a Tablespace in Online Mode	23-4
Changing the Encryption Algorithm of an Encrypted Tablespace - Rekey	23-5
Encrypting a Tablespace in Offline Mode	23-5
Decrypting a Tablespace in Offline Mode	23-6

Part VII Middleware Provisioning

24 Overview of Middleware Provisioning

Introduction to Middleware Provisioning	24-1
Oracle Fusion Middleware Provisioning Terminology	24-4
Supported Use Cases for Middleware Provisioning Procedures	24-6
Provisioning Middleware Domains and Oracle Homes	24-7
Scaling Up/Out/ WebLogic Server, SOA, Service Bus, and WebCenter Domains	24-15
Deploying / Redeploying / Undeploying Java EE Applications	24-21
Provisioning Coherence Nodes and Clusters	24-22
Provisioning SOA Artifacts	24-22
Provisioning Service Bus Resources	24-23

25 Provisioning Fusion Middleware Domain and Oracle Homes

Getting Started with Fusion Middleware Provisioning	25-2
Different Approaches to Launch the Provision Fusion Middleware Deployment Procedure	25-2
High-Level Steps for Middleware Provisioning	25-3
Step1: Creating a Profile	25-3
Step2: Running Provision Fusion Middleware Procedure to Provision the Profile	25-4
Prerequisites for Provisioning from the Middleware Provisioning Profiles	25-5
Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile	25-5
Prerequisites for Provisioning the WebLogic Domain Profile	25-6
Using Custom Scripts Stored in the Software Library	25-7
Using Custom Scripts with Input Parameters	25-8
Using Custom Scripts Without Inputs Parameters	25-9
Creating Middleware Provisioning Profiles	25-10
Creating a Provisioning Profile Based on an Installation Media	25-10
Creating a Provisioning Profile Based on an Oracle Home	25-12
Creating a Provisioning Profile Based on a WebLogic Domain	25-13
Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile	25-16
Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile	25-21
Provisioning a Fusion Middleware Domain from an Existing Oracle Home	25-29
Customizing the Destination Environment from an Existing Oracle Home	25-33
Cloning from an Existing WebLogic Domain Based-Profile	25-39
Customizing the Destination Environment from an Existing WeLogic Domain Based-Profile	25-44

26 Provisioning the SOA Domain and Oracle Homes

Getting Started with Provisioning SOA Domain and Oracle Home	26-2
Source Environment and Destination Environment after SOA Provisioning	26-2
Source and Destination Environments for a Fresh SOA Provisioning Use Case	26-3
Source and Destination Environments for SOA Cloning Use Case	26-4
Supported Versions of SOA for Provisioning	26-4
Before you Begin Provisioning SOA Domain and Oracle Home	26-5
Create Middleware Roles and Assign Privileges to them	26-5
Setting Named Credentials and Privileged Credentials for the Middleware Targets	26-6
(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database	26-6
Use Case 1: First Time Provisioning of a SOA Domain	26-6

Use Case 2: Provisioning from a SOA Oracle Home Based Provisioning Profile	26-7
Use Case 3: Cloning from a Provisioning Profile based on an Existing SOA Domain	26-7
Use Case 4: Provisioning from an Existing SOA Home	26-8
Use Case 5: Scaling Up an Existing SOA Domain	26-8
Use Case 6 - Cloning a database in WebLogic Domain	26-8

27 Provisioning the Service Bus Domain and Oracle Homes

Getting Started with Provisioning Service Bus Domain and Oracle Home	27-1
Supported Versions of Service Bus for Provisioning	27-2
Before you Begin Provisioning Service Bus Domain and Oracle Home	27-3
Create Middleware Roles and Assign Privileges	27-3
Setting Named Credentials and Privileged Credentials for the Middleware Targets	27-3
(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database	27-4
Use Case 1: First Time Provisioning of a Service Bus Domain	27-4
Use Case 2: Provisioning from a Service Bus Home Based Provisioning Profile	27-4
Use Case 3: Cloning from a Provisioning Profile based on an Existing Service Bus Domain	27-5
Use Case 4: Provisioning from an Existing Service Bus Home	27-5
Use Case 5: Scaling Up an Existing Service Bus Domain	27-6

28 Provisioning the Oracle WebCenter Domain and Oracle Homes

Getting Started with Provisioning WebCenter Domain and Oracle Home	28-1
About WebCenter Topologies Supported in Enterprise Manager	28-2
Source Environment and Destination Environment after WebCenter Provisioning	28-4
Source and Destination Environments for a Fresh WebCenter Provisioning Use Case	28-5
Source and Destination Environments for WebCenter Cloning Use Case	28-5
Supported Versions of WebCenter for Provisioning	28-6
Before you Begin Provisioning WebCenter Domain and Oracle Home	28-6
Create Middleware Roles and Assign Privileges to them	28-7
Setting Named Credentials and Privileged Credentials for the Middleware Targets	28-7
(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database	28-8
Use Case 1: First Time Provisioning of a WebCenter Portal with Lock-downs	28-8
Use Case 2: Provisioning a WebCenter Home	28-8
Use Case 3: Cloning an Existing WebCenter Portal Environment	28-9
Use Case 4: Provisioning from an Existing WebCenter Home	28-9

29 Middleware Provisioning using the EM CLI

Creating Middleware Provisioning Profiles	29-1
Creating a WebLogic Domain Profile	29-1
Creating an Oracle Home Profile	29-3
Creating an Installation Media Profile	29-5
Submitting the Procedure using EM CLI	29-6
Listing Middleware Provisioning Profiles	29-7
Listing All the Profiles	29-7
Listing All the WebLogic Domain Profiles	29-8
Listing All the Oracle Home Profiles	29-8
Listing All the Installation Media Profiles	29-9
Describing Provisioning Profiles	29-9
Describing a WebLogic Domain Profile	29-10
Describing an Oracle Home Profile	29-11
Describing an Installation Media Profile	29-12
Describing an Exported Domain/Partiton Profile	29-12
Deleting Provisioning Profiles	29-13
Describing scaledown_domain usage	29-14
Describing deploy_jvmd usage	29-15

30 Middleware Profiles Using REST APIs

Creating Middleware Provisioning Profiles	30-1
Creating a WebLogic Domain Profile	30-1
Creating an Oracle Home Profile	30-3
Creating an Installation Media Profile	30-4
Listing Middleware Provisioning Profiles	30-7
Listing All the Profiles	30-7
Listing WebLogic Domain Profiles	30-8
Listing Oracle Home Profile	30-9
Listing Installation Media Profiles	30-10
Describing Provisioning Profiles	30-10
Describing a WebLogic Domain Profile	30-11
Describing an Oracle Home Profile	30-14
Describing an Installation Media Profile	30-15
Deleting Provisioning Profiles	30-16
Deleting a WebLogic Domain Profile	30-16
Deleting an Oracle Home Profile	30-17

31 Scaling Up / Scaling Out Fusion Middleware Domains

Getting Started	31-1
Prerequisites	31-1
Running the Scale Up / Scale Out Middleware Deployment Procedure	31-3
WebLogic Domain Scaling Up: Scale up/Scale out Fusion Middleware Domain Page	31-4
Weblogic Domain Scaling Up / Scaling Out : Schedule Page	31-7
Middleware Provisioning and Scale Up / Scale Out Best Practices	31-7

32 Export and Import of Domain Partitions

About Exporting and Importing WebLogic Domain Partitions	32-1
Exporting WebLogic Domain Partition	32-2
Importing WebLogic Domain Partition	32-3

33 Migrating a Domain to a Partition

Pre-Requisites and Limitations	33-1
Exporting WebLogic Domain	33-2
Importing WebLogic Domain Partition	33-2

34 Deploying / Redeploying / Undeploying Java EE Applications

Getting Started with the Deploy/Undeploy Java EE Application Deployment Procedure	34-1
Deploying, Undeploying, or Redeploying Java EE Applications	34-2
Supported Releases for Deploy/Undeploy Java EE Applications Deployment Procedure	34-3
Prerequisites for Deploying/Undeploying Java EE Applications	34-3
Creating a Java EE Application Component	34-3
Java EE Applications Deployment Procedure	34-5
Deploying a Java EE Application	34-5
Redeploying a Java EE Application	34-9
Undeploying a Java EE Application	34-12
Deploying / Redeploying / Undeploying Java EE Applications Using EMCLI	34-13
Deploying a Java EE Application Using EMCLI	34-13
Undeploying a Java EE Application Using EMCLI	34-16
Redeploying a Java EE Application Using EMCLI	34-19

35 Provisioning Coherence Nodes and Clusters

Getting Started	35-1
Supported Releases	35-2
Deploying Coherence Nodes and Clusters	35-2
Prerequisites	35-2
Creating a Coherence Component	35-3
Deployment Procedure	35-4
Adding a Coherence Node	35-7
Sample Scripts	35-10
Troubleshooting	35-16

36 Provisioning SOA Artifacts and Composites

Getting Started with SOA Artifacts Provisioning	36-1
Understanding SOA Artifacts Provisioning	36-2
Deployment Procedures, Supported Releases, and Core Components Deployed	36-4
Provisioning SOA Artifacts	36-4
Provisioning SOA Artifacts from a Reference Installation	36-5
Provisioning SOA Artifacts from Gold Image	36-7
Deploying SOA Composites	36-9

37 Provisioning Service Bus Resources

Getting Started with Provisioning Service Bus Resources	37-1
Supported Releases	37-2
Provisioning Service Bus Resources from Service Bus Domain	37-2
Understanding the Export Modes for Service Bus Resources	37-5
Provisioning Service Bus Resources from Oracle Software Library	37-6

Part VIII Bare Metal Server Provisioning

38 Provisioning Bare Metal Servers

Getting Started with Provisioning Bare Metal Servers	38-1
Overview Of Bare Metal Provisioning	38-2
Accessing Bare Metal Provisioning Page	38-3
Provisioning Environment for Bare Metals	38-3
Software Library and its Entities	38-4
Boot Server	38-4
Stage Server	38-4

Reference Host	38-4
RPM Repository	38-4
Provisioning Bare Metal	38-5
Supported Releases of Linux	38-5
Setting Up Infrastructure for Bare Metal Provisioning	38-5
Setting Up Stage Server	38-5
Prerequisites to Setup a Stage Server	38-6
Setting up a Stage Server and Accessing the Management Agent files	38-6
Setting Up Boot Server and DHCP Server	38-8
Setting Up RPM Repository	38-10
Setting Up RHEL 5/Oracle Linux 5 RPM Repository	38-10
Exposing RPM Repository through HTTP or FTP	38-11
Configuring Stage Server	38-11
Configuring Boot Server	38-12
Configuring DHCP Server	38-13
Configuring RPM Repository	38-13
Checklist for Boot Server, Stage Server, RPM Repository, and Reference Host	38-14
Configuring Software Library Components	38-14
Creating Operating System Component	38-15
Creating Disk Layout Component	38-16
Creating an Oracle Virtual Server Component	38-17
Prerequisites For Provisioning Bare Metal Servers and Oracle VM Servers	38-19
Provisioning Bare Metal Servers	38-19
Provisioning Oracle VM Servers	38-23
Viewing Saved Plans	38-26
Using Saved Plans for Provisioning Linux Operating Systems on Bare Metal Servers	38-26

Part IX Host Management

39 Overview of Host Management

Host Statistics	39-1
Diagnosing Host Problems	39-2
Viewing Targets on the Host	39-2
Storage Statistics and History	39-2

40 Setting Up the Environment to Monitor Hosts

Required Installations	40-1
For Linux Hosts - Installing YAST (Only if the host is not patched with latest 13.2)	40-2

Setting Up Credentials	40-3
Setup Needed for Host Monitoring	40-3
Viewing Monitoring Configuration	40-3
Setting Up Monitoring Credentials	40-4
Target Setup Needed for Host Administration	40-4

41 Customizing Your Host Monitoring Environment

Customizing the Host Home Page	41-1
Using Groups	41-2

42 Monitoring Hosts

Overall Monitoring	42-1
CPU Details	42-1
Memory Details	42-2
Disk Details	42-2
Program Resource Utilization	42-2
Log File Alerts	42-2
Metric Collection Errors	42-2
Storage Details	42-3
Storage Utilization	42-3
Overall Utilization	42-3
Provisioning Summary	42-4
Consumption Summary	42-4
ASM	42-4
Databases	42-5
Disks	42-5
File Systems	42-5
Volumes	42-6
Vendor Distribution	42-7
Storage History	42-7
Storage Layers	42-8
Storage Refresh	42-8
Ksplice for Oracle Linux	42-9

43 Administering Hosts

Configuration Operations on Hosts	43-1
Configuring File and Directory Monitoring Criteria	43-1
Configuring Generic Log File Monitor Criteria	43-2
Configuring Program Resource Utilization Monitoring Criteria	43-3

Administration Tasks	43-4
Services	43-5
Default System Run Level	43-6
Network Card	43-7
Host Lookup Table	43-7
NFS Client	43-8
User and Group Administration (Users)	43-9
User and Group Administration (Groups)	43-10
Using Tools and Commands	43-11
Enabling Sudo and Power Broker	43-11
Executing the Host Command Using Sudo or PowerBroker	43-12
Using Remote File Editor	43-12
Adding Host Targets	43-13
Running Host Command	43-13
Accessing Host Command	43-14
Executing Host Command Using Sudo or Power Broker	43-14
Execute Host Command - Multiple Hosts	43-14
Target Properties	43-15
Execute Host Command - Group	43-16
Execute Host Command - Single Host	43-17
Load OS Script	43-17
Load From Job Library	43-18
Execution History	43-18
Execution Results	43-18
Miscellaneous Tasks	43-18
Enabling Collection of WBEM Fetchlet Based Metrics	43-19
Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts	43-19
Adding and Editing Host Configuration	43-20

Part X Patch Management

44 Patching Software Deployments

Overview of the New Patch Management Solution	44-1
Overview of the Current Patch Management Challenges	44-1
About the New Patch Management Solution	44-2
Overview of Patch Plans	44-4
About Patch Plans	44-4
About Types of Patch Plans	44-6
About the Create Plan Wizard	44-6
Overview of Patch Templates	44-8

About Patch Templates	44-8
About the Edit Template Wizard	44-8
Supported Targets, Releases, and Deployment Procedures for Patching	44-9
Overview of Supported Patching Modes	44-13
Overview of Patching in Online and Offline Mode	44-13
Overview of Patching in In-Place and Out-of-Place Mode	44-14
Overview of Patching in Rolling and Parallel Mode	44-17
Understanding the Patching Workflow	44-19
Setting Up the Infrastructure for Patching	44-19
Meeting Basic Infrastructure Requirements for Patching	44-20
Creating Administrators with the Required Roles for Patching	44-20
Setting Up the Infrastructure for Patching in Online Mode (Connected to MOS)	44-22
Enabling Online Mode for Patching	44-22
Registering the Proxy Details for My Oracle Support	44-22
Setting Up the Infrastructure for Patching in Offline Mode (Not Connected to MOS)	44-23
Enabling Offline Mode for Patching	44-24
Downloading Enterprise Manager Catalog Zip File From Another Host With Internet Connectivity	44-24
Uploading Enterprise Manager Catalog Zip File from your Host With No Internet Connectivity	44-24
Uploading Patches to Oracle Software Library	44-25
Analyzing the Environment and Identifying Whether Your Targets Can Be Patched	44-28
Identifying the Patches to Be Applied	44-29
About Patch Recommendations	44-30
About Knowledge Articles for Patching	44-33
About Service Requests for Patching	44-33
Searching for Patches on My Oracle Support	44-33
Searching for Patches in Oracle Software Library	44-34
Applying Patches	44-35
Creating a Patch Plan	44-36
Accessing the Patch Plan	44-38
Analyzing, Preparing, and Deploying Patch Plans	44-39
Switching Back to the Original Oracle Home After Deploying a Patch Plan	44-49
Saving Successfully Analyzed or Deployed Patch Plan As a Patch Template	44-50
Creating a Patch Plan from a Patch Template and Applying Patches	44-51
Patching Oracle Grid Infrastructure Targets	44-52
Patching Oracle Exadata	44-53
Patching Oracle Data Guard Targets	44-57
Oracle Data Guard Patching Workflow	44-57
Oracle Data Guard Patching Scenarios	44-58

Patching Oracle Identity Management Targets	44-62
Patching Oracle Siebel Targets	44-62
Patching Oracle Service Bus	44-63
Rollback of Oracle Service Bus Patch	44-65
Deploying WebLogic Patches Along with SOA or Oracle Service Bus Patches In A Single Patch Plan	44-66
Diagnosing and Resolving Patching Issues	44-68
Workaround for Errors	44-68
Workaround for Missing Property Errors	44-69
Workaround for Unsupported Configuration Errors	44-70
Common Patching Issues	44-70
Resolving Patching Issues	44-72
Rolling Back Patches	44-73
Additional Patching Tasks You Can Perform	44-73
Viewing or Modifying a Patch Template	44-74
Saving a Deployed Patch Plan as a Patch Template	44-74
Downloading Patches from a Patch Template	44-75
Deleting a Patch Plan	44-75
Deleting a Patch Template	44-76
Converting a Nondeployable Patch Plan to a Deployable Patch Plan	44-76
Associating Additional Targets to a Patch in a Patch Plan	44-76
Manually Staging the Patching Root Component	44-78
Restricting Root User Access for Patching	44-78
Resolving Patch Conflicts	44-78
Analyzing the Results of Patching Operations	44-79
Customizing Patching Deployment Procedures	44-79
Customizing a Static Patching Deployment Procedure	44-79
Customizing a Dynamic Patching Deployment Procedure	44-80
Pausing the Patching Process While Patching Targets in Rolling Mode	44-80
Rolling Back Patches	44-81
End-to-End Use Case: Patching Your Data Center	44-82
Patching Database as a Service Pools	44-82

45 Patching Linux Hosts

Overview of Patching Linux Hosts	45-1
About the Deployment Procedure for Patching Linux Hosts	45-2
Supported Linux Releases	45-2
Setting Up Infrastructure for Linux Patching	45-2
Prerequisites for Using the Linux Patching Feature	45-3
Setting Up the RPM Repository for Linux Patching	45-3
Prerequisites for Setting Up the RPM Repository	45-3

Setting Up the RPM Repository for Patching	45-5
Setting Up Linux Patching Group for Compliance Reporting	45-7
Prerequisites for Setting Up Linux Patching Group	45-7
Setting Up a Linux Patching Group	45-8
Patching Linux Hosts	45-9
Applying Patches on a Linux Patching Group Based on Compliance	45-10
Applying Ad Hoc or Emergency Patches on Linux Hosts	45-12
Managing Linux Configuration Files	45-14
Overview of Linux Configuration Files	45-14
Prerequisites for Managing Configuration Files	45-15
Creating a Linux Configuration File Channel	45-15
Uploading Linux Configuration Files to a Particular Channel	45-15
Prerequisites for Uploading Linux Configuration Files	45-15
Uploading Linux Configuration Files	45-15
Importing Linux Configuration Files from One Channel to Another	45-16
Prerequisites for Importing Linux Configuration Files	45-16
Importing Linux Configuration Files	45-16
Deploying Linux Configuration Files From a Particular Channel	45-16
Prerequisites for Deploying Linux Configuration Files	45-16
Deploying Linux Configuration Files	45-17
Deleting a Linux Configuration File Channel	45-17
Prerequisites for Deleting a Linux Configuration File Channel	45-17
Deleting Linux Configuration File Channels	45-17
Oracle Grid Infrastructure and Oracle RAC Configuration Support	45-17
Additional Linux Patching Tasks You Can Perform	45-19
Viewing Linux Patching Compliance History	45-20
Prerequisites for Viewing Linux Patching Compliance History	45-20
Viewing Linux Patching Compliance History	45-20
Patching Non-Compliant Linux Packages	45-20
Prerequisites for Patching Non-Compliant Linux Packages	45-21
Patching Non-Compliant Linux Packages	45-21
Rolling Back Linux Patch Update Sessions or Deinstalling Packages	45-21
Prerequisites for Rolling Back Linux Patch Update Sessions or Deinstalling Packages	45-21
Rolling Back Linux Patch Update Sessions or Deinstalling Packages	45-22
Registering a Custom Package Channel	45-22
Prerequisites for Registering a Custom Package Channel	45-22
Registering a Custom Package Channel	45-22
Cloning a Package Channel	45-23
Prerequisites for Cloning a Package Channel	45-23
Cloning a Package Channel	45-23

Copying Packages from One Channel to Another	45-24
Prerequisites for Copying Packages from One Channel to Another	45-24
Copying Packages from One Channel to Another	45-24
Adding Custom Packages to a Channel	45-25
Prerequisites for Adding Custom Packages to a Channel	45-25
Adding Custom Packages to a Channel	45-25
Deleting a Package Channel	45-25
Prerequisites for Deleting a Package Channel	45-26
Deleting a Package Channel	45-26

46 Oracle Linux Home

Oracle Linux Home Target	46-1
Oracle Linux Home page	46-1
General Region	46-2
Overview of Incidents and Problems	46-2
Host Flux	46-3
CPU	46-3
Memory	46-4
Oracle Linux Patching	46-5
Patchable Linux Groups	46-5
Patchable Linux Hosts	46-6
Ksplice for Oracle Linux	46-8
Ksplice Metrics	46-8
Ksplice Patching	46-9
Ksplice Linux Hosts Page	46-10
Oracle Linux Home	46-12
Target Navigation Tree (TNT) of "Oracle Linux Home"	46-12

47 Database Fleet Maintenance

About Database Fleet Maintenance	47-1
Getting Started with Fleet Maintenance	47-2
Meeting the Prerequisites	47-2
Discovering Configuration Pollution	47-4
Creating a Gold Image	47-5
Retrieving a List of Available Gold Images	47-8
Verifying if Image is Applicable	47-9
Deleting an Image or Versions of an Image	47-10
Subscribing the Targets to the Selected Image	47-11
Verifying the Subscription	47-12

Unsubscribing Targets from an Image	47-13
Creating an Updated Version of the Image	47-14
Changing Version Status to Current	47-14
Staging and Deploying the Software	47-15
Migrating the Listeners	47-17
Deploying Container Databases	47-18
Updating the Database / Cluster	47-18
Rolling Back Cluster / Database	47-20
Non-Rolling Applicable Patches	47-21
Updating and Rolling Back Standby Databases	47-21
Node Wise RAC Database / Cluster Update	47-23
Update De-mystified	47-25
Upgrading Oracle Homes	47-26
Cleaning Up Oracle Homes	47-27
Custom Pre and Post Scripts for Fleet Operations	47-28
Custom Script Considerations	47-28
Software Library Entity For Uploading The Script	47-29
How to Get URN for SWLIB Entity	47-29
Pre/Post Script Usage	47-30
Database Fleet Maintenance - RESTful APIs	47-30

48 Performing Engineered System Software Updates

Overview of Exadata System Software Update	48-1
Configuring Options for Exadata Component Software Updates	48-3
Updating Exadata Database Servers	48-5
Updating Exadata Storage Servers	48-5
Updating Exadata Infiniband Switches	48-7
Rolling Backup Deployed Software Updates	48-8
Patching Oracle Identity Management Targets	48-8
Overview of Exalytics System Software Update	48-9
Configuring the Options for Oracle Exalytics Updates	48-9
Updating Oracle Exalytics Compute Nodes	48-10
Updating Oracle Exalytics Business Intelligence Instance	48-11

Part XI Configuration, Compliance, and Change Management

49 Managing Configuration Information

Overview of Configuration Management	49-1
Overview of Configuration Searches	49-3

Managing Configuration Searches	49-3
Searching for a Configuration Search	49-3
Running a Configuration Search	49-4
Editing a Configuration Search	49-4
Deleting a Configuration Search	49-5
Importing or Exporting a Configuration Search	49-5
Creating a Configuration Search	49-5
Creating a New Configuration Search	49-5
Creating a Configuration Search from an Existing Configuration Search	49-7
Creating a Configuration Search Using SQL	49-7
Overview of Configuration Browser	49-7
Viewing Configuration Data	49-8
Working with Saved Configurations	49-10
Working with Inventory and Usage Details	49-11
Overview of Configuration History	49-12
Accessing Configuration History	49-13
Working with Configuration History	49-13
Searching History	49-13
Annotating Configuration Changes	49-15
Scheduling a History Search and Creating a Notification List	49-15
Saving History to a File	49-16
Saving Configuration History	49-16
Creating a Search Using SQL	49-16
Viewing History Job Activity	49-16
Overview of Comparisons and Templates	49-17
About Comparison Templates	49-17
Working with Comparison Templates	49-18
Creating or Editing a Comparison Template	49-18
Managing Comparison Templates	49-20
Specifying Rules	49-21
Creating a Value Constraint Rule	49-22
Creating a Matching Rule	49-22
Creating a Rule for Including and Excluding Configuration Items	49-22
About Rules Expression and Syntax	49-23
Understanding Rules by Example	49-25
Matching Rule Examples	49-25
Ignore Rule Examples	49-26
About Comparisons	49-28
Considerations Before Creating a Comparison	49-28
Steps in Setting Up a Drift or Consistency Comparison	49-29
About One-Time Comparisons	49-29

About Configuration Drift	49-31
About Configuration Consistency	49-32
About the Definition Library	49-34
Setting Up a Comparison Template	49-34
Creating Notifications for Comparisons	49-35
Working with Comparison Results	49-36
About Consistency Management (System) Comparison Results	49-36
About Drift (Target) Comparison Results	49-37
Synchronizing Configuration Extension Files	49-37
Comparison and Drift Management BI Publisher Reports	49-39
Overview of Configuration Extensions and Collections	49-40
Working with Configuration Extensions	49-41
Creating a Custom Target Type	49-41
Creating or Editing a Configuration Extension	49-41
Using the Files & Commands Tab	49-42
Using the SQL Tab	49-43
Setting Up Credentials When Creating a Configuration Extension	49-44
Setting Up Rules	49-44
Managing Configuration Extensions	49-45
About Configuration Extensions and Versioning	49-47
About Configuration Extensions and Privileges	49-47
About Configuration Extensions and Deployment	49-48
Deploying and Undeploying Configuration Extensions	49-48
Editing a Deployment of Configuration Extensions	49-49
Viewing a Configuration Collection	49-50
Extending Configuration Data Collections	49-50
Extending Existing Target Collections	49-50
Adding New Target Data Collections	49-51
Using Configuration Extensions as Blueprints	49-52
Overview of Parsers	49-53
Managing Parsers	49-53
About XML Parsers	49-54
About the Default XML Parser	49-54
About the Generic XML Parser	49-55
XML Parser Examples	49-56
About Format-Specific Parsers	49-57
Database Query Parser Parameters	49-59
Database Query Paired Column Parser Parameters	49-59
Directory Parser Parameters	49-60
E-Business Suite Parser Parameters	49-60
Galaxy CFG Parser Parameters	49-60

MQ-Series Parser Parameters	49-61
Siebel Parser Parameters	49-61
Unix Installed Patches Parser Parameters	49-61
Unix Recursive Directory List Parser Parameters	49-61
About Columnar Parsers	49-62
Columnar Parser Parameters	49-63
About Properties Parsers	49-64
Basic Properties Parser Parameters	49-65
Advanced Properties Parser Parameters	49-66
Advanced Properties Parser Constructs	49-68
Using Parsed Files and Rules	49-72
Sample XML File Parsing and Rule Application	49-72
Sample Non-XML File Parsing and Rule Application	49-73
Sample SQL Query Parsing and Rule Application	49-75
Overview of Relationships	49-77
Overview of Configuration Topology Viewer	49-78
About Configuration Topology Viewer	49-79
Examples of Using Topology	49-79
Viewing a Configuration Topology	49-80
Determining System Component Structure	49-81
Determining General Status of Target's Configuration Health	49-81
Getting Configuration Health/Compliance Score of a Target	49-81
Analyzing a Problem and Viewing a Specific Issue in Detail	49-82
About Dependency Analysis	49-82
About Impact Analysis	49-83
Creating a Custom Topology View	49-83
Deleting a Custom Topology View	49-84
Excluding Relationships from a Custom Topology View	49-84
Including Relationships to a Target in a Custom Topology View	49-84
Creating a Relationship to a Target	49-85
Deleting a Relationship from a Target	49-85
Controlling the Appearance of Information on a Configuration Topology Graph	49-86

50 Managing Compliance

Overview of Compliance	50-1
Terminology Used in Compliance	50-2
Accessing the Compliance Features	50-4
Roles and Privileges Needed to Use the Compliance Features	50-4
Evaluating Compliance	50-6
Accessing Compliance Statistics	50-8

Using the Compliance Dashboard Effectively	50-8
Viewing Compliance Summary Information	50-9
Viewing Target Compliance Evaluation Results	50-9
Viewing Compliance Framework Evaluation Results	50-10
Managing Violations	50-11
Investigating Compliance Violations and Evaluation Results	50-12
Investigating Violations of Repository Compliance Standard Rules and Targets Causing Violations	50-13
Viewing All the Violations Reported for Your Enterprise	50-13
Examples of Viewing Violations	50-14
Investigating Evaluation Errors	50-17
Analyzing Compliance Reports	50-17
Overview of Compliance Score and Importance	50-18
Compliance Score of a Compliance Standard Rule -Target	50-19
Real-time Monitoring Rule Compliance Score	50-19
Compliance Score of a Compliance Standard for a Target	50-20
Compliance Framework Compliance Score	50-20
Parent Node Compliance Score	50-21
Investigating Real-time Observations	50-22
Viewing Observations	50-22
Viewing Observations By Systems	50-22
Viewing Observations By Compliance Framework	50-23
Viewing Observations By Search	50-24
Viewing Details of an Incident	50-25
Operations on Observations During Compliance Evaluation	50-25
Manually Setting an Observation As Authorized Or Not Authorized	50-25
Notifying a User When an Observation Occurs	50-26
Notifying a User When an Authorized Observation Occurs	50-27
Configuring Compliance Management	50-27
About Compliance Frameworks	50-27
Operations on Compliance Frameworks	50-29
Creating a Compliance Framework	50-29
Creating Like a Compliance Framework	50-31
Editing a Compliance Framework	50-31
Deleting a Compliance Framework	50-32
Exporting a Compliance Framework	50-32
Importing a Compliance Framework	50-33
Browsing Compliance Frameworks	50-33
Searching Compliance Frameworks	50-34
Browsing Compliance Framework Evaluation Results	50-34
Searching Compliance Framework Evaluation Results	50-34

Browsing Compliance Framework Errors	50-34
Searching Compliance Framework Errors	50-35
Verifying Database Targets Are Compliant with Compliance Frameworks	50-35
About Compliance Standards	50-36
Operations on Compliance Standards	50-38
Creating a Compliance Standard	50-38
Creating Like a Compliance Standard	50-41
Editing a Compliance Standard	50-41
Deleting a Compliance Standard	50-41
Exporting a Compliance Standard	50-42
Importing a Compliance Standard	50-42
Browsing Compliance Standards	50-43
Searching Compliance Standards	50-43
Browsing Compliance Standard Evaluation Results	50-43
Searching Compliance Standard Evaluation Results	50-44
Browsing Compliance Standard Errors	50-44
Searching Compliance Standard Errors	50-44
Associating a Compliance Standard with Targets	50-45
Associating a Compliance Standard with a Group Target	50-46
Viewing Real-time Monitoring Compliance Standard Warnings	50-46
Enabling Security Metrics	50-47
Considerations When Creating Compliance Standards	50-47
About Compliance Standard Rule Folders	50-48
Creating Rule Folders	50-48
Managing Rule Folders in a Compliance Standard	50-48
About Compliance Standard Rules	50-49
Operations on Compliance Standards Rules	50-51
Creating a Repository Compliance Standard Rule	50-52
Creating a WebLogic Server Signature Compliance Standard Rule	50-54
Creating a Real-time Monitoring Compliance Standard Rule	50-58
Creating an Agent-side Rule	50-65
Creating a Manual Rule	50-66
Creating a Missing Patches Compliance Standard Rule	50-67
Creating a Configuration Consistency Rule	50-70
Creating Configuration Drift Rule	50-71
Creating Like a Compliance Standard Rule	50-72
Editing a Compliance Standard Rule	50-73
Deleting a Compliance Standard Rule	50-73
Exporting a Compliance Standard Rule	50-74
Importing a Compliance Standard Rule	50-74
Browsing Compliance Standard Rules	50-74

Searching Compliance Standard Rules	50-75
Using Corrective Actions	50-75
Real-time Monitoring Facets	50-76
About Real-time Monitoring Facets	50-77
Facet Entity Types	50-77
Facet Patterns	50-78
Operations on Facets	50-79
Viewing the Facet Library	50-79
Creating and Editing Facets	50-80
Creating and Editing Facet Folders	50-82
Deleting a Facet	50-82
Using Create Like to Create a New Facet	50-83
Importing and Exporting Facets	50-84
Changing Base Facet Attributes Not Yet Used In a Rule	50-84
Examples	50-85
Creating Repository Rule Based on Custom Configuration Collections	50-85
Creating Compliance Standard Agent-side and Manual Rules	50-88
Suppressing Violations	50-95
Clearing Violations	50-97

51 Managing Enterprise Data Governance

Overview of Enterprise Data Governance	51-1
About Enterprise Data Governance	51-1
What Are Protection Policies?	51-2
What Are Application Signatures?	51-2
Using Enterprise Data Governance	51-2
The Enterprise Data Governance Dashboard	51-3
Working with Sensitive Database Discovery Results	51-3
Working with Metadata Discovery Jobs	51-3
Creating a Metadata Discovery Job	51-4
Managing Automatic Metadata Discovery	51-4
Managing Metadata Discovery Results	51-5
Working with Data Discovery Jobs	51-5
Creating a Data Discovery Job	51-5
Managing Data Discovery Results	51-6
Creating Custom Application Signatures	51-7

52 Managing Database Schema Changes

Overview of Change Management for Databases	52-1

Using Schema Baselines	52-2
Overview of Scope Specification	52-3
About Capturing a Schema Baseline Version	52-4
About Working With A Schema Baseline Version	52-4
About Working With Multiple Schema Baseline Versions	52-4
Exporting and Importing Schema Baselines	52-5
Creating Directory Objects for Export and Import	52-5
Using Schema Comparisons	52-6
Defining Schema Comparisons	52-6
About Working with Schema Comparison Versions	52-8
Using Schema Synchronizations	52-9
About Defining Schema Synchronizations	52-9
Creating a Synchronization Definition from a Comparison	52-11
Working with Schema Synchronization Versions	52-12
About the Schema Synchronization Cycle	52-12
Creating Additional Synchronization Versions	52-15
Using Change Plans	52-15
About Working with Change Plans	52-16
Creating a Change Plan	52-16
Creating and Applying a Change Plan From a Schema Comparison	52-16
Using External Clients to Create and Access Change Plans in Cloud Control	52-19
Submitting Schema Change Plans From SQL Developer Interface	52-20
Using Database Data Comparison	52-21
Requirements for Database Data Comparisons	52-21
Comparing Database Data and Viewing Results	52-23

53 Additional Setup for Real-time Monitoring

Overview of Real-Time Monitoring	53-1
Overview of Resource Consumption Considerations	53-2
OS File Monitoring Archiving	53-2
OS File Read Monitoring	53-2
Creating Facets That Have Very Broad Coverage	53-2
Cloud Control Repository Sizing	53-3
Configuring Monitoring Credentials	53-3
Preparing To Monitor Linux Hosts	53-4
OS File Monitoring	53-4
Debugging Kernel Module Or Other File Monitoring Issues	53-6
Preparing To Monitor Windows Hosts	53-7
Verifying Auditing Is Configured Properly	53-9
Subinac External Requirements	53-9

Preparing To Monitor Solaris Hosts	53-9
Enabling BSM Auditing	53-9
Enabling BSM Auditing Using Solaris Versions 9 and 10	53-10
Enabling BSM Auditing Using Solaris 11	53-10
Managing Audit Log Files	53-11
Preparing to Monitor AIX Hosts	53-11
Installation Prerequisite for AIX 5.3	53-11
Administering AIX Auditing	53-12
Verifying AIX System Log Files for the OS User Monitoring Module	53-13
Preparing To Monitor the Oracle Database	53-13
Setting Auditing User Privileges	53-13
Specifying Audit Options	53-14
Setting Up Change Request Management Integration	53-15
BMC Remedy Action Request System 7.1 Integration	53-15
Installing and Customizing Remedy ARS	53-15
Overview of the Repository Views Related to Real-time Monitoring Features	53-21
Modifying Data Retention Periods	53-26
Real-time Monitoring Supported Platforms	53-27
OS User Monitoring	53-28
OS Process Monitoring	53-29
OS File Monitoring	53-30
OS Windows Registry Monitoring	53-33
OS Windows Active Directory User Monitoring	53-33
OS Windows Active Directory Computer Monitoring	53-34
OS Windows Active Directory Group Monitoring	53-34
Oracle Database Table Monitoring	53-35
Oracle Database View Monitoring	53-35
Oracle Database Materialized View Monitoring	53-36
Oracle Database Index Monitoring	53-36
Oracle Database Sequence Monitoring	53-37
Oracle Database Procedure Monitoring	53-37
Oracle Database Function Monitoring	53-38
Oracle Database Package Monitoring	53-38
Oracle Database Library Monitoring	53-39
Oracle Database Trigger Monitoring	53-39
Oracle Database Tablespace Monitoring	53-39
Oracle Database Cluster Monitoring	53-40
Oracle Database Link Monitoring	53-40
Oracle Database Dimension Monitoring	53-40
Oracle Database Profile Monitoring	53-40
Oracle Database Public Link Monitoring	53-41

Oracle Database Public Synonym Monitoring	53-41
Oracle Database Synonym Monitoring	53-41
Oracle Database Type Monitoring	53-42
Oracle Database Role Monitoring	53-42
Oracle Database User Monitoring	53-43
Oracle Database SQL Query Statement Monitoring	53-43

54 Overview of Change Activity Planner

Before Getting Started	54-1
Change Activity Planner Roles and Privileges	54-1
Change Activity Planner Terminology	54-2
Plan	54-2
Task Definition	54-3
Task Group	54-5
Task	54-5
Creating a Change Activity Plan	54-6
Creating a Task Definition	54-8
Creating a Task Group	54-11
Operations on Change Activity Plans	54-11
Creating a Plan Like Another Plan	54-12
Editing a Plan	54-12
Deleting a Plan	54-13
Deactivating a Plan	54-13
Exporting Plans	54-14
Printing Plans	54-14
Changing the Owner of a Plan	54-14
Managing a Change Activity Plan	54-15
Summary Tab	54-16
Tasks Tab	54-17
Comments and Audit Trail Tab	54-18
Viewing My Tasks	54-18
Example of Using Change Activity Planner	54-20
Automating Activity Planning	54-20
Additional Steps in Automating Activity Planning	54-21
Using Change Activity Planner for Patching	54-22

Part XII Deployment Procedures

55 About Deployment Procedures

Overview of the Provisioning Page	55-1
Granting Roles and Privileges to Administrators	55-3
Granting Roles and Privileges to Administrators on the Deployment Procedure	55-4
Granting Roles and Privileges to Administrators on Software Library	55-6
Components of a Procedure	55-6
Target List	55-6
Procedure Variables	55-7
Phases and Steps	55-8
Types of Phases	55-8
Types of Procedure Steps	55-8
Performing Tasks on Procedure Steps	55-10
Creating a Procedure	55-11
Adding Rolling or Parallel Phase	55-11
Adding Steps	55-13
Managing Deployment Procedures	55-17
Viewing, Editing, Deleting a Procedure	55-18
Editing and Saving Permissions of a Procedure	55-18
Tracking the Procedure Execution and Status of Deployment Procedures	55-18
Rescheduling a Procedure	55-20
Reverting a Procedure	55-20
Setting Step Level Grace Period	55-20
Creating, Saving, and Launching User Defined Deployment Procedure (UDDP)	55-21
Step 1: Creating User Defined Deployment Procedure	55-21
Step 2: Saving and Launching User Defined Deployment Procedure with Default Inputs	55-22
Saving and Launching the Deployment Procedure with Lock Down	55-23
Step 3: Launching and Running the Saved User Defined Deployment Procedure	55-30
Step 4: Tracking the Submitted User Defined Deployment Procedure	55-31
Procedure Instance Execution Page	55-31
Comparison Between the Existing Design and the New Design for Procedure Instance Execution Page	55-32
Overview of the Procedure Instance Execution Page	55-33
Investigating a Failed Step for a Single or a Set of Targets	55-35
Retrying a Failed Step	55-36
Creating an Incident	55-36
Viewing the Execution Time of a Deployment Procedure	55-36
Searching for a Step	55-36
Downloading a Step Output	55-36

56 Customizing Deployment Procedures

About Deployment Procedure Customization Types	56-1
Customizing a Deployment Procedure	56-3
Editing the Rolling and Parallel Phase of a Deployment Procedure	56-3
Editing a Job Step of a Deployment Procedure	56-4
Editing a Directive Step of a Deployment Procedure	56-4
Editing a Component Step of a Deployment Procedure	56-4
Editing a File Transfer Step of a Deployment Procedure	56-6
Editing a Host Command Step of a Deployment Procedure	56-6
Editing a Manual Step of a Deployment Procedure	56-7
A Workflow Example for Assigning Values to Deployment Procedure Variables at Runtime	56-8
Step 1: Creating a Perl Script to Assign Values to Deployment Procedure Variables at Runtime	56-10
Step 2: Uploading TestPingAndDPvariable.pl to Software Library	56-10
Step 3: Creating a Deployment Procedure	56-11
Step 4: Launching the Deployment Procedure, and Providing the Variable Values at Runtime	56-12
Step 5: Verifying the Deployment Procedure Variable Values	56-12
Changing Deployment Procedure Error Handling Modes	56-14
Setting Up email Notifications Regarding the Status of Deployment Procedures	56-15
Configuring an Outgoing Mail (SMTP) Server In Enterprise Manager	56-16
Adding E-mail Addresses for Enterprise Manager Notifications	56-17
Copying Customized Provisioning Entities from One Enterprise Manager Site to Another	56-17
Prerequisites for Copying Customized Provisioning Entities from One Enterprise Manager Site to Another	56-18
Copying Customized Provisioning Entities from One Enterprise Manager Site to Another	56-18
A Workflow Example for Customizing a Directive	56-19
Creating and Uploading a Copy of a Default Directive	56-19
Customizing a Deployment Procedure to Use the New Directive	56-20
Running the Customized Deployment Procedure	56-20

Part XIII Additional Information

A Using Enterprise Manager Command Line Interface

Overview	A-1

Prerequisites	A-2
Enterprise Manager Command Line Interface Verbs	A-2
Provisioning EM CLI Verbs	A-2
Patching EM CLI Verbs	A-5
Software Library EM CLI Verbs	A-9
Provisioning Using EM CLI	A-11
Creating the Properties File to Submit a Deployment Procedure	A-12
Step 2: Entering New Values in The Properties File	A-14
Step 3: Submitting the Procedure With The Updated Properties File as Input	A-14
Step 4: Verifying The Status Of the Procedure	A-15
Using Properties File from an Existing Execution of a Deployment Procedure	A-16
Launching a Procedure using an Existing Saved Procedure	A-18
Saving a Procedure Configuration of a Procedure	A-18
Updating the Procedure Configuration of a Procedure	A-19
Provisioning Pluggable Databases	A-19
Creating a New Pluggable Database	A-19
Provisioning a Pluggable Database Using a Snapshot Profile	A-21
Migrating a Non-Container Database as a Pluggable Database	A-23
Unplugging and Dropping a Pluggable Database	A-24
Patching Using EM CLI	A-24
Before You Begin Patching	A-25
Patching Using EM CLI	A-25
Creating a New Properties File for Patching Targets	A-26
Using the Properties File of an Existing Patch Plan to Patch the targets	A-30
WorkFlow Examples Using EM CLI Commands	A-31
Provisioning Oracle Database Software	A-32
Provisioning Oracle WebLogic Server	A-33
Prerequisites for Provisioning Oracle WebLogic Server	A-33
Provisioning Oracle WebLogic Server Using the Provisioning Profile	A-33
Scaling Up or Scaling Out Middleware Deployment Procedure	A-35
Provisioning User Defined Deployment Procedure	A-41
Prerequisites for Provisioning User Defined Deployment Procedure	A-41
Adding Steps and Phases to User Defined Deployment Procedure Using GUI	A-41
Using EM CLI commands to Run an Instance of the Procedure	A-42
Patching WebLogic Server Target	A-42
Creating a New Generic Component by Associating a Zip File	A-47
Step 1: Identifying the Parent Folder in Software Library	A-47
Step 2: Creating a Generic Component Entity	A-48
Step 3: Associating a Zip File to the Generic Component	A-50
Step 4: Verifying the Newly Created Entity	A-51

Migrate and Remove a Software Library Storage Location	A-51
Step 1: Adding a Destination Storage Location for Migrating Files	A-51
Step 2: Migrate and Remove an existing storage location	A-52
Adding ATS Service Test from Using EM CLI	A-52
Limitations of Using Enterprise Manager Command Line Interface	A-54

B Checking Host Readiness Before Provisioning or Patching

Setting Up User Accounts Before Provisioning	B-1
Configuring SSH	B-2
Shell Limits	B-2
Root Setup (Privilege Delegation)	B-2
Environment Settings	B-2
Kernel Requirements	B-3
Node Time Requirements	B-4
Package Requirements	B-4
Memory and Disk Space Requirements	B-4
Network & IP Address Requirements	B-5
Storage Requirements	B-6
Installation Directories and Oracle Inventory	B-7

C Using emctl partool Utility

Overview of Provisioning Archive Files	C-1
Overview of emctl partool Utility	C-1
Checking Oracle Software Library	C-3
Exporting Deployment Procedures	C-3
Obtaining Deployment Procedure's GUID	C-3
Creating PAR File	C-4
Importing PAR Files	C-5
Importing Using Command Line Interface	C-5
Importing Specific PAR File	C-5
Importing All PAR Files	C-6
Importing Using Cloud Control Console	C-6

D Understanding PXE Booting and Kickstart Technology

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

E End-to-End Use Case: Patching Your Data Center

The Challenge of Patching Your Data Center	E-1
The Enterprise Manager Solution	E-1
Identify the Patches Relevant to Your Data Center	E-2
Prepare, Test, and Certify the Patch Rollout Plan	E-2
Create a Change Activity Plan to Roll Out the Patches	E-2
Monitor the Progress and Report the Status of the Change Activities	E-2
Executing the Example Scenario	E-2
Create Administrators with the Required Roles	E-3
Set Up the Infrastructure	E-3
Analyze the Environment and Identify Whether Your Targets Can Be Patched	E-3
Identify the Relevant Patches	E-4
Create a Patch Plan, Test the Patches, and Certify the Patches	E-4
Create a Change Activity Plan to Roll Out the Patches	E-4
Roll Out the Patches	E-4
Check and Report the Status of the Change Activities	E-5
Verify If the Targets Have Been Patched	E-5

F Troubleshooting Issues

Troubleshooting Database Provisioning Issues	F-1
Grid Infrastructure Root Script Failure	F-1
Issue	F-1
Description	F-1
Solution	F-1
SUDO Error During Deployment Procedure Execution	F-2
Issue	F-2
Description	F-2
Solution	F-2
Prerequisites Checks Failure	F-2
Issue	F-2
Cause	F-2
Solution	F-2
Oracle Automatic Storage Management (Oracle ASM) Disk Creation Failure	F-3
Issue	F-3
Cause	F-3
Solution	F-3
Oracle ASM Disk Permissions Error	F-3
Issue	F-3
Description	F-3
Solution	F-3

Specifying a Custom Temporary Directory for Database Provisioning	F-3
Incident Creation When Deployment Procedure Fails	F-4
Issue	F-4
Solution	F-4
Reading Remote Log Files	F-4
Retrying Failed Jobs	F-4
Issue	F-4
Solution	F-4
Troubleshooting Patching Issues	F-5
Oracle Software Library Configuration Issues	F-5
Error Occurs While Staging a File	F-5
Error Occurs While Uploading a Patch Set	F-6
OPatch Update Job Fails When Duplicate Directories Are Found in the Software Library	F-6
My Oracle Support Connectivity Issues	F-7
Error Occurs While Testing the Proxy Server That Supports Only Digest Authentication	F-7
Host and Oracle Home Credential Issues	F-8
Cannot Create Log Files When You Set Privileged Credentials as Normal Oracle Home Credentials	F-8
Collection Issues	F-9
Missing Details in Plan Wizard	F-10
Cannot Add Targets to a Patch Plan	F-11
Patch Recommendation Issues	F-12
Patch Recommendations Do Not Appear After Installing Oracle Management Agent on Oracle Exadata Targets	F-12
Patch Plan Issues	F-13
Patch Plan Becomes Nondeployable and Fails	F-13
Instances Not to Be Migrated Are Also Shown as Impacted Targets for Migration	F-14
Cluster ASM and Its Instances Do Not Appear as Impacted Targets While Patching a Clusterware Target	F-14
Recovering from a Partially Prepared Plan	F-15
Error #1009 Appears in the Create Plan Wizard While Creating or Editing a Patch Plan	F-16
Analysis Succeeds But the Deploy Button is Disabled	F-16
Patch Plan Fails When Patch Plan Name Exceeds 64 Bytes	F-16
Out-of-Place Patching Fails for 11.2.0.3 Exadata Clusterware	F-17
Patch Plan Analysis Issues	F-18
Patch Plan Remains in Analysis State Even After the Deployment Procedure Ends	F-18
Patch Plan Analysis Fails When the Host's Node Name Property Is Missing	F-19
Link to Show Detailed Progress on the Analysis Is Not Actionable	F-20

Raising Service Requests When You Are Unable to Resolve Analysis Failure Issues	F-20
User Account and Role Issues	F-21
Out-of-Place Patching Errors Out If Patch Designers and Patch Operators Do Not Have the Required Privileges	F-21
Troubleshooting Linux Patching Issues	F-21
Troubleshooting Linux Provisioning Issues	F-22
Frequently Asked Questions on Linux Provisioning	F-25
Refreshing Configurations	F-27
Refreshing Host Configuration	F-27
Refreshing Oracle Home Configuration	F-28
Reviewing Log Files	F-29
OMS-Related Log Files	F-29
Management Agent-Related Log Files	F-30
Advanced Options	F-30
On the OMS Side	F-30
On the Management Agent Side	F-30

Index

Preface

The Lifecycle Management Guide introduces you to the lifecycle management solutions offered by Oracle Enterprise Manager Cloud Control (Cloud Control), and describes in detail how you can use the discovery, provisioning, patching, and configuration and compliance management features to manage your data center.

Audience

This guide is primarily meant for administrators who want to use the discovery, provisioning, patching, and configuration and compliance management features offered by Cloud 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 default as well custom deployment procedures, patch plans, and patch templates to manage the enterprise configuration.

Scope and Coverage

Oracle Enterprise Manager Cloud Control Lifecycle Management Guide describes features pertaining to the following plug-in and platform releases:

- Oracle Database Plug-in (13.3.0.0.0)
- Oracle Fusion Middleware Plug-in (13.3.0.0.0)
- Oracle Cloud Framework Plug-in (13.3.0.0.0)
- Cloud Services Management Plug-in (13.3.0.0.0)
- Enterprise Manager Cloud Control 13c Release 3(13.3.0.0.0)

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=do-cacc>.

Access to Oracle Support

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

Related Documents

For more information, see the following books in the Cloud Control documentation library:

- *Oracle Enterprise Manager Cloud Control Basic Install Guide*
- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*
- *Oracle Enterprise Manager Cloud Control Upgrade Guide*
- *Oracle Enterprise Manager Cloud Control Administrator's Guide*

For the latest releases of these and other Oracle documentation, check the Oracle Technology Network at the following URL:

<http://www.oracle.com/technetwork/indexes/documentation/index.html>

Cloud Control also provides extensive online Help. Click **Help** at the top-right corner of any Cloud Control page to display the online help window.

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.
<i>italic</i>	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 contains the following chapters:

- [Introduction to Lifecycle Management](#)
- [Setting Up Your Infrastructure](#)

1

Introduction to Lifecycle Management

This chapter covers the following:

- [Overview of the New Lifecycle Management Solutions](#)
- [Information Map for Lifecycle Management Solutions](#)

Overview of the New Lifecycle Management Solutions

In today's world, with the cloud infrastructure, numerous low cost servers and software deployments on those servers have brought in a fresh set of lifecycle management challenges. The challenges range from discovering and monitoring the health of existing software deployments to provisioning new software deployments and maintaining them over a period of time.

Besides that, other problems include difficulty in managing consistency and compatibility across these software deployments and operating systems, managing configuration changes, and managing security vulnerabilities that lead to lack of compliance.

These lifecycle management challenges eventually force you to engage more human resources and devote significant amount of time in managing the data center operations.

Oracle Enterprise Manager Cloud Control (Cloud Control) offers lifecycle management solutions that help you meet all lifecycle management challenges easily by automating time-consuming tasks related to discovery, initial provisioning and cloning, patching, configuration management, ongoing change management, and compliance management.

[Figure 1-1](#) illustrates the lifecycle management solutions offered by Cloud Control.

Figure 1-1 Lifecycle Management Solutions

Discovery	Discovery involves identifying hosts and their software deployments, adding them as manageable targets, and monitoring and managing their health.
Provisioning	Provisioning involves repeatable, reliable, automated, unattended and scheduled mass-deployment of software, applications or servers across different platforms, environments, and locations.
Patching	Patching involves maintaining the software over a period of time with bug fixes, enhancements, and latest features offered by the software vendor.
Configuration Management	Configuration management involves discovering, collecting, monitoring, comparing, and customizing configuration data for all managed targets in the enterprise.
Compliance Management	Compliance Management involves defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage.
Change Management	Change Management involves consolidating and propagating changes in database during development, customization, or application upgrades, to ensure database compliance across enterprise.
Enterprise Data Governance	Enterprise Data Governance involves identifying, managing, tracking, and securing sensitive data in the data center.
Change Activity Planner	Change Activity Planner involves planning, managing, and monitoring operations within your data center. These operations involve dependencies and coordination across teams and business owners, as well as multiple processes.

Table 1-1 describes each of these lifecycle management solutions.

Table 1-1 Lifecycle Management Solutions

Solution Area	Coverage
Discovery	- Automatically discovers software deployments using IP scanning techniques (NMAP). - Converts unmanaged software deployments to managed targets in Cloud Control so that their health can be monitored. - Offers an integrated workflow for deploying Oracle Management Agents and discovering targets on selected auto-discovered hosts.
Provisioning	- Discovers bare metal servers and live target servers. - Provisions Linux operating system on bare metal servers (hypervisors and virtual machines). - Associates patching templates with provisioning so that patches can be applied automatically once the operating system is provisioned. - Provisions Oracle Databases, Oracle Real Application Clusters (Oracle RAC), Oracle Grid Infrastructure (for standalone servers and clustered environments), and Pluggable Databases. - Provisions clusters in virtualized Exadata systems. - Supports mass upgrade of single instance, Oracle RAC, and Oracle RAC One database instances one at a time. - Provisions Oracle WebLogic Server, Oracle Coherence, Oracle SOA Suite, SOA Artifacts, Service Bus, Oracle WebCenter, Business Intelligence, and Java EE applications.
Patching	- Offers an integrated patching workflow with My Oracle Support—access to recommendations, search patches, and so on. - Orchestrates patching workflow using <i>Patch Plans</i>, including automated selection of deployment procedures and analysis of the patch conflicts. - Validates patches for applicability in your environment, validates patch plans, and automatically receives patches to resolve conflicts. - Helps you save successfully analyzed or deployable patch plans as patch templates, which contain a predetermined set of patches and deployment options saved from the source patch plan. - Offers out-of-place patching (only for standalone databases), in-place patching, and rolling and parallel patching modes, both in offline and on-line mode.
Change Management	- Captures database object definitions and initialization parameters at different points in time. - Compares databases; compares baselines. - Propagates changes from database definitions and initialization parameters captured in a baseline or from a database to a target database. - Specifies, groups, and packages object metadata changes. Creates change plans from ad hoc changes, comparison-based differences, or developer tools. - Compares data between a local and remote database, and determines how seed data customizations will be affected by application upgrades.

Table 1-1 (Cont.) Lifecycle Management Solutions

Solution Area	Coverage
Configuration Management	• Searches configuration data across the enterprise. • Displays configuration data in the context of a single managed entity—configuration item types and properties, system configuration data, system target relationships, custom configuration data. • Monitors change activity across the enterprise—includes changes both to configurations and to relationships, which are associations that exist among managed entities. • Compares configurations of a particular target type using comparison templates, which enable you to ignore the obvious differences and set alerts on critical issues that need immediate attention. • Identifies files and other configuration data that Cloud Control does not already collect from well-known target types or from a target type introduced as part of the custom configuration definition. • Creates new relationships between managed entities using the Topology Viewer or a generic system target type. Helps you perform dependency analysis and impact analysis on assets in your enterprise using the Topology Viewer.
Compliance Management	• Evaluates the compliance of targets and systems as they relate to your business best practices for configuration, security, and storage. • Advises of how to change configuration to bring your targets and systems into compliance. • Helps you define, customize, and manage Compliance frameworks, Compliance standards, Compliance standard rules. • Helps you test your environment against the criteria defined for your company or regulatory bodies using these self-defined entities.
Enterprise Data Governance	• Provides the means to identify databases within the enterprise that potentially contain sensitive data, and then to evaluate the data within these candidates to determine if sensitive data exists. • Uses metadata discovery to identify databases containing objects that are protected by security features known as Protection Policies. • Discovers sensitive database candidates by identifying application signatures, a set of database objects such as schemas, tables, and views that are unique to a specific application. • Performs metadata discovery automatically whenever a database target is discovered. This feature can be disabled if you want more control over when and how the metadata discovery job runs. • Enables you to associate a sensitive database candidate with a new or existing Application Data Model (ADM) and set sensitive columns for the ADM.
Change Activity Planner	• Enables you to plan, manage, and monitor operations within your data center. These operations involve dependencies and coordination across teams and business owners, as well as multiple processes. • Provides you the ability to create plans comprising of one or more tasks. Tasks can be associated with operations like a patch template, a compliance standard, or a manual job. • Enables you to monitor all managed plans. This helps you to identify any issues that may delay the activity plan completion deadline. • Prints plans that can be used for reporting purposes. Information includes overall summary across all plans, plan summary within a given plan, overall tasks across all tasks across plans, and task summary across tasks within a given plan.

 Note:

The provisioning and patch management solutions are essentially based on deployment procedures, which are Oracle-supplied predesigned procedures that help you accomplish the provisioning and patching tasks. Deployment procedures contain a hierarchical sequence of steps, where each step might contain a sequence of other steps. Essentially, they encapsulate the workflow of all the tasks that need to be performed for a provisioning or patching operation. For more information about deployment procedures, see [About Deployment Procedures](#). For information about the default deployment procedure that you must use for your provisioning or patching operation, refer to the respective chapters.

Information Map for Lifecycle Management Solutions

[Table 1-2](#) lists the chapters and sections relevant to the various lifecycle management solutions offered by Cloud Control. Consider this an information roadmap to learn about the solution and perform the required operations.

Table 1-2 Information Map for Lifecycle Management Solutions

Domain	Target Name	Solution Area	Reference Links
Database	Single-Instance Database	Discovery	Discovering Hosts Automatically and Adding Targets Manually Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Creating and Provisioning Oracle Databases Provisioning Oracle Grid Infrastructure for Oracle Databases Creating Databases
		Upgrade	Upgrading Databases
		Patching	Patching Software Deployments
		Change Management	Managing Database Schema Changes
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Enterprise Data Governance	Managing Enterprise Data Governance
		Change Activity Planner	Overview of Change Activity Planner
	Oracle Real Application Server (Oracle RAC)	Discovery	Discovering Hosts Automatically and Adding Targets Manually Discovering Hosts Manually and Adding Targets Manually

Table 1-2 (Cont.) Information Map for Lifecycle Management Solutions

Domain	Target Name	Solution Area	Reference Links
		Provisioning	- Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases - Provisioning Oracle Real Application Clusters for 12c and 18c - Extending Oracle Real Application Clusters - Deleting or Scaling Down Oracle Real Application Clusters
		Patching	Patching Software Deployments
		Change Management	Managing Database Schema Changes
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Enterprise Data Governance	Managing Enterprise Data Governance
		Change Activity Planner	Overview of Change Activity Planner
	Oracle RAC One Database	Discovery	- Discovering Hosts Automatically and Adding Targets Manually - Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Provisioning Oracle Real Application Clusters One (Oracle RAC One) Node Databases
		Change Management	Managing Database Schema Changes
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Enterprise Data Governance	Managing Enterprise Data Governance
		Change Activity Planner	Overview of Change Activity Planner
	Pluggable Database	Discovery	- Discovering Hosts Automatically and Adding Targets Manually - Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Managing Pluggable Databases Using Enterprise Manager
	Oracle Database Replay Client	Discovery	- Discovering Hosts Automatically and Adding Targets Manually - Discovering Hosts Manually and Adding Targets Manually

Table 1-2 (Cont.) Information Map for Lifecycle Management Solutions

Domain	Target Name	Solution Area	Reference Links
Fusion Middle-ware		Provisioning	Provisioning Oracle Database Re-play Client
		Change Management	Managing Database Schema Changes
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Enterprise Data Governance	Managing Enterprise Data Governance
		Change Activity Planner	Overview of Change Activity Planner
	Oracle Fusion Middleware	Discovery	• Discovering Hosts Automatically and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	• Provisioning Fusion Middleware Domain and Oracle Homes • Scaling Up / Scaling Out Fusion Middleware Domains
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Change Activity Planner	Overview of Change Activity Planner
	Java EE Applications	Discovery	• Discovering Hosts Automatically and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Deploying / Redeploying / Undeploying Java EE Applications
		Configuration Management	Managing Configuration Information
		Compliance Management	Managing Compliance
		Change Activity Planner	Overview of Change Activity Planner
	Coherence Nodes and Clusters	Discovery	• Discovering Hosts Automatically and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Provisioning Coherence Nodes and Clusters
		Configuration Management	Managing Configuration Information

Table 1-2 (Cont.) Information Map for Lifecycle Management Solutions

Domain	Target Name	Solution Area	Reference Links
	Oracle SOA Ar- tifacts and Composites	Compliance Manage- ment	Managing Compliance
		Change Activity Plan- ner	Overview of Change Activity Plan- ner
		Discovery	• Discovering Hosts Automatical- ly and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Provisioning SOA Artifacts and Composites
		Configuration Man- agement	Managing Configuration Information
		Compliance Manage- ment	Managing Compliance
		Change Activity Plan- ner	Overview of Change Activity Plan- ner
	Service Bus	Discovery	• Discovering Hosts Automatical- ly and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Provisioning Service Bus Resources
		Configuration Man- agement	Managing Configuration Information
		Compliance Manage- ment	Managing Compliance
		Change Activity Plan- ner	Overview of Change Activity Plan- ner
	Others	Compliance Manage- ment	Managing Compliance
		Change Activity Plan- ner	Overview of Change Activity Plan- ner
Supported Op- erating Sys- tems for Patching and Provisioning	Oracle Linux, Red Hat Enter- prise Linux (RHEL), SuSE Linux (SLES)	Discovery	• Discovering Hosts Automatical- ly and Adding Targets Manually • Discovering Hosts Manually and Adding Targets Manually
		Provisioning	Provisioning Bare Metal Servers
		Patching	Patching Linux Hosts
		Configuration Man- agement	Managing Configuration Information
		Compliance Manage- ment	Managing Compliance
		Change Activity Plan- ner	Overview of Change Activity Plan- ner

2

Setting Up Your Infrastructure

This chapter describes the infrastructure requirements you must meet before you start using the lifecycle management features. This chapter is essentially for administrators or designers who create the infrastructure. The requirements described in this chapter have to be performed only once.

This chapter covers the following:

- [Getting Started with Setting Up Your Infrastructure](#)
- [Setting Up Oracle Software Library](#)
- [Setting Up Credentials](#)
- [Creating Enterprise Manager User Accounts](#)
- [\(Optional\) Setting Up My Oracle Support](#)
- [\(Optional\) Configuring Self-Update](#)
- [\(Optional\) Setting Up E-mail Notifications](#)
- [\(Optional\) Setting Restricted Accesses for the Root Components](#)

Getting Started with Setting Up Your Infrastructure

This chapter helps you get started by providing an overview of all the steps involved in setting up your infrastructure. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully set up your infrastructure for carrying out all the lifecycle management tasks, including Patching and Provisioning.

[Figure 2-1](#) is a pictorial representation of the sequence of steps you must perform in order to setup your infrastructure.

Figure 2-1 Setting Up Your Infrastructure WorkFlow



Click the reference links provided against the steps in the [Table 2-1](#) for more information on each section.

Table 2-1 Getting Started with Setting Up Your Infrastructure

Step	Description	Reference Links
Step 1	Setting Up Software Library	Setting Up Oracle Software Library
Step 2	Setting Up Credentials	Setting Up Credentials
Step 3	Creating Enterprise Manager User Accounts	Creating Enterprise Manager User Accounts
Step 4	Setting Up My Oracle Support Credentials	(Optional) Setting Up My Oracle Support
Step 5	<i>Additional /Value Add setup (optional)</i> Configuring Self-Update	(Optional) Configuring Self-Update
Step 6	<i>Additional /Value Add setup (optional)</i> Setting Up E-Mail Notifications	(Optional) Setting Up E-mail Notifications

**Note:**

Ensure that the OMS is patched appropriately to the required level. For information about the patches that need to be applied on the Enterprise Manager Cloud Control Management Server (OMS) for using the Provisioning and Patching features, see My Oracle Support note 427577.1.

Setting Up Oracle Software Library

Oracle Software Library (Software Library) is one of the core features offered by Oracle Enterprise Manager Cloud Control (Cloud Control). Technically, it is a storage location that stores certified software entities such as software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. In addition to storing them, it also enables you to maintain versions, maturity levels, and states of these software entities.

To access the Software Library console page, in Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching** and then, click **Software Library**. On the Software Library home page, as shown in [Figure 2-2](#), there are two types of folders: Oracle-owned folders (marked by a lock symbol) and User-owned folders.

Figure 2-2 Software Library Console

Software Library Page Refreshed Feb 21, 2017 1:33:39 AM PST

Software Library maintains entities that represent software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. You can pick any of the Oracle-supplied entities, customize them or create a custom one of your own. Once defined, these reusable entities can be referenced from a Deployment Procedure to automate the patching, provisioning, or deployment of the associated software.

Actions View View Edit... Delete... Find Name Search

Name	Type	Subtype	Revision	Status	Maturity	Owner	Description
Software Library						ORACLE	Root Folder for Software Library entities
Adapters						ORACLE	Directives and Components belonging to the registered adapters
Bare Metal Provisioning						ORACLE	Bare Metal Provisioning directory
Business Intelligence Provisioning Utilities						ORACLE	SW Library Entities belonging to BI Provisioning
Cloud Custom Action Utilities						ORACLE	Cloud Custom Action
Cloud						ORACLE	Cloud
Coherence Node Provisioning						ORACLE	Coherence Node Provisioning Entities
Common Provisioning Utilities						ORACLE	Directives belonging to Common Provisioning (SIDB and RACPROV and any other)
CompositeDeploy						ORACLE	CompositeDeploy Entities
Core URL Certificate Monitoring						ORACLE	Core URL Certificate Monitoring entities
CRSRAC Provisioning						ORACLE	Directives and Components belonging to CRSRAC Provisioning (SICSRAC and RACPROV)
CVU Prerequisite-fixup components						ORACLE	CVU Prerequisite-fixup components belonging to DB Provisioning
Recovery Appliance						ORACLE	Directives and Components used for Recovery Appliance.
Database Provisioning CVU Downloads						SYSMAN	
Database Provisioning Profiles						SYSMAN	
DB Consolidation						ORACLE	DB Consolidation area folder display name
DB Provisioning						ORACLE	Directives and Components belonging to DB Provisioning
Database as a Service						ORACLE	Database templates base location
Diagnostic Tools						ORACLE	Diagnostic Tools Directory
dummy						SYSMAN	

To start using the Software Library to create and manage entities, the Software Library Storage Locations must be configured. System Administrators are responsible for configuring the Software Library storage locations, following which the Software Library becomes usable.

Cloud Control offers the following types of storage locations:

- **Upload File Locations:** These locations are configured for storing files uploaded by Software Library as part of creating or updating an entity. The Upload File Locations support two storage options:
 1. OMS Shared File System
 2. OMS Agent File System
- **Referenced File Locations:** These are locations that allow you to leverage the organization's existing IT infrastructure (like file servers, web servers, or storage systems) for sourcing software binaries and scripts. Such locations allow entities to refer to files without having to upload them explicitly to a Software Library storage. Referenced File Locations support three storage options:
 1. HTTP Locations
 2. NFS Locations
 3. Management Agent Locations

You can configure the storage locations from the Administration Console. To do so, in Cloud Control, from **Setup** menu, select **Provisioning and Patching**, then select **Software Library**. The Software Library Administration Page as shown in [Figure 2-3](#) appears:

Figure 2-3 Software Library Administration

Software Library: Administration Page Refreshed Feb 21, 2017 1:32:22 AM PST ↺

Software Library > Software Library: Administration

The administration console enables you to configure and administer the Software Library storage locations and cache nodes.

[Upload File Locations](#) [Referenced File Locations](#) [Cache Nodes](#)

Configure the storage locations that can be used for uploading files for Software Library entities.

Storage Type: OMS Shared File System

Configure file system locations on the OMS hosts. Make sure the locations are accessible by all the OMS instances, typically mounted or shared locations. You can optionally configure the common credential to be used by the Software Library for reading/writing from/to a location.

Actions ▼ View ▼ [+ Add...](#) [Edit...](#) [X Migrate and Remove](#)

Name	Status	Location	Associated Entities	Credential	Total Space	Available Space	Used Space	Deleted Entities Used Space	Last Refreshed
swlib	Active	/scratch/...view_storage/...j...	Show	Change Credential Clear credential	219.8...	17.32...	15.89...	0 Bytes	Feb 20, 2017 9:04:34 PM ...

 See Also:

For information on configuring Software Library, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*

 Note:

To run the procedure on a Windows host which involves executing some Software Library entities (for example, directive scripts), you (**the Windows user**) must be granted the following privileges:

- Act as part of the operating system
- Adjust memory quotas for a process
- Logon as batch job
- Replace a process level token

If not, the execution of the directive steps in the procedure may fail.

Setting Up Credentials

To perform any of the provisioning and patching tasks in Enterprise Manager Cloud Control, you need to set up Named Credentials for normal operating system user account (*Oracle*) and Named Credentials for privileged user accounts (*root*).

A Named Credential specifies a user's authentication information on a system. Named credentials can be a username and password pair such as the operating system login credentials, or the Oracle home owner credentials primarily used for performing operations such as running jobs, patching and other system management tasks.

Enterprise Manager Cloud Control enables you to register the system credentials as Named Credentials for normal user (*Oracle*). Alternately, if you have root privileges, you can even register the root account details as Named Credentials for the privileged users. Once they are registered as Named Credentials, you can save them as Preferred Credentials if you want.

The advantages of saving the credentials are:

- You do not have to expose the credential details to all the users.
- It saves time and effort as you do not have to specify the user name and password every time for each Oracle home or host machine, you can instead select a named profile that will use the saved credentials.

For more information on Named Credentials, see *Oracle Enterprise Manager Cloud Control Security Guide*.

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 the *Oracle* account credentials or the *root* account credentials may not be sufficient. Under such circumstances, use authentication utilities to run some steps within the Deployment Procedure with the privileges of another user. The authentication utilities supported by Enterprise Manager Cloud Control are SUDO and PowerBroker. This support is offered using the Privilege Delegation mechanism available in Enterprise Manager Cloud Control.

For a conceptual overview of Privilege Delegation and the authentication tools supported by it, see *Oracle Enterprise Manager Cloud Control Security Guide*.

[Table 2-2](#) lists the use cases pertaining to credentials, and describes the steps to be performed for setting up credentials for provisioning. Select the use case that best matches with the situation you are in, and follow the suggested instructions.

Table 2-2 Setting Up Enterprise Manager Credentials

Use Case	Steps to be performed
If you do not have direct access or the required credentials for the normal operating system user account (<i>Oracle</i>) OR If you do not have direct access or the required credentials for the privileged account (<i>root</i>).	Do the following: - Set up the Privilege Delegation as follows: - Create Privilege Delegation (PDP) Template either for SUDO or PowerBroker. To do so, see Creating Privilege Delegation section in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>. - Apply the created template on the Management Agents of the target hosts. - Create Named Credentials for normal operating system user account (<i>Oracle</i>) with privileges to run as SUDO or PowerBroker, for more information see Creating Named Credentials section in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>. OR Create Named Credentials for privileged users account (<i>root</i>) with privileges to run as SUDO or PowerBroker, for more information see Creating Privileged Credentials section in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>. - Save the Named credential for normal operating system account or the named credentials for the privileged user account as Preferred Credential. To do so, see Saving Preferred Credentials for Hosts and Oracle Homes and Saving Preferred Credentials to Access My Oracle Support sections in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>.

Table 2-2 (Cont.) Setting Up Enterprise Manager Credentials

Use Case	Steps to be performed
If you have direct access or the required credentials for the normal operating system user account (<i>Oracle</i>)	Do the following: 1. Create Named Credentials for normal operating system user account (<i>Oracle</i>), for more information see Creating Named Credentials section in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>.
OR	OR
If you have direct access or the required credentials for the privileged account (<i>root</i>).	Create Named Credentials for privileged user accounts (<i>root</i>) Credentials, for more information see Creating Privileged Credentials section in <i>Oracle Enterprise Manager Cloud Control Security Guide</i> . 2. Save the Named credential for normal operating system account or the named credentials for the privileged user account as Preferred Credential. To do so, see Saving Preferred Credentials for Hosts and Oracle Homes and Saving Preferred Credentials to Access My Oracle Support sections in <i>Oracle Enterprise Manager Cloud Control Security Guide</i>.


Creating Enterprise Manager User Accounts

This section describes the following:

- [Overview of User Accounts](#)
- [Creating Designer User Account](#)
- [Creating Operator User Account](#)

Overview of User Accounts

From the Cloud Control, you can create and manage new Enterprise Manager Administrator accounts. Each administrator account includes its own login credentials, as well as a set of roles and privileges that are assigned to the account.



Based on the accesses, the users can be classified as follows:

- Super Administrator
- Designers (EM_ALL_DESIGNER)
- Operators (EM_ALL_OPERATOR)

Super Administrators

Super Administrators are powerful Cloud Control administrators with full access privileges on all targets. They are responsible for creating and administering accounts within the Cloud Control environment. For example, Super Administrators create the Designer and Operator roles, and grant these roles to different users and groups within their enterprise.

Designers

Designers are lead administrators with increased privileges on Deployment Procedures and Software Library. Starting with Cloud Control, designers can create deployment procedure templates using the **Lock down** feature, and save these templates to enforce standardization and consistency. Operator privileges are granted on these templates so that administrators who login as Operators can launch these templates, and run the Deployment Procedure successfully. Doing this ensures that the procedures are less error prone and more consistent.

For more information about saving deployment procedures using lock downs, see [Saving and Launching the Deployment Procedure with Lock Down](#).

Designers are responsible for performing all the design-time activities such as:

- Creating the provisioning profiles in the Software Library.
- Creating components, directives, and images, and storing them in Oracle Software Library.

- Customizing the default deployment procedures according to the needs of the organization.
- Creating patch plans and patch templates.

The predefined Oracle role for a Designer is `EM_ALL_DESIGNER`, this role in turn includes fine grained roles where you can specifically set `EM_PROVISIONING_DESIGNER` for provisioning tasks, and `EM_PATCH_DESIGNER` for patching tasks. For more information about privilege grants to Designers, see [Granting Roles and Privileges to Administrators](#).

Operators

Operators are administrators who have restricted privileges on a Deployment Procedure and Software Library. Normally, operators can view and submit a deployment procedure. The Designer user may also grant the Operator the necessary privileges on any targets or entities.

Operators use the infrastructure created by designers and perform run-time activities such as:

- Accessing the provisioning profiles present in the Software Library for provisioning procedures.
- Launching software deployments to provision software on selected targets.
- Patching software deployments using patch plans and patch templates.

The predefined Oracle role for an Operator is `EM_ALL_OPERATOR`, this role in turn includes fine grained roles where you can specifically set `EM_PROVISIONING_OPERATOR` for provisioning tasks, and `EM_PATCH_OPERATOR` for patching tasks. For more information about privilege grants to Operators, see [Granting Roles and Privileges to Administrators](#).



Note:

Designers can choose to perform both design-time and run-time activities, but operators can perform only run-time activities.

Creating Designer User Account

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

1. In Cloud Control, from the **Setup** menu, select **Security**, then select **Administrators**.
2. On the Administrators page, click **Create**.
3. In the Create Administrator wizard, do the following:
 - a. On the Properties page, specify the name *Designer* and provide a password. Leave the other fields blank, and click **Next**.
 - b. On the Roles page, select **EM_ALL_DESIGNER**, and click **Next**.

 Note:

You can alternately restrict the Designer access to either Provisioning or Patching domains. For granting privileges explicitly for Provisioning, select the EM_PROVISION_DESIGNER role. Similarly, for granting designer privileges explicitly for Patching, select the EM_PATCH_DESIGNER role.

- c. On the Target Privileges page, select the targets privileges that must be granted to a Designer user account. For information about the target privileges available to an Administrator with Designer role, see [Granting Roles and Privileges to Administrators on the Deployment Procedure](#).
- d. On the Resource Privileges page, select the privileges to be explicitly granted for each of the resource types.
- e. On the Review page, review the information you have provided for this user account, and click **Finish**.

Creating Operator User Account

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

1. In Cloud Control, from the **Setup** menu, select **Security**, then select **Administrators**.
2. On the Administrators page, click **Create**.
3. 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 **EM_ALL_OPERATOR**, and click **Next**.

 Note:

You can alternately restrict the Operator access to either Provisioning or Patching domains. For granting privileges explicitly for Provisioning, select the EM_PROVISION_OPERATOR role. Similarly, for granting designer privileges explicitly for Patching, select the EM_PATCH_OPERATOR role.

- c. On the Target Privileges page, select the targets privileges that must be granted to an Operator user account. For information about the target privileges available to an Administrator with Operator role, see [Granting Roles and Privileges to Administrators on the Deployment Procedure](#).
- d. On the Resource Privileges page, select the privileges to be explicitly granted for each of the resource types.
- e. On the Review page, review the information you have provided for this user account, and click **Finish**.

(Optional) Setting Up My Oracle Support

For Cloud Control to connect to My Oracle Support for Agent Patching, patching other targets, MOS related tasks, and for Self-Update tasks, you must ensure that you set the proxy server settings and register the details. To do so, follow the instructions outlined in [Registering the Proxy Details for My Oracle Support](#).



Note:

Beginning with Enterprise Manager Cloud Control 12c Release 3 (12.1.0.3), My Oracle Support accesses support.oracle.com directly. This means that you must provide network access to this URL, or grant proxy access to it from any client that will access My Oracle Support.

(Optional) Configuring Self-Update

The Self Update feature enables you to obtain information about updates to Cloud Control components. The Self Update home page can be used to obtain information about new updates and provides a common workflow to review, download and apply the updates. The Self Update console automatically informs you whenever new updates that are applicable to your installation are made available by Oracle.

Software Library components and directives that you can use for provisioning and patching are called provisioning entities. A Provisioning bundle refers to a specific provisioning or patching area, such as database provisioning or FMW provisioning through which Cloud Control delivers updates to customers.



Note:

Ensure that the user has VIEW_ANY_SELFUPDATE privileges.

For applying Oracle-supplied updates to provisioning entities, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Extensibility**, then select **Self Update**.
2. Schedule to download provisioning bundle. The Self-update framework downloads the bundle to a well-defined location. For more information about Self-Update, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
3. From the **Actions** menu, select **Subscribe** to ensure that you receive notification whenever a provisioning bundle is available for download.
4. In the Updates Home page, select update of Type **Provisioning Bundle** and from the **Actions** menu, select **Open**.
5. Apply the provisioning bundle updates manually. Follow instructions as per selected provisioning bundle to apply the update manually.
6. In the Updates Home page, verify that the update is applied.

(Optional) Setting Up E-mail Notifications

Cloud 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 email notifications, see [Customizing Deployment Procedures](#).

(Optional) Setting Restricted Accesses for the Root Components

This section describes how you can perform certain lifecycle management tasks like Provisioning and Patching in Enterprise Manager using restricted access. In order to run some root commands, you either need to provide the restricted root access to a user (usually Management Agent user) or run the command that require root access manually.

Primarily, this section covers:

- [Root Components for Patching](#)
- [Root Components for Provisioning](#)

Root Components for Patching

This section covers:

- [Manually Staging the Root Components](#)
- [Restricting the Root User Access](#)

Manually Staging the Root Components

By default, the root component is automatically staged on the target host (at the location defined by %emd_emstagedir%). However, if you want to manually stage the root component at a custom location before initiating the patching process, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Navigate to **Patching And Provisioning**, and then **Components**.
3. Select **Root Dispatcher**, and then click **View**.
4. Click **Stage**.
5. Specify the host (where you want to stage the root dispatcher), and the dispatcher location, that is, the location on the specified host where you want to stage the root dispatcher. Click **Submit** to manually stage the root dispatcher.
6. Navigate to **Patching And Provisioning**, and then **Components**.
7. Select **Root Scripts**, and then click **View**.

8. Click **Stage**.
9. Specify the host (where you want to stage the root component), and the dispatcher location, that is, the location on the specified host where you want to stage the root component. Click **Submit** to manually stage the root component.

 Note:

Before you manually stage the root components at a custom location, ensure that the user has the *read* and *execute* permissions on the dispatcher location and on `root_dispatcher.sh`.

10. Decompress the file staged in the previous step.

 Note:

After staging the root component manually, you must specify this in the patch plan that you create for applying the required Oracle patches. Access the Deployment Options page of the patch plan that you created. In the Where to Stage section, select **No (already staged)** for **Stage Root Component**, and specify the dispatcher location where you have staged the root component manually. If the dispatcher location is a shared location, select **Dispatcher Location Shared**.

Restricting the Root User Access

In Enterprise Manager Cloud Control, you can provide restricted *root* access to the Management Agent user, such that the Management Agent user can only run certain commands as *root*, and not all commands. You can ensure this by editing the `etc/sudoers` policy file or, optionally in LDAP.

To patch the Oracle database and middleware targets, the Management Agent user must be able to run the `perl`, `root_dispatcher.sh` and `id` entities as the *root* user. These entities form a part of the patching root component.

To provide restricted *root* access to the Management Agent user, such that this user can only run the `perl`, `root_dispatcher.sh` and `id` entities as the *root* user, edit the `etc/sudoers` file such that it has the following contents:

```
Cmnd_Alias <command_alias> = \
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-
    FAULT_SUBACTION DEFAULT_ACTION perl *, \
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-
    FAULT_SUBACTION DEFAULT_ACTION id, \
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-
    FAULT_SUBACTION DEFAULT_ACTION <dispatcher_loc>/root_dispatcher.sh *

aime ALL=(root) <command_alias>
```

<agent_home> represents the Management Agent home.

 Note:

If the validation of privileged credentials from the EM Patch Plan wizard is not required then do the following:

1. Set the value of *check_privileged_sudo_credentials* to false in the *mgmt_patching_properties* table in the EM Repository.
2. The */etc/sudoers* can only contain the following:

```
Cmdnd_Alias <command_alias> = \
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNC-
    TIONALITY DEFAULT_SUBACTION DEFAULT_ACTION <dispatcher_loc>/
    root_dispatcher.sh *
aime ALL=(root) <command_alias>
```

 Note:

<dispatcher_loc> represents the root dispatcher location where the root component is staged. By default, the root component is automatically staged at %emd_emstagedir%. If you chose a custom location for the automatic staging, or staged the root component manually at a custom location, ensure that you specify this location for <dispatcher_loc>. Else, specify the default location, that is, %emd_emstagedir%.

 Note:

For information on commands to be run as *root* for patching RAC and Grid Infrastructure databases, see [Oracle Grid Infrastructure and Oracle RAC Configuration Support](#).

Root Components for Provisioning

This section covers the following:

- [Manually Staging the Root Components](#)
- [Restricting the Root User Access](#)

Manually Staging the Root Components

By default, the root component is automatically staged on the target host (at the location defined by %emd_emstagedir%). However, if you want to manually stage the root component at a custom location before initiating the provisioning process, follow these steps:

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

2. Navigate to **Patching And Provisioning**, and then **Components**.
3. Select **Root Dispatcher**, and then click **View**.
4. Click **Stage**.
5. Specify the host (where you want to stage the root dispatcher), and the dispatcher location, that is, the location on the specified host where you want to stage the root dispatcher. Click **Submit** to manually stage the root dispatcher.
6. Navigate to **Patching And Provisioning**, and then **Components**.
7. Select **Root Scripts**, and then click **View**.
8. Click **Stage**.
9. Specify the host (where you want to stage the root component), and the dispatcher location, that is, the location on the specified host where you want to stage the root component. Click **Submit** to manually stage the root component.

 Note:

Before you manually stage the root components at a custom location, ensure that the user has the *read* and *execute* permissions on the dispatcher location and on `root_dispatcher.sh`.

10. Decompress the file staged in the previous step.

 Note:

After staging the root components manually, you must access the Select Software Locations page from the Provision Oracle Database wizard. In the **Root Dispatcher Location**, specify the same dispatcher location where you have staged the root components manually, and select **Select this option if all the root scripts are staged to ROOT_DISPATCH_LOC** already.

If you have *not* manually staged the root scripts component, then you can use the database provisioning workflow to specify the details. In the Select Software Locations page of the Provision Oracle Database wizard, enter a location where you want to stage the root scripts. If you do not provide the location details, then a standard enterprise manager stage location will be used to stage the root components.

Restricting the Root User Access

In Enterprise Manager Cloud Control, you can provide restricted *root* access to the Management Agent user, such that the Management Agent user can only run certain commands as *root*, and not all commands. You can ensure this by editing the `etc/sudoers` policy file. Since provisioning involves invoking all the commands on the target host using Management Agent, you must include `nmosudo` in the `sudoers` file with the set of commands that is possible to run. To provide restricted *root* access to the

Management Agent user, edit the `etc/sudoers` file such that it has the following contents:

```
Cmnd_Alias <command_alias> = \  
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-  
FAULT_SUBACTION DEFAULT_ACTION perl *, \  
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-  
FAULT_SUBACTION DEFAULT_ACTION <dispatcher_loc>/root_dispatcher.sh *,\  
    <agent_home>/sbin/nmosudo DEFAULT_PLUGIN DEFAULT_FUNCTIONALITY DE-  
FAULT_SUBACTION DEFAULT_ACTION id  
  
aime ALL=(root) PROVISIONING_DISPATCHER
```

<command_alias> represents the alias that describes the entities that can be run by aime as *root*.

<agent_home> represents the Management Agent home.

<dispatcher_loc> represents the root dispatcher location where the root component is staged. By default, the root component is automatically staged at `%emd_emstagedir%`. If you chose a custom location for the automatic staging, or staged the root component manually at a custom location, ensure that you specify this location for <dispatcher_loc>. Else, specify the default location, that is, `%emd_emstagedir%`.

Part II

Discovery

This part contains the following chapter:

- [Discovering Hosts and Software Deployments](#)

3

Discovering Hosts and Software Deployments

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

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

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

Discovering Hosts Automatically and Adding Targets Manually

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

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

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

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

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

Discovering Hosts Manually and Adding Targets Manually

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

consume resources on the Oracle Management Agent to perform discovery when it is not needed.

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

Part III

Database Provisioning

This part contains the following chapters:

- [Overview of Database Provisioning](#)
- [Creating and Provisioning Oracle Databases](#)
- [Provisioning Oracle Grid Infrastructure for Oracle Databases](#)
- [Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases](#)
- [Provisioning Oracle Real Application Clusters One \(Oracle RAC One\) Node Databases](#)
- [Provisioning Oracle Real Application Clusters for 12c and 18c](#)
- [Extending Oracle Real Application Clusters](#)
- [Deleting or Scaling Down Oracle Real Application Clusters](#)
- [Provisioning Oracle Database Replay Client](#)
- [Provisioning Oracle Standby Databases](#)
- [Managing Pluggable Databases Using Enterprise Manager](#)
- [Cloning Oracle Databases and Pluggable Databases](#)
- [Database High Availability REST API](#)

4

Overview of Database Provisioning

Provisioning involves repeatable, reliable, automated, unattended, and scheduled mass deployment of software, applications, or servers across different platforms, environments, and locations.


This chapter provides an overview of the database provisioning feature in Oracle Enterprise Manager Cloud Control (Cloud Control), supported targets and deployment procedures offered by Cloud Control, and the infrastructure you need to set up to get started with database provisioning. In particular, this chapter covers the following:

- [Introduction to Database Provisioning](#)
- [Supported Use Cases and Targets Provisioned Using Database Provisioning Procedures](#)
- [Setting Up Database Provisioning](#)

Introduction to Database Provisioning

The Provisioning solution is an important part of Lifecycle Management solution offered by Cloud Control. As part of the database provisioning solution, Cloud Control enables you to provision Oracle Databases (also known as single-instance databases) and Oracle Real Application Clusters (RAC) databases, extend or delete Oracle RAC nodes, provision Oracle Real Application Clusters One node databases, provision pluggable databases, and also upgrade Oracle single-instance databases in a scalable and automated manner. The following diagram shows the database provisioning solution in Cloud Control.

Figure 4-1 Database Provisioning Solution in Cloud Control



For this release, database provisioning features are as follows:

Designer and Operator Roles

Cloud Control offers clearly defined administrator roles such as *Designers* and *Operators*. With a *Designer* role, you can lock down one or more fields in the deployment procedures, thus allowing the operators to run the procedure only with standard, pre-set configurations.

Locking Down Feature in Designer Role

The locking down feature in Database Provisioning enables Designers to lock down the set of variables, such as host targets, credentials, Oracle homes to be provisioned, and others, in the deployment procedure wizard. This enforces standard deployments and minimizes errors in configurations during mass deployment. The operator can then deploy the procedure that the designer configures and saves in the Procedure Library, see [Saving and Launching the Deployment Procedure with Lock Down](#).

Provisioning Profiles and Database Templates

You can create Provisioning Profiles to be used in database provisioning to ensure standardization in deployments and to minimize errors. You can also create database templates from the Cloud Control Console to be used in your provisioning activities.

Creating Databases Using Cloud Control

Cloud Control now enables you to create databases from the Cloud Control console. This ensures that you can use a single interface for provisioning and creating databases, see [Creating Databases](#).

Easy to Navigate Database Provisioning Wizards

Designers and Operators can easily use and navigate through the enhanced Database Provisioning wizards in Cloud Control.

Self Update

Using the Self Update feature, you can automatically download and install updates to your provisioning entities, see [\(Optional\) Configuring Self-Update](#).

Database Provisioning Console for all Database Provisioning Activities

The Database Provisioning console is a starting point for your database provisioning activities. The console displays information about provisioning setup, profiles, deployment procedures, and information about getting started with provisioning.

Supported Use Cases and Targets Provisioned Using Database Provisioning Procedures

Cloud Control enables you to perform database provisioning using deployment procedures. A deployment procedure is a set of predefined steps that complete the task of provisioning, see [About Deployment Procedures](#).

[Table 4-1](#) lists the database deployment procedures offered by Cloud Control and the various targets that can be provisioned.

Table 4-1 Database Deployment Procedures and Targets Provisioned

Deployment Procedure	Targets Provisioned
Provision Oracle Database	• Oracle Database (single instance) 10g Release 1 and above • Oracle Grid Infrastructure 11g Release 2 and above • Oracle Automatic Storage Management (Oracle ASM) 12c Release 1 and above • Oracle Restart (Single Instance High Availability) 12c Release 1 and above
Provision Oracle Real Application Clusters	• Oracle Real Application Clusters (Oracle RAC) 10g Release 1 and above • Oracle RAC One Node 11g Release 2 and above • Oracle Grid Infrastructure 11g Release 2 and above
Create Oracle Database	• Oracle Database (single-instance database) 10g Release 1 and above • Oracle Real Application Clusters (Oracle RAC) 11g Release 1 and above • Oracle RAC One Node 11g Release 1 and above
Provision Pluggable Database Note: You cannot manually invoke this deployment procedure from the Database Provisioning page. This deployment procedure is invoked internally when you use the Provision Pluggable Database console or EM CLI to provision pluggable databases.	Pluggable Databases (available only as a part of Oracle Database 12c Release 1 and above)
Provision Oracle Clusterware / Oracle RAC for UNIX and RDBMS versions 10g/11g Release 1 (applicable for UNIX platform)	• Oracle Real Application Clusters (Oracle RAC) 10g Release 1 up to 11g Release 1 • Oracle Clusterware 10g Release 1 up to 11g Release 1 • Oracle Automatic Storage Management (Oracle ASM) 10g Release 1 up to 11g Release 1
Provision Oracle Clusterware / Oracle RAC for Windows and RDBMS versions 10g/11g (applicable for Windows platform)	• Oracle Real Application Clusters (Oracle RAC) 10g Release 2 up to 11g Release 1 • Oracle Clusterware 10g Release 2 up to 11g Release 1 • Oracle Automatic Storage Management (Oracle ASM) 10g Release 2 up to 11g Release 1
Extend/Scale Up Oracle Real Application Clusters	Oracle Real Application Clusters (Oracle RAC) 10g Release 1 and above
Delete/Scale Down Oracle Real Application Clusters	Oracle Real Application Clusters (Oracle RAC) 10g Release 1 and above
Provision Oracle Database Client	Oracle Database Client 10g Release 2 and above

Table 4-2 lists various use cases for database provisioning deployment procedures.

Table 4-2 use cases for Database Deployment Procedures

Deployment Procedure	use case	Link
Provision Oracle Database	- Provisioning and Creating Single-Instance Databases - Provisioning Single-Instance Database with Oracle Automatic Storage Management - Provisioning Single-Instance Database Software Only - Provisioning Oracle Grid Infrastructure with Single-Instance Database and Configuring Database with Oracle Automatic Storage Management - Provisioning Oracle Grid Infrastructure and Single-Instance Database Software Only	Provisioning and Creating Oracle Databases Provisioning Oracle Databases with Oracle Automatic Storage Management Provisioning Oracle Database Software Only Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management Provisioning Oracle Grid Infrastructure and Oracle Database Software Only
Provision Oracle Real Application Clusters	- Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management - Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster - Provisioning Oracle Real Application Clusters Database with File System on a New Cluster	Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster Provisioning Oracle Real Application Clusters Database with File System on a New Cluster
Create Oracle Database	- Creating Single-Instance Database - Create Oracle Real Application Clusters Database - Creating Oracle Real Application Clusters One database	Creating an Oracle Database Creating Oracle Real Application Clusters Database Creating Oracle Real Application Clusters One Node Database
Provision Pluggable Database	- Creating a New Pluggable Database - Plugging In an Unplugged Pluggable Database - Cloning a Pluggable Database - Migrating a Non-Container Database as a Pluggable Database - Unplugging and Dropping a Pluggable Database	Creating a New Pluggable Database Using Enterprise Manager Plugging In an Unplugged Pluggable Database Using Enterprise Manager Cloning a Pluggable Database Using Enterprise Manager Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager Unplugging and Dropping a Pluggable Database Using Enterprise Manager
Provision Oracle Clusterware / Oracle RAC for Windows and RDBMS versions 10g/11g	- Cloning a Running Oracle Real Application Clusters - Provisioning Oracle Real Application Clusters Using Gold Image	Cloning a Running Oracle Real Application Clusters Provisioning Oracle Real Application Clusters Using Gold Image
Provision Oracle Clusterware / Oracle RAC for UNIX and RDBMS versions 10g/11g	Provisioning Oracle Real Application Clusters Using Archived Software Binaries	Provisioning Oracle Real Application Clusters Using Archived Software Binaries

Table 4-2 (Cont.) use cases for Database Deployment Procedures

Deployment Procedure	use case	Link
Extend/Scale Up Oracle Real Application Clusters	Extending Oracle Real Application Clusters	Extending Oracle Real Application Clusters
Delete/Scale Down Oracle Real Application Clusters	Deleting Oracle Real Application Clusters	Deleting the Entire Oracle RAC Scaling Down Oracle RAC by Deleting Some of Its Nodes
Provision Oracle Database Client	- Cloning a Running Oracle Database Replay Client - Provisioning Oracle Database Replay Client Using Gold Image - Provisioning Oracle Database Replay Client Using Installation Binaries	Cloning a Running Oracle Database Replay Client Provisioning an Oracle Database Replay Client Using Gold Image Provisioning an Oracle Database Replay Client Using Installation Binaries

Setting Up Database Provisioning

You can provision Oracle Databases, Oracle Real Application Clusters Databases, and Oracle RAC One Node Databases using database templates, installation media, or database entities, or you can use provisioning profiles to standardize deployments.

This section explains the following:

- [Meeting Basic Infrastructure and Host Requirements](#)
- [Understanding Administrator Privileges for Provisioning Database](#)
- [Prerequisites for Designers](#)
- [Prerequisites for Operators](#)
- [Creating Database Provisioning Profiles](#)
- [Viewing, Creating, and Deleting Database Provisioning Profiles Using EMCLI](#)
- [Creating Installation Media](#)
- [Creating Database Templates](#)
- [Uploading Database Templates to Software Library](#)
- [Creating Database Provisioning Entities](#)
- [Downloading Cluster Verification Utility](#)

Note:

- Primary group of the OS user owning the DB/Grid software installation and Primary group of OS user owning Enterprise Manager Management Agent installation must be same.

Meeting Basic Infrastructure and Host Requirements

To satisfy these requirements, you must do the following:

- Meet the basic infrastructure requirements as described in [Setting Up Your Infrastructure](#).
- Ensure that the host is set up for database provisioning entities. For more information about host readiness, see [Checking Host Readiness Before Provisioning or Patching](#).
- If you plan to provision database software on a Microsoft Windows host, you must ensure that Cygwin is installed on the host, before provisioning the database software, see *Enterprise Manager Cloud Control Basic Installation Guide*.

Understanding Administrator Privileges for Provisioning Database

[Table 4-3](#) describes the roles and the minimum privileges required for using database deployment procedures. These roles are default roles available in Cloud Control. You need not create them, but you must explicitly create administrators based on these roles. For instructions, see [Creating Enterprise Manager User Accounts](#).

Table 4-3 Privileges for Using Deployment Procedures

Role	Target Privileges	Resource Privileges	Implementation Recommendation
EM_PROVISIONING_DESIGNER	Operator any target	• Resource Type: Deployment Procedure Privilege: Create, Manage Launch Access • Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity	Required when you want to grant and restrict access to deployment procedures.
EM_PROVISIONING_OPERATOR	• Operator any target • Launch DP Permission	• Resource Type: Deployment Procedure Privilege: Create, Manage Launch Access • Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity	Required when you want to launch deployment procedures.

Prerequisites for Designers

Following are the prerequisites for designers to start database provisioning:

- Ensure that you meet the mandatory infrastructure requirements described in [Setting Up Your Infrastructure](#).
- Discover and monitor the destination hosts in Cloud Control. For this purpose, you need the latest version of Oracle Management Agent (Management Agent) on the

destination hosts. For more information refer to the Oracle Cloud Control Installation and Basic Configuration Guide. Ensure that the agents are installed in the same location on all hosts.

- Set up the Oracle Software Library (Software Library). Ensure that the installation media, database templates, or provisioning entities are available in the Software Library. For information about creating them, see [Setting Up Database Provisioning](#). Alternatively, use a provisioning profile to store the database template, see [Creating Database Provisioning Profiles](#).
- Store the operating system credentials of the destination hosts as preferred credentials in Oracle Management Repository (Management Repository) or use Named Credentials.

If you are using SUDO, PowerBroker, see [Setting Up Credentials](#) for information on setting up these authentication utilities.

- Ensure that the operating system groups corresponding to the following roles already exist on the hosts you select for provisioning. 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.

The Oracle Database user (typically *oracle*) must be a member of the following groups:

- Inventory Group (OINSTALL) as the primary group
- Database Administrator (OSDBA)
- Database Operator (OSOPER)

The Grid Infrastructure user (typically *grid*) must be a member of the following groups:

- Inventory Group (OINSTALL) as the primary group
- ASM Database Administrator (ASMDBA)
- ASM Instance Operator (ASMOPER)
- ASM Instance Administrator (OSASM)
- Ensure that you use an operating system user that has write permission on the following locations:
 - 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 you have Operator -Any Target privileges in Cloud Control.
- For provisioning Oracle Real Application Clusters Databases (Oracle RAC), the following are additional prerequisites:

- Meet the hardware, software, and network requirements for Oracle Grid Infrastructure and Oracle RAC installation on the target hosts.
- The Oracle RAC Database user must be a member of the group ASM Database Administrator (ASMDBA).

Prerequisites for Operators

Operators who run the deployment procedures must meet the following prerequisites:

- Ensure that as an operator, you have permissions to view credentials (set and locked by the designer), view targets, submit jobs, and launch deployment procedures.
- Ensure that the operating system groups corresponding to the following roles already exist on the hosts you select for provisioning. The operating system users of these groups automatically get the respective privileges.
 - Inventory Group (OINSTALL)
 - ASM Database Administrator (ASMDBA)
 - ASM Instance Operator (ASMOPER)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)
 - ASM Instance Administrator (OSASM)
- Ensure that you have Operator -Any Target privileges in Cloud Control.

Creating Database Provisioning Profiles

Provisioning Profile is an entity that contains software bits and configuration. When a provisioning profile is created from an existing installation, it provides the flexibility to clone either Grid Infrastructure (with software or configuration) and Oracle Database (with software or configuration) to a gold standard pre-determined by the administrator. You can create database templates using provisioning profiles. A designer or administrator can create a database provisioning profile as a one-time activity that can be used by operators for mass deployment. Using provisioning profile enables standardization in deployments and reduces the need for rescheduling deployments by avoiding errors while configuring deployment procedures.



Note:

You do not require out of box profiles anymore. Provisioning profiles for 11.2.0.4 Gold Image can be created using the gold image flow. You can also create installation media based profiles for any version of grid infrastructure and database.


If a database is used as a reference for a Gold Image, the new profile will contain database data. If the reference database is not in ARCHIVE LOG MODE, then the reference database will be shutdown and restarted during the process.

To create database provisioning profile, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, in the Profiles section, click **Create**. The Create Database Provisioning Profile wizard is launched.
3. On the Reference Target page, click the search icon.
4. On the Search and Select: Targets window, select the reference target from which you want to create the provisioning profile, and then click **Select**.
5. In the Reference target page, the Include Operation allows you to select the components you want to include in the provisioning profile. Depending on the reference host configuration, you can select to include Database Oracle Home, Grid Infrastructure Oracle Home, and their related configuration as part of the provisioning profile as below:
 - **Database Oracle Home** to include Oracle database gold image in the profile
 - **Grid Infrastructure Oracle Home** and its Configuration Properties to include Grid infrastructure gold image and its configuration in the profile
 - **Data Content** to include Oracle database template (or Data) in the profile where you can **Select Structure and Data** to include physical and structural files from the database or **Structure Only** to include only the structural files in the template.

 Note:

- While including the Data Content, select the Database Template option present under either Structure and Data or Structure Only, options, you cannot select both.
- For instructions on creating profiles for RMAN backups, see [Cloning Oracle Databases and Pluggable Databases](#)



The screenshot shows the 'Create Database Provisioning Profile: Reference Target' wizard. The 'Reference Target' step is active, indicated by a blue dot in the progress bar at the top. The wizard has four steps: Reference Target, Content Options, Profile, and Review. The 'Reference Target' section contains the following fields and options:

- Name:** racupg
- Type:** Cluster Database
- Reference Node:** racdb01.us.oracle.com
- Include:**
 - ☒ Database Oracle Home (C:\ORACLE\product\11.2.0\jdk\bin\java.exe, 8179)
 - ☒ Grid Infrastructure Oracle Home (C:\ORACLE\product\11.2.0\jdk\bin\java.exe, 7120)
 - ☒ Data Content
 - ☐ Structure and Data
 - ☒ Structure Only

At the bottom, there are two buttons: 'Create' and 'Database Template'.

6. In the Credentials section, select **Preferred Credentials** or **Named Credentials**. If using Named Credentials, select the credentials or click + to specify new Named Credentials.

In the Add Credentials window, specify the **User Name** and **Password**. Select **Set as Preferred Credentials** if you want to set these as the Preferred Credentials. Click **Add**.

Click **Next**.

Credentials			
Target	Credential Type	Credential	Credential Name
addDatabase as oracle user	Database Home	Preferred Credentials ▼	Normal Host Credentials
addDatabase as oracle user	Grid Infrastructu...	Preferred Credentials ▼	Normal Host Credentials
Inventory	Cluster Database	Preferred Credentials ▼	SYSDBA Database Credentials

7. On the Profile page, do the following:
 - In the Profile Information section, enter a unique profile name of your choice.
For example:
Cluster Profile [time when created]
 - Retain or edit the default details such as **Profile Location** where you want to store the provisioning profile in the Software Library, **Name**, **Description**, **Version**, **Vendor**, **Notes**, and the **Name of Components** included in the profile.
 - In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
 - In the Working Directory section, specify the temporary working directory to be used during database provisioning profile creation.
 - In the Software Library Storage section, select the **Software Library Location Type** and **Software Library Location Name**.
 - Click **Next**.
8. In the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation. The Deployment Instance Name is generated with the profile name and user name.
9. Once you have submitted the provisioning profile creation procedure, ensure that the provisioning profile appears in the Database Provisioning page.

Viewing, Creating, and Deleting Database Provisioning Profiles Using EMCLI

This section explains the EMCLI commands used to view, create, and delete database provisioning profiles using EMCLI verbs. The following topics are covered here:

- [Viewing Database Provisioning Profiles Using EMCLI](#)

- [Creating Database Provisioning Profiles Using EMCLI](#)
- [Deleting Database Provisioning Profiles Using EMCLI](#)

Viewing Database Provisioning Profiles Using EMCLI

Use the following EMCLI verb to know the different provisioning profiles and the required input parameters for them:

```
emcli describe_dbprofile_input
```

This command generates a response file for the different types of profiles.

Creating Database Provisioning Profiles Using EMCLI

To create a provisioning profile, use the following EMCLI verb:

```
emcli create_dbprofile -input_file=data:"<Prop file name>"
```

For example:

```
emcli create_dbprofile -input_file=data:"/tmp/profile.txt"
```

This command takes in a property file that completely describes the type of profile that will be created and the options used. To know the input properties that need to be specified in the input file, execute the following command:

```
emcli describe_dbprofile_input [-data_mode={EXPORT/DBCA_TEMPLATE/RMAN/STORAGE_SNAPSHOT}]
```

After submitting a create profile procedure, perform the following steps to review the provisioning profile details:

1. Use the following EMCLI verb to get the running provisioning profile instance:

```
emcli get_instances
```

2. Use the GUID from the previous step to get the response file. For example:

```
emcli get_instance_data -instance=<GUID> >/tmp/profile.txt
```

3. The input properties are listed in /tmp/profile.txt. For example:

```
# Input properties are:
GI_GOLD_IMAGE_ENTITY_NAME=ust Profile 02-04-2014 08:03 PM - Grid infrastructure
gold image
GI_GOLD_IMAGE_TYPE=GOLD_IMAGE
PROFILE_DESC=Grid Infrastructure Home Reference Profile 02-04-2014 08:03 PM from
clustname
PROFILE_LOCATION=Grid Infrastructure Home Provisioning Profiles/11.2.0.2.0/
linux_x64
PROFILE_NAME=Cluster clustname Profile 02-04-2014 08:03 PM
PROFILE_NOTES=Host Name: h1.example.com
Cluster: slxaclust
PROFILE_VENDOR=Oracle
PROFILE_VERSION=11.2.0.2.0
REFERENCE_DATABASE=clustname
REFERENCE_DATABASE_TYPE=cluster
REF_GI_CREDENTIALS=AIME_NORMAL:SYSMAN
REF_HOST_CREDENTIALS=AIME_NORMAL:SYSMAN
REF_NODE_SELECTED=h1.example.com
STORAGE_NAME_FOR_SOFTWARE=swlib
```



```
STORAGE_TYPE_FOR_SOFTWARE=0msShared  
WORKING_DIRECTORY=/tmp
```

Deleting Database Provisioning Profiles Using EMCLI

To delete a provisioning profile and its subcomponents, follow these steps:

1. Use the following EMCLI verb to list the database profiles created:

```
emcli list_dbprofiles
```

2. Use the following EMCLI verb to delete the database profile:

```
emcli delete_dbprofile -comp_loc= "<db profile name and location>"
```

For example:

```
emcli delete_dbprofile -comp_loc="Grid Infrastructure Home Provisioning Profiles/  
11.2.0.2.0/linux_x64/Cluster clustername Profile 02-04-2014 08:03 PM"
```

comp_loc is the combination of the database profile name and the location of the profile.

3. To check the status of the profile deletion, run the following EMCLI command:

```
emcli get_instance_status -instance=<GUID> -xml -details -showJobOutput
```

Creating Installation Media

To create installation media that can be used for database provisioning, follow these steps:

1. Create a temporary location `mkdir /tmp/installmedia`.
2. Navigate to the following URL:
<http://www.oracle.com/technetwork/database/enterprise-edition/downloads/index.html>
3. Click the **See All** link for the operating system on which you want to provision the database.
4. Select **Accept License Agreement**.
5. Download zip files 1 and 2 for Database and Grid Infrastructure software to the temporary directory created earlier.
6. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching** and then select **Software Library**.
7. In Software Library, select the directory where you want to create the installation media component for the database.
8. From the **Actions** menu, select **Create Entity**, then select **Component**.
9. In the Create Entity: Component dialog, select Subtype as **Installation Media** and click **Continue**.
10. In the Create Installation Media: Describe page, enter the Name and Description for the component, and click **Next**.
11. In the Create Installation Media: Configure page, select **Product Version**, **Platform**, and **Product** from the list.

For Product, select **Oracle Database** for Oracle Database, **Oracle Grid Infrastructure** for Grid Infrastructure software, and **Oracle Client** for Oracle Database Client.

Click **Next**.

12. In the Create Installation Media: Select Files page, select **Upload Files**.
 - a. In the Specify destination section, choose a Software Library storage location as the **Upload Location** for the database software.
 - b. In the Specify Source section, select **File Source** as Agent Machine and select the host from which you want to upload the files.
 - c. Click **Add**.
 - d. In the Remote File Browser, click **Login As**.
 - e. Select the Host Credentials and click **OK**.
 - f. Navigate to the temporary directory and select the zipped database files that you created.
 - g. Click **Add** and then click **OK**.

Click **Next**.

13. In the Create Installation Media: Review page, review the details you have provided and click **Save and Upload** to save and upload the installation media files to Software Library.

Creating Database Templates

Cloud Control allows you to create database templates that you can use for cloning or creating additional databases. To create database templates, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. In the Databases page, click on the database from which you want to create a template.
3. In the Database home page, from the **Oracle Database** menu, select **Provisioning**, then select **Create Database Template**.
4. In Template Type page, select:
 - **Structure as well as data** to include physical data files and structural information in the template. User-defined schemas and data will be included in the template. Databases created from this type of template will be identical to the source database.
 - **Structure** to include structural information about the source database including tablespace options, initialization parameters, and data files. User-defined schemas and data will not be included in the template.

Select host credentials. You can select **Preferred Credentials**, **Named Credentials**, or **Enter Credentials**.

Click **Next**.

5. In the Template Options page, specify the **Template Name** and **Description**. Specify the template location:
 - Select **Store Template in Software Library** to specify the **Storage Type** and **Location** on the OMS Agent File System or Shared File System.

- Select **Store Template on the Managed Host** to store template at ORACLE_HOME/assistants/dbca/templates in the target Oracle home.

Specify the database file locations. Select:

- **Use Oracle Flexible Architecture** to convert the location of files in the template to OFA.
- **Maintain File Location** if you want the location of the files in the template to be identical to the source database.

Click **Next**.

6. In the Schedule page, specify the job name and schedule. If you want to run the job immediately, then retain the default selection, that is, One Time (Immediately). If you want to run the job later, then select One Time (Later) and provide time zone, start date, and start time details. You can also select to blackout the database during the template creation process. Click **Next**.
7. In the Review page, review the details you have provided for the job and if you are satisfied with the details, then click **Submit Job** to run the job 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.
8. In the Jobs page, verify that the job has successfully completed and the template has been created as specified.



Note:

You can also use Database Configuration Assistant (DBCA) for creating database templates.

You can edit and customize the database template you create and then upload the customized template to the Software Library, see [Uploading Database Templates to Software Library](#).

Uploading Database Templates to Software Library

You can edit and customize your database templates and then upload them to Software Library as follows:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library Home page, select the folder where you want to upload the database template.
3. From the **Actions** menu, select **Create Entity**, then select **Component**. Alternatively, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
4. From the Create Entity: Component dialog box, select **Database Template** and click **Continue**.
Cloud Control displays the Create DatabaseTemplate page.
5. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

Click **+Add** to attach the database template. Select the template as the Source file in the format *templatename.dbt* or *templatename.dbc*. Retain the File Name as displayed. Ensure that the file size is less than 2 MB.

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

6. On the Select Files page, add all the database template related files.

Select Upload Files to upload all the database template files as follows:

- a. In the Specify Destination section, choose the Software Library location where you want to upload the files.
- b. In the Specify Source section, select the location where you have stored the template files. The location can be your local machine or the agent machine.
- c. Click **+Add** to upload the database template files.

For Structure template, again add the *templatename.dbt* file. In case of Structure And Data template, upload the *templatename.dbc*, *datafiledump.dfb* and the *controlfile.ctl* files. Mark the *templatename.dbc* file as the Main File.

Select Refer Files to refer files from an existing referenced file storage location. Select the Referenced File Location and add the source file.

7. On the Review page, review the details and then click **Save and Upload** to create the component and upload the binary to Software Library.

Creating Database Provisioning Entities

You can create and store provisioning entities in the Software Library to be used for provisioning Oracle databases. Cloud Control allows you to create the following types of database provisioning entities:

- Oracle Database Clone
- Oracle Clusterware Clone

The following subsections explain how to create these provisioning entities:

- [Creating an Entity for Oracle Database Software from a Shared NFS location](#)
- [Creating an Oracle Database Clone from an External Storage](#)
- [Creating an Oracle Clusterware Clone from a Reference Home](#)
- [Creating an Oracle Clusterware Clone from an External Storage](#)

Creating an Entity for Oracle Database Software from a Shared NFS location

To create an entity for Oracle Database Software from a shared NFS location, perform the following steps:

Step 1 - Create a Referenced File Location

To create a referenced file location, perform the following steps:

1. From the **Enterprise** menu in EM select **Provisioning and Patching** and then select **Software Library**.
2. Click the **Referenced File Locations** tab and select **NFS** as the Storage type.
3. Click **Add**.
4. In the Add NFS Location dialog provide the following details:
 - a. Name for the location - For example, NFS_SHARED_LIBRARY
 - b. NFS Server - IP or host name of the NFS server
 - c. Location – Absolute path of the NFS share where you intend to store the Software Binaries/Archives/Gold Images.

Step 2 - Create an Oracle Home Zip Archive

To create an Oracle home zip archive, perform the following steps:

1. Create a zip archive of the Oracle home that you want to use for creating DB Gold Image, run the following command:
 - a. `cd $ORACLE_HOME`
 - b. `zip -r -9 -1 <path to the new zip file> ./*`
 For example : `cd $ORACLE_HOME; zip -r -9 -1 /scratch/vol1/dbhome121010.zip ./*`
2. Copy the zip archive to the NFS Share configured as the **Referenced File Location**.

Step 3 - Create a Software Entity

To create a software entity, perform the steps below:

1. From the **Enterprise** menu in EM, select **Provisioning and Patching** , and then select **Software Library**.
2. Select a location where you want to create Software Entity which will represent your DB gold Image.
3. From the **Actions** menu, select **Create Entity**, and then select **Component**.
4. Select **Oracle Database Software Clone** as the Subtype.
5. In the Oracle Database Software Clone: Describe section, provide a meaning full name and description for the Gold Image. For example, 'Oracle DB Home 12.1.0', and then click **Next**.
6. In the Oracle Database Software Clone: Configure section:
 - a. In the **Create Component From:** field, select **Existing Oracle Home Archive**.
 - b. Select **Reference Location Type** as NFS.
 - c. Select the NFS Referenced File Location name created earlier in the **Reference Location Name**.
 - d. Provide the path to the Oracle home archive zip file staged in the NFS Share relative to the NFS Share path in the **Oracle Archive Home Location**.
7. Select appropriate **Product**, **Version** and **Platform**, and click **Next**.
8. Review the details and click **Save and Upload**.

Creating an Oracle Database Clone from an External Storage

To create an Oracle Database Clone from an external storage, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library Home page, select any custom folder and create the database clone component.
3. From the **Actions** menu, select **Create Entity**, then select **Component**. Alternatively, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
4. From the Create Entity: Component dialog box, select **Oracle Database Software Clone** and click **Continue**.

Cloud Control displays the Create Oracle Database Software Clone page.

5. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

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

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

6. On the Configure page, from the Create Component from menu, select **Existing Oracle Home Archive** and do the following:
 - a. In the Oracle Home Archive section, select a external storage location from where you can refer to the database clone software. From the **External Storage Location Name** menu, select the location name, see [Setting Up Oracle Software Library](#).

In **Oracle Home Archive Location**, enter the exact path, which is basically the relative path from the configured location, of the archive file residing on the external storage location. Ensure that the archive file is a valid zip file.

Note:

To create the zip file of an Oracle Home, use the following syntax:

```
<ZIP PATH>/zip -r -S -9 -1 <archiveName.zip> <directory or list  
of files to be archived> -x <patterns to exclude files>
```

- b. In the Oracle Home Properties section, select the **Product**, **Version**, **Platform**, and **RAC Home** values, as these configuration properties are particularly useful to search or track an entity.

7. On the Review page, review the details and then click **Save and Upload** to create the component and upload the binary to Software Library.

Creating an Oracle Clusterware Clone from a Reference Home

To create an Oracle Clusterware Clone from a reference home, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library Home page, select any custom folder and create the database clone component.
3. From the **Actions** menu, select **Create Entity**, then select **Component**. Alternatively, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
4. From the Create Entity: Component dialog box, select **Oracle Clusterware Clone** and click **Continue**.

Cloud Control displays the Create Oracle Clusterware Clone: Describe page.

5. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

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

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

6. On the Configure page, from the Create Component from menu, select **Reference Home** and do the following:
 - a. In the Reference Oracle Home section, click the magnifier icon to select the desired Oracle Clusterware Oracle home from the list of Clusterware homes running on the host machine.

The **Oracle Home Location** and **Host** fields are populated with the selected values.

- b. In the Oracle Home Credentials section, select the credential type you want to use for accessing the targets you manage, see [Setting Up Credentials](#).
- c. In the Working Directory and Files to Exclude section, enter a **Working Directory** on the host on which you have write permissions, so that the cloned zip file can be created and placed there temporarily.

The **Files to exclude** field is pre-populated with certain types of files or patterns that will be excluded from the cloned zip file. However, you can customize this list based on your requirement.

- d. In the Software Library Upload Location section, select a configured storage location from the list where you want to place the Oracle Clusterware clone software, see [Setting Up Oracle Software Library](#).

7. On the Review page, review the details, and then click **Save and Upload** to create the component and upload the binary to the Software Library.

Creating an Oracle Clusterware Clone from an External Storage

To create an Oracle Clusterware Clone from a external storage location, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library Home page, select any custom folder and create the database clone component.
3. From the **Actions** menu, select **Create Entity**, then select **Component**. Alternatively, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
4. From the Create Entity: Component dialog box, select **Oracle Clusterware Clone** and click **Continue**.

Cloud Control displays the Create Oracle Clusterware Clone: Describe page.

5. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

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

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

6. On the Configure page, from the Create Component from menu, select **Existing Oracle Home Archive** and do the following:
 - a. In the Oracle Home Archive section, select a external storage location from where you can refer to the Oracle Clusterware clone software. From the **External Storage Location Name** menu, select the location name, see [Setting Up Oracle Software Library](#).

In **Oracle Home Archive Location**, enter the exact path, which is basically the relative path from the configured location, to the archive file residing on the external storage location. Ensure that the archive file is a valid zip file.

Note:

To create the zip file of an Oracle Home, use the following syntax:

```
<ZIP PATH>/zip -r -S -9 -1 <archiveName.zip> <directory or list  
of files to be archived> -x <patterns to exclude files>
```

- b. In the Oracle Home Properties section, select the **Product**, **Version**, and **Platform** values, as these configuration properties are particularly useful to search or track an entity.

7. On the Review page, review the details, and then click **Save and Upload** to create the component and upload the binary to Software Library.

Downloading Cluster Verification Utility

Cluster Verification Utility (CVU) performs system checks in preparation for installation, patch updates, or other system changes. You can synchronize cluster verification utility (CVU) binaries with Software Library.

Enterprise Manager, by default, provides a routine job that is scheduled daily to download binaries from My Oracle Support if corresponding binaries in the Software Library need to be updated. If your Enterprise Manager deployment is behind a firewall or a DMZ such that the HTTP connection to My Oracle Support is disabled, the routine job will skip its execution. In this case, you can manually download the CVU binaries corresponding to your platform from OTN or My Oracle Support using patch 16766985 as source. You can then synchronize these manually downloaded Cluster Verification Utility (CVU) binaries to Software Library as follows:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, click **CVU Setup**.
3. In the **Upload Prerequisites to Software Library** page, select one of the following:
 - a. **Local Machine** to select the CVU binaries from your local computer.
 - b. **Agent Machine** to select the CVU binaries from the agent machine.
4. Click **OK**. This will update the Software Library with the latest cluster verification utility binaries.

5

Creating and Provisioning Oracle Databases


This chapter explains how you can create databases (Oracle single-instance, Oracle RAC, and Oracle RAC One Node) or provisioning single-instance databases. In particular, this chapter covers the following:

- [Oracle Database Topology](#)
- [Creating Databases](#)
- [Provisioning Databases](#)

Oracle Database Topology

[Figure 5-1](#) shows a typical Oracle Database (single-instance database) topology that you can provision using Cloud Control.

Figure 5-1 Oracle Database Topology



The topology shows a 12.1.0.2 RDBMS managed by Single-Instance High Availability (SIHA) component of Grid Infrastructure 12.1.0.2. The software components of the topology are:

- Oracle High Availability daemons running from Grid Infrastructure home.
- Single-Instance Oracle ASM running from Grid Infrastructure home.
- Single-Instance Oracle database running from an Oracle Database Oracle home.

The hardware components of the topology are:

- A database host with a public interface.
- A dedicated storage network that links to the ASM disk raid.

Creating Databases

This section explains how you can create databases using Oracle Enterprise Manager Cloud Control (Cloud Control). It covers the following:

- [Getting Started with Creating Databases](#)
- [Creating an Oracle Database](#)
- [Creating Oracle Real Application Clusters Database](#)
- [Creating Oracle Real Application Clusters One Node Database](#)

Note:

This section also provides information about creating single-instance, Oracle Real Application Clusters (Oracle RAC), and Oracle Real Application Clusters One Node (Oracle RAC One Node) container databases.

You can create a container database on a host only if Oracle Database 12c Release 1 (12.1), or higher, is installed on the host. For more information on container databases, see *Oracle Database Administrator's Guide*.

Getting Started with Creating Databases

This section helps you get started with this chapter by providing an overview of the steps involved in creating databases. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully create a database using Cloud Control. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 5-1 Getting Started with Creating Oracle Databases

Step	Description	Reference Links
Step 1	Selecting the Use Case This chapter covers a few use cases for creating databases. Select the use case that best matches your requirements.	• To learn about creating Oracle Single Instance Database, see Creating an Oracle Database. • To learn about creating Oracle RAC Database, see Creating Oracle Real Application Clusters Database. • To learn about creating Oracle RAC One Node Database, see Creating Oracle Real Application Clusters One Node Database.

Table 5-1 (Cont.) Getting Started with Creating Oracle Databases

Step	Description	Reference Links
Step 2	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, and setting up of Oracle Software Library.	• To learn about prerequisites in creating Oracle Database, see Prerequisites for Creating an Oracle Database. • To learn about prerequisites in creating Oracle RAC Database, see Procedure for Creating an Oracle Real Application Clusters Database. • To learn about prerequisites in creating Oracle RAC One Node Database, see Prerequisites for Creating an Oracle RAC One Node Database.
Step 3	Running the Deployment Procedure Run the Deployment Procedure to successfully create the database.	• To create Single Instance Database, see Procedure for Creating an Oracle Database. • To create Oracle RAC Database, see Procedure for Creating an Oracle Real Application Clusters Database. • To create Oracle RAC One Node Database, see Procedure for Creating an Oracle Real Application Clusters One Node Database.

Creating an Oracle Database

This section provides information about creating an Oracle Database (also called single-instance database).



Note:

You can also use the information provided in this section to create a single-instance container database.

You can create a container database on a host only if Oracle Database 12c Release 1 (12.1), or higher, is installed on the host. For more information on container databases, see *Oracle Database Administrator's Guide*.

This section covers the following:

- [Prerequisites for Creating an Oracle Database](#)
- [Procedure for Creating an Oracle Database](#)

Prerequisites for Creating an Oracle Database

To create single-instance databases using Cloud Control, ensure that you meet the following prerequisites:

1. Ensure that you meet the infrastructure requirements explained in [Setting Up Your Infrastructure](#).

2. Ensure that you have created and stored a database template in the Software Library or Oracle Home. For information about creating database templates, see [Creating Database Templates](#).
3. Oracle Home for the database you want to create must be installed and you need to have credentials of the owner of the Oracle Home. If the Create Database wizard is launched from the Provision Database deployment procedure wizards, the Oracle Home need not be installed earlier. In such cases, the validations for Oracle Home will be skipped during the procedure interview and will be performed during execution of the deployment procedure.
4. The database plug-in that supports the corresponding database version should be deployed on OMS and Agent. For information about deploying plug-ins, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
5. Ensure that you have sufficient space to create the database, and that you have write permissions to the recovery file location.
6. If you are using a template from the Software Library for database creation, you must have Write permission to the Staging Location.
7. If you are using Automatic Storage Management (ASM) as storage, ASM instances and diskgroups must be configured prior to creating database.
8. The Cloud Control user creating the database template must have CONNECT_ANY_TARGET privilege in Cloud Control.

Procedure for Creating an Oracle Database

To create an Oracle database, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select the Create Oracle Database Deployment Procedure and click **Launch**. The Create Oracle Database wizard is launched.
3. In the Database Version and Type page, select the database **Version** and select **Oracle Single Instance Database**.

In the Hosts section, specify hosts and Oracle Home to provision the database. You can also specify Host Credentials and Common Oracle Home across all hosts. The Host Credentials can be Named or Preferred Credentials.

Click the plus (+) icon to add the host. Select the host and specify **Oracle Home**. Select **Host Credentials** or add new. Click the plus icon to add new credentials and specify **User Name**, **Password**, and **Run Privileges** and save the credentials.

Click **Next**.

4. In the Database Template page, choose the database template location. The location can be Software Library or Oracle Home. The template selected must be compatible with the selected Oracle Home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must exist on all hosts where you want to create the database.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle Home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

5. In the Identification and Placement page, specify database configuration details. Specify **Global Database Name** and **SID** prefix.

In the Database Consolidation section, select **Create As Container Database** if you want to create a container database. By default, an empty container database is created. If you want to add one or more pluggable databases to that container database, then select **Create a Container Database with one or more PDBs**, and set the number of PDBs.

If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all the cloned PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name accountsPDB, then the PDBs are created with the names accountsPDB1, accountsPDB2, accountsPDB3, accountsPDB4, and accountsPDB5.

Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP database accounts.

For database version 12.1 or higher, for Microsoft Windows operating systems, the database services will be configured for the Microsoft Windows user specified during Oracle home installation. This user will own all services run by Oracle software. In the Oracle Home Windows User Credentials section, specify the host credentials for the Microsoft Windows user account to configure database services. Select existing named credentials or specify new credentials. To specify new credentials, provide the user name and password. You can also save these credentials and set them as preferred credentials.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

 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 Cloud 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.

6. In the Storage Locations page, select the storage type, whether File System or Automatic Storage Management (ASM).

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 or /u01/product/db/oradata.

If you want to use ASM, then select **Automatic Storage Management (ASM)**, 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.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify up to nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

7. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

8. In the Additional Configuration Options, all the available listeners running from the Oracle Home and Grid Infrastructure listeners are listed. You can either select a listener or create a new one. You can select multiple listeners to register with the database. To create a new listener, specify the **Listener Name** and **Port**. Select database schemas and specify custom scripts, if any. Select custom scripts from the host where you are creating the database or from Software Library. If you have selected multiple hosts, you can specify scripts only from Software Library.

If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

9. In the Schedule page, specify a Deployment Instance name and a schedule for the deployment. If you want to run the procedure immediately, then retain the default selection, that is Immediately. If you want to run the procedure later, then select Later and provide time zone, start date, and start time details. Click **Next**.
10. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.

Click **Analyze** to check for prerequisites and to ensure that all the necessary requirements for provisioning are met.

11. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Creating Oracle Real Application Clusters Database

This section provides information about creating Oracle Real Application Clusters Database.

Note:

You can also use the information provided in this section to create a Oracle Real Application Clusters container database.

You can create a container database on a host only if Oracle Database 12c Release 1 (12.1), or higher, is installed on the host. For more information on container databases, see *Oracle Database Administrator's Guide*.

This section covers the following:

- [Prerequisites for Creating an Oracle Real Application Clusters Database](#)
- [Procedure for Creating an Oracle Real Application Clusters Database](#)

Prerequisites for Creating an Oracle Real Application Clusters Database

To create an Oracle RAC databases using Cloud Control, ensure that you meet the following prerequisites:

1. Ensure that you meet the mandatory infrastructure requirements explained in [Setting Up Your Infrastructure](#).
2. Ensure that you have created and stored the database template in the Software Library or Oracle Home. For information about creating database templates, see [Creating Database Templates](#).
3. Oracle Home for the database you want to create must be installed and you need to have credentials of the owner of the Oracle Home. If the Create Database wizard is launched from the Provision Database deployment procedure wizards, the Oracle Home need not be installed earlier. In such cases, the validations for Oracle Home will be skipped during the procedure interview and will be performed during execution of the deployment procedure.
4. The database plug-in that supports the corresponding database version should be deployed on OMS and Agent. For information about deploying plug-ins, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
5. Ensure that you have sufficient space to create the database, and that you have write permissions to the recovery file location.
6. If you are using a template from the Software Library for database creation, you must have Write permission to the Staging Location.

7. If you are creating Oracle Real Application Clusters database, you must have Grid Infrastructure installed and configured. If the Create Database wizard is launched from the Provision Database deployment procedure wizards, Grid Infrastructure need not be installed and configured. In such cases, the validations for Grid Infrastructure will be skipped during the procedure interview and will be performed during execution of the deployment procedure.
8. If you are using Automatic Storage Management (ASM) as storage, ASM instances and diskgroups must be configured prior to creating database.
9. The Cloud Control user creating the database template must have `CONNECT_ANY_TARGET` privilege in Cloud Control.

Procedure for Creating an Oracle Real Application Clusters Database

To create an Oracle Real Application Clusters (Oracle RAC) database, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select the Create Oracle Database Deployment Procedure and click **Launch**. The Create Oracle Database wizard is launched.
3. In the Database Version and Type page, select the database **Version** and select **Oracle Real Application Clusters (Oracle RAC) Database**.

In the Cluster section, select the Cluster and Oracle Home. Select a reference host to perform validations to use as reference to create database on the cluster.

Select **Cluster Credentials** or add new. Click the plus icon to add new credentials and specify **User Name**, **Password**, and **Run Privileges** and save the credentials.

Click **Next**.

4. In the Database Template page, choose the database template location. The location can be Software Library or Oracle Home. The template selected must be compatible with the selected Oracle Home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle Home. The default location is `ORACLE_HOME/assistants/dbca/templates`.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

5. In the Identification and Placement page, select the type of Oracle RAC database, whether Policy Managed or Admin Managed. Also, specify **Global Database Name** and **SID** prefix.

For admin-managed database, select nodes on which you want to create the cluster database. You must specify the node selected as the reference node in the Database Version and Type page.

For policy-managed database, select the server pools to be used for creating the database, from the list of existing server pools, or choose to create a new server pool. Policy-managed databases can be created for database versions 11.2 and higher. For database versions lower than 11.2, you will need to select nodes to create the Oracle RAC database.

In the Database Consolidation section, select **Create As Container Database** if you want to create a container database. By default, an empty container database is created. If you want to add one or more pluggable databases to that container database, then select **Create a Container Database with one or more PDBs**, and set the number of PDBs.

If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all the cloned PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name accountsPDB, then the PDBs are created with the names accountsPDB1, accountsPDB2, accountsPDB3, accountsPDB4, and accountsPDB5.

Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP.

For database version 12.1 or higher, for Microsoft Windows operating systems, the database services will be configured for the Microsoft Windows user specified during Oracle home installation. This user will own all services run by Oracle software. In the Oracle Home Windows User Credentials section, specify the host credentials for the Microsoft Windows user account to configure database services. Select existing named credentials or specify new credentials. To specify new credentials, provide the user name and password. You can also save these credentials and set them as preferred credentials.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

6. In the Storage Locations page, select the storage type, whether File System or Automatic Storage Management (ASM).

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored. These locations must be on shared storage such as cluster file system location or ASM diskgroups.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

In the Archive Log Settings section, select **Enable Archiving** to enable archive logging. In the Specify Archive Log Locations, you can specify up to nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

7. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

8. In the Additional Configuration Options page, select custom scripts from the Software Library. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click the **Lock** icon to lock the field. Click **Next**.

9. In the Schedule page, specify a Deployment Instance name and a schedule for the deployment. If you want to run the procedure immediately, then retain the default selection, that is Immediately. If you want to run the procedure later, then select Later and provide time zone, start date, and start time details. Click **Next**.
10. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment. Click **Analyze** to check for prerequisites and to ensure that all the necessary requirements for provisioning are met.
11. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Creating Oracle Real Application Clusters One Node Database

This section provides information about creating Oracle Real Application Clusters One Node Database (also called as Oracle RAC One Node Database).

 Note:

You can also use the information provided in this section to create a Oracle Real Application Clusters One Node container database.

You can create a container database on a host only if Oracle Database 12c Release 1 (12.1), or higher, is installed on the host. For more information on container databases, see *Oracle Database Administrator's Guide*.

This section covers the following:

- [Prerequisites for Creating an Oracle RAC One Node Database](#)
- [Procedure for Creating an Oracle Real Application Clusters One Node Database](#)

Prerequisites for Creating an Oracle RAC One Node Database

To create an Oracle RAC One databases using Cloud Control, ensure that you meet the following prerequisites:

1. Ensure that you meet the infrastructure requirements explained in [Setting Up Your Infrastructure](#).
2. Ensure that you have created and stored the database template in the Software Library or Oracle Home. For information about creating database templates, see [Creating Database Templates](#).
3. Oracle Home for the database you want to create must be installed and you need to have credentials of the owner of the Oracle Home. If the Create Database wizard is launched from the Provision Database deployment procedure wizards, the Oracle Home need not be installed earlier. In such cases, the validations for Oracle Home will be skipped during the procedure interview and will be performed during execution of the deployment procedure.
4. The database plug-in that supports the corresponding database version should be deployed on OMS and Agent. For information about deploying plug-ins, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
5. Ensure that you have sufficient space to create the database, and that you have write permissions to the recovery file location.
6. If you are using a template from the Software Library for database creation, you must have Write permission to the Staging Location.
7. If you are creating Oracle Real Application Clusters database, you must have Grid Infrastructure installed and configured. If the Create Database wizard is launched from the Provision Database deployment procedure wizards, Grid Infrastructure need not be installed and configured. In such cases, the validations for Grid Infrastructure will be skipped during the procedure interview and will be performed during execution of the deployment procedure.

8. If you are using Automatic Storage Management (ASM) as storage, ASM instances and diskgroups must be configured prior to creating database.
9. The Cloud Control user creating the database template must have CONNECT_ANY_TARGET privilege in Cloud Control.

Procedure for Creating an Oracle Real Application Clusters One Node Database

To create an Oracle Real Application Clusters One Node database, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select the Create Oracle Database Deployment Procedure and click **Launch**. The Create Oracle Database wizard is launched.
3. In the Database Version and Type page, select the database **Version** and select **Oracle RAC One Node Database**.

In the Cluster section, select the cluster and Oracle Home. Select a reference host to perform validations to use as reference to create database on the cluster.

Select **Cluster Credentials** or add new. Click the plus icon to add new credentials and specify **User Name**, **Password**, and **Run Privileges** and save the credentials.

Click **Next**.

4. In the Database Template page, choose the database template location. The location can be Software Library or Oracle Home. The template selected must be compatible with the selected Oracle Home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle Home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

5. In the Identification and Placement page, select nodes on which you want to create the cluster database. Specify **Global Database Name** and **SID** prefix. Select the type of Oracle RAC database, whether Policy Managed or Admin Managed. Specify the **Service Name**.

In the Database Consolidation section, select **Create As Container Database** if you want to create a container database. By default, an empty container database is created. If you want to add one or more pluggable databases to that container database, then select **Create a Container Database with one or more PDBs**, and set the number of PDBs.

If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all the cloned PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name `accountsPDB`, then the PDBs are created with the names `accountsPDB1`, `accountsPDB2`, `accountsPDB3`, `accountsPDB4`, and `accountsPDB5`.

Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP database accounts.

For database version 12.1 or higher, for Microsoft Windows operating systems, the database services will be configured for the Microsoft Windows user specified during Oracle home installation. This user will own all services run by Oracle software. In the Oracle Home Windows User Credentials section, specify the host credentials for the Microsoft Windows user account to configure database services. Select existing named credentials or specify new credentials. To specify new credentials, provide the user name and password. You can also save these credentials and set them as preferred credentials.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

6. In the Storage Locations page, select the storage type, whether File System or Automatic Storage Management (ASM).

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify up to five locations.

In the Recovery Files Location section, select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

In the Archive Log Settings section, select **Enable Archiving** to enable archive logging. In the Specify Archive Log Locations, you can specify up to nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

7. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

8. In the Additional Configuration Options page, select custom scripts from the Software Library. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options. Click on the Lock icon to lock the field. Click **Next**.
9. In the Schedule page, specify a Deployment Instance name and a schedule for the deployment. If you want to run the procedure immediately, then retain the default selection, that is Immediately. If you want to run the procedure later, then select Later and provide time zone, start date, and start time details. Click **Next**.
10. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment. Click **Analyze** to check for prerequisites and to ensure that all the necessary requirements for provisioning are met.
11. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Provisioning Databases

This section explains how you can provision or mass-deploy single-instance databases in an unattended, repeatable, and reliable manner, using Oracle Enterprise Manager Cloud Control (Cloud Control). It covers the following:

- [Getting Started with Provisioning Oracle Databases](#)
- [Provisioning and Creating Oracle Databases](#)
- [Provisioning Oracle Databases with Oracle Automatic Storage Management](#)
- [Provisioning Oracle Database Software Only](#)

Getting Started with Provisioning Oracle Databases

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Oracle Databases. Consider this section to be a docu-

mentation map to understand the sequence of actions you must perform to successfully provision single-instance databases. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 5-2 Getting Started with Provisioning Oracle Databases

Step	Description	Reference Links
Step 1	Understanding Oracle Database Topology Understand the Database Provisioning feature that is offered by Cloud Control for provisioning single-instance databases.	To learn about Oracle Database topology, see Oracle Database Topology .
Step 2	Selecting the Use Case This chapter covers a few use cases for provisioning Oracle Database. Select the use case that best matches your requirements.	- To learn about provisioning and configuring Oracle Database, see Provisioning and Creating Oracle Databases. - To learn about provisioning Oracle Database with Automatic Storage Management, see Provisioning Oracle Databases with Oracle Automatic Storage Management. - To learn about provisioning Oracle Database software, see Provisioning Oracle Database Software Only.
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, and setting up of Oracle Software Library.	- To learn about the prerequisites in provisioning and configuring Oracle Database, see Prerequisites for Provisioning Databases. - To learn about the prerequisites in provisioning Oracle Database with Automatic Storage Management, see Prerequisites for Provisioning Oracle Databases with Oracle Automatic Storage Management. - To learn about the prerequisites in provisioning Oracle Database software, see Prerequisites for Provisioning Oracle Database Software Only.
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle Database.	- To provision and configure Oracle Database, see Procedure for Provisioning Databases. - To provision Oracle Database with Automatic Storage Management, see Procedure for Provisioning Databases. - To provision Oracle Database software, see Procedure for Provisioning Oracle Database Software Only.

Provisioning and Creating Oracle Databases

This section describes how you can provision and create Oracle Databases.

In particular, this section covers the following:

- [Prerequisites for Provisioning Databases](#)

- [Procedure for Provisioning Databases](#)

Prerequisites for Provisioning Databases

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Databases

To run the deployment procedure for provisioning a database, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the **Provision Oracle Database** Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select either **Deploy and configure Grid Infrastructure**, **Deploy the software only** or **Do not provision**.
- To deploy Database software select either **Deploy software only** or **Deploy and create a new database**, which creates a new database and configures it after installing the standalone Oracle Database.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, the various configuration options are displayed. Provide values for the Setup Hosts, Deploy Software, Configure Grid Infrastructure, and Create Database tasks.
5. Click on the Setup Hosts link.
6. In the Specify OS Users page, specify the operating system user for the Oracle Home for the database.

Note:

To use no root credentials, see [Using Non-Root Credentials for Provisioning Oracle Databases](#).

For Oracle Home User for the database, select the Normal User and Privileged User to be added to the OS group.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**.

7. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.

- Inventory Group (OINSTALL)
- Database Administrator (OSDBA)
- 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://docs.oracle.com/database/122/index.html>

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://docs.oracle.com/database/122/index.html>

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a completed status.

8. Click on the Deploy Software link.
9. In the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database.

In the Source section, select the Software Library location for **Oracle Database** binaries.

 Note:

For Windows operating systems, if the Oracle Database component selected is of version 12.1 or higher, you can install all services as a named Oracle service user with limited privileges. This will enhance security for database services.

In the Windows Security option section, you can configure the option for an existing user or add a user and specify the User Name and Password. Select **Decline Security** option if you want all the services to be installed and configured as an administrative user.

In the Destination location, specify the following:

- **Oracle Base for Database**, 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 differ-

ent from the Oracle home directory where the database software will be installed.

- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

In the Additional Parameters section, specify the **Working 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. For **Database Installer Parameters**, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle database. 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.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a completed status.

10. Click on the Create Databases link.

11. In the Database Template page, choose the database template location. The location can be Software Library or Oracle Home. The template selected must be compatible with the selected Oracle Home version.

If you choose **Select Template from Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must exist on all hosts where you want to create the database.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you choose **Select Template from Oracle Home**, select the template from the Oracle home. The default location is `ORACLE_HOME/assistants/dbca/templates`.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

12. In the Identification and Placement page, specify database configuration details. Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP database accounts. You can choose to use the same or different administrative passwords for these accounts.

 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 Cloud 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.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

- 13.** In the Storage Locations page, select the storage type, whether File System or Automatic Storage Management (ASM).

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 or /u01/product/db/oradata.

If you want to use ASM, then select **Automatic Storage Management (ASM)**, 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.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database

files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify up to nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

15. In the Additional Configuration Options, all the available listeners running from the Oracle Home are listed. You can either select a listener or create a new one. You can select multiple listeners to register with the database. To create a new listener, specify the **Listener Name** and **Port**. Select database schemas and specify custom scripts, if any. Select custom scripts from the host where you are creating the database or from Software Library. If you have selected multiple hosts, you can specify scripts only from Software Library.

If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

16. Review the details you have provided for creating the database and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a completed status.
17. Click the Compliance Standards link.
18. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.

19. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

20. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
21. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
22. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
23. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly created databases appear as Cloud Control targets.

Provisioning Oracle Databases with Oracle Automatic Storage Management

This section describes how you can provision single-instance databases with Oracle Automatic Storage Management (Oracle ASM).

In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle Databases with Oracle Automatic Storage Management](#)
- [Procedure for Provisioning Databases](#)

Prerequisites for Provisioning Oracle Databases with Oracle Automatic Storage Management

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Databases

To provision a single-instance database with Oracle Automatic Storage Management (Oracle ASM), follow these steps:

1. Log in as a designer, and from the Enterprise menu, select **Provisioning and Patching**, then select **Database Provisioning**.

2. In the Database Procedures page, select the **Provision Oracle Database** Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select either **Deploy and configure Grid Infrastructure**, **Deploy the software only** or **Do not provision**.
- To deploy Database software select either **Deploy software only** or **Deploy and create a new database**, which creates a new database and configures it after installing the standalone Oracle Database.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system user for the Oracle Home for the database.

 Note:

To use no root credentials, refer to [Using Non-Root Credentials for Provisioning Oracle Databases](#).

For Oracle Home User for the database, select the Normal User and Privileged User to be added to the OS group.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - Database Administrator (OSDBA)
 - 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://docs.oracle.com/database/122/index.html>

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://docs.oracle.com/database/122/index.html>

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a completed status.

7. Click on the Deploy Software link.
8. In the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database.

In the Source section, select the Software Library location for **Oracle Database** binaries.

In the Destination location, specify the following:

- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

In the Additional Parameters section, specify the **Working 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. For **Database Installer Parameters**, 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.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a completed status.

9. Click on the Create Databases link.
10. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you choose **Select Template from Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must exist on all hosts where you want to create the database.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you choose **Select Template from Oracle Home**, select the template from the Oracle home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

11. In the Identification and Placement page, specify database configuration details. Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP database accounts. You can choose to use the same or different administrative passwords for these accounts.

 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 Cloud 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.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

12. In the Storage Locations page, select the storage type as **Automatic Storage Management (ASM)** 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.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify upto nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

13. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Additional Configuration Options, all the available listeners running from the Oracle Home are listed. You can either select a listener or create a new one. You can select multiple listeners to register with the database. To create a new listener, specify the **Listener Name** and **Port**. Select database schemas and specify custom scripts, if any. Select custom scripts from the host where you are creating the database or from Software Library. If you have selected multiple hosts, you can specify scripts only from Software Library.

If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

15. Review the details you have provided for creating the database and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a completed status.
16. Click the Compliance Standards link.
17. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.
18. In the Configure page, click **Next**.
19. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.
20. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only pre-requisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.
Click **Next**.
21. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
22. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
23. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
24. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly created databases appear as Cloud Control targets.

Provisioning Oracle Database Software Only

This section provides information about provisioning single-instance database software.

In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle Database Software Only](#)
- [Procedure for Provisioning Oracle Database Software Only](#)

Prerequisites for Provisioning Oracle Database Software Only

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Oracle Database Software Only

Follow these steps:

1. Log in as a designer, and from the Enterprise menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle Database Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select **Deploy Database software** to provision single-instance databases.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system user for the Oracle Home for the database.



Note:

To use no root credentials, refer to [Using Non-Root Credentials for Provisioning Oracle Databases](#).

For Oracle Home User for the database, select the Normal User and Privileged User to be added to the OS group.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)

If these groups 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://docs.oracle.com/database/122/index.html>

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://docs.oracle.com/database/122/index.html>

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a completed status.

7. Click on the Deploy Software link.
8. In the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database.

In the Source section, select the Software Library location for **Oracle Database** binaries.

 Note:

For Windows operating systems, if the Oracle Database component selected is of version 12.1 or higher, you can install all services as a named Oracle service user with limited privileges. This will enhance security for database services.

In the Windows Security option section, you can configure the option for an existing user or add a user and specify the User Name and Password. Select **Decline Security** option if you want all the services to be installed and configured as an administrative user.

In the Destination location, specify the following:

- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

In the Additional Parameters section, specify the **Working 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. For **Database Installer Parameters**, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle database. 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.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a completed status.

9. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.
Click **Next**.
10. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
11. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
12. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
13. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly created databases appear as Cloud Control targets.

Provisioning Oracle Database Software from a Shared NFS location

To provision Oracle Database Software from a shared NFS location, perform the following steps:

Before you proceed with the steps below ensure the Preferred credentials (both Named and Privileged) and Privilege delegation are configured for all the destination hosts.

1. From the **Enterprise** menu, select **Provisioning and Patching**, and then select **Database Provisioning**.
2. Select **Provision Oracle Database** deployment procedure and click **Launch**.
3. In the Provision Oracle Database: Select Hosts screen, provide the details as appropriate and add destination hosts.



Note:

Do not select Provisioning Profile in this step.

4. Click **Next**.
5. Click on **Setup Hosts** and provide the required details.
6. Click on **Deploy Software**.
7. In the Source section, select the appropriate software entity that was created, by clicking on the search icon for **Oracle Database**.

8. Provide other details as appropriate, and click **Next**.
9. On the Provision Oracle Database: Configure screen, click **Next**.
10. Complete the Provision Oracle Database wizard flow by providing appropriate details in the remaining steps and click **Submit** in the Review screen.

NFS share containing the DB Gold Image will be mounted on all the destination hosts selected for Oracle Home provisioning.

Using Non-Root Credentials for Provisioning Oracle Databases

Non-root credentials are supported for provisioning Oracle databases. To use this feature, perform the following steps:

1. On the Specify OS users page, select **Override Preferred Credentials**. On the Specify OS users dialogue box that appears, create the normal name credential, and then set Run Privilege to **None**. Click **OK**.
2. Select the new normal name credential for both Normal user and Privileged user.
3. Click **Submit**.

When the database provisioning process reaches the step which requires root credentials, the process will stop. You will need to run the command line manually. To do this, set the environment to \$AGENT_HOME, and then run the command line copy from the Instructions field for the following two steps:

- Execute fixups manually
 - Execute Root scripts manually
4. After the command line is run manually using the root user for both the steps, click **Confirm**. The database provisioning process then continues till it completes.

6

Provisioning Oracle Grid Infrastructure for Oracle Databases

This chapter explains how you can mass-deploy Oracle Grid Infrastructure for Oracle databases (also called as single-instance databases) in an unattended, repeatable, and reliable manner, using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with Provisioning Oracle Grid Infrastructure for Oracle Databases](#)
- [Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management](#)
- [Provisioning Oracle Grid Infrastructure and Oracle Database Software Only](#)

Getting Started with Provisioning Oracle Grid Infrastructure for Oracle Databases

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Oracle Grid Infrastructure for single-instance databases. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision Oracle Grid Infrastructure with single-instance databases. 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	Selecting the Use Case This chapter covers a few use cases for provisioning Oracle Grid Infrastructure. Select the use case that best matches your requirements.	• To learn about provisioning Grid Infrastructure and Oracle databases and configuring database with Oracle ASM, see Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management. • To learn about provisioning Grid Infrastructure and Oracle database software only, see Provisioning Oracle Grid Infrastructure and Oracle Database Software Only.

Table 6-1 (Cont.) Getting Started with Provisioning Oracle Grid Infrastructure

Step	Description	Reference Links
Step 2	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, and setting up of Oracle Software Library.	- To learn about the prerequisites for provisioning Grid Infrastructure and Oracle databases and configuring database with Oracle ASM, see Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM. - To learn about the prerequisites for provisioning Grid Infrastructure and single-instance database software only, see Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only.
Step 3	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle Grid Infrastructure.	- To provision Grid Infrastructure and Oracle databases and configuring database with Oracle ASM, follow the steps explained in Procedure for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM. - To provision Grid Infrastructure and Oracle database software only, follow the steps explained in Procedure for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only.

Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management

This section describes how you can provision Oracle Grid Infrastructure and single-instance databases with Oracle Automatic Storage Management (Oracle ASM).

In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM](#)
- [Procedure for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM](#)

Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM

To provision Oracle grid infrastructure and Oracle databases with Oracle ASM, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle Database Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select **Deploy software and configure**.
- To deploy Database software select either **Deploy and create a new database**, which creates a new database and configures it after installing the standalone Oracle Database, or **Deploy software only**, which installs a new database software.

Click the **Lock** icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system user for the Oracle Home for the database and Grid Infrastructure. Specify the Normal User and Privileged User to be added to the OS groups.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)
 - ASM Instance Administrator (OSASM)
 - ASM Database Administrator (ASMDBA)
 - ASM Instance Operator (ASMOPER)

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>

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a completed status.

7. Click the **Deploy Software** link.
8. In the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.

In the Destination location, specify the following:

- **Oracle Base for Grid Infrastructure**, a location on the destination host where the diagnostic and administrative logs, and other logs associated with the Grid Infrastructure can be stored.
- **Grid Infrastructure Home**, 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.
- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

For Grid Infrastructure, Oracle Base is /u01/app/user and Oracle Home is %ORACLE_BASE%/sihome. You can use %ORACLE_BASE% and %GI_ORACLE_BASE% to specify the relative paths which will be interpolated to their respective values.

In the Additional Parameters section, specify the **Working 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. For **Installer Parameters**, 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.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a completed status.

9. Click the **Configure Grid Infrastructure** link.
10. In the Configure GI page, of ASM Storage section, click **Add** to add an ASM Disk Group. In the **Add/Edit Disk Group** dialog box, specify the **Disk Group Name**,

Disk List, and specify the redundancy as **Normal**, **High**, **Flex**, **External**, or **Extended**. Click **OK**.

For ASM 11.2 and higher, specify the **Disk Group Name** for storing the parameter file. In GIMR Disk Group section, select the Disk Group to store the GIMR files. From 19c onwards, you can choose to setup GIMR by selecting the **Configure GIMR** check box. If it is not selected, GIMR will not be configured. Specify the **ASM Password** for ASMSNMP and SYS users. Specify the Listener Port for registering the ASM instances.

As a designer, you can click on the Lock icon to lock these fields. These fields will then not be available for editing in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the storage options for the Grid Infrastructure and database, the Configure Grid Infrastructure task will have a completed status.

11. Click the **Create Databases** link.

12. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you choose **Select Template from Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must exist on all hosts where you want to create the database.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you choose **Select Template from Oracle Home**, select the template from the Oracle home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

13. In the Identification and Placement page, specify database configuration details. Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP database accounts. You can choose to use the same or different administrative passwords for these accounts.

 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 Cloud 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.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Storage Locations page, select the storage type as **Automatic Storage Management (ASM)**, 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.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify upto nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

15. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

16. In the Additional Configuration Options, all the available listeners running from the Oracle Home are listed. You can either select a listener or create a new one. You can select multiple listeners to register with the database. To create a new listener, specify the **Listener Name** and **Port**. Select database schemas and specify custom scripts, if any. Select custom scripts from the host where you are creating the database or from Software Library. If you have selected multiple hosts, you can specify scripts only from Software Library.

If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

17. Review the details you have provided for creating the database and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a completed status.
18. Click the **Compliance Standards** link.
19. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.
20. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.
21. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

22. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
23. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
24. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
25. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly provisioned databases appear as Cloud Control targets.

**Note:**

If the Deployment Procedure fails, then review log files described in [Reviewing Log Files](#).

Provisioning Oracle Grid Infrastructure and Oracle Database Software Only

This section describes how you can provision Oracle Grid Infrastructure and Oracle Database software.

In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only](#)
- [Procedure for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only](#)

Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Oracle Grid Infrastructure and Oracle Database Software Only

To provision Oracle Grid Infrastructure and Oracle Database software, follow these steps:

1. Log in as a designer, and from the Enterprise menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Oracle Database Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select **Deploy software and configure** to provision single-instance high-availability databases.
- To deploy Database software select **Deploy software only**, which installs a new database software.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system user for the Oracle Home for the database and Grid Infrastructure. Specify the Normal User and Privileged User to be added to the OS groups.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)
 - ASM Instance Administrator (OSASM)
 - ASM Database Administrator (ASMDBA)
 - ASM Instance Operator (ASMOPER)

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>

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a completed status.

7. Click on the Deploy Software link.
8. In the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.

In the Destination location, specify the following:

- **Oracle Base for Grid Infrastructure**, a location on the destination host where the diagnostic and administrative logs, and other logs associated with the Grid Infrastructure can be stored.
- **Grid Infrastructure Home**, 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.
- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database.

For Grid Infrastructure, Oracle Base is /u01/app/user and Oracle Home is %ORACLE_BASE%/sihome. You can use %ORACLE_BASE% and %GI_ORACLE_BASE% to specify the relative paths which will be interpolated to their respective values.

In the Additional Parameters section, specify the **Working 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. For **Installer Parameters**, 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.

Click on the Lock icon against the fields that you do not want to be edited in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a completed status.

9. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only pre-

requisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

10. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
11. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
12. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

 Note:

If the Deployment Procedure fails, then review log files described in [Reviewing Log Files](#).

Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases

This chapter explains how you can mass-deploy Oracle Grid Infrastructure and Oracle Real Application Clusters (Oracle RAC) for clustered environments in an unattended, repeatable, and reliable manner. In particular, this chapter covers the following:

- [Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases](#)
- [Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management](#)
- [Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster](#)
- [Provisioning Oracle Real Application Clusters Database with File System on a New Cluster](#)

Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases

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. 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 Databases

Step	Description	Reference Links
Step 1	Understanding Oracle RAC Topology Understand the Oracle Real Application Clusters Database topology provisioned by Cloud Control.	To learn about the topology, see Oracle Real Application Clusters Database Topology and Oracle Flex Clusters and Flex ASM .


Table 7-1 (Cont.) Getting Started with Provisioning Oracle Grid Infrastructure and Oracle RAC Databases

Step	Description	Reference Links
Step 2	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 requirements.	- To learn about provisioning Oracle Grid Infrastructure and Oracle RAC database and configuring ASM and Database, see Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management. - To learn about provisioning Oracle RAC database with File System on an existing cluster, see Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster. - To learn about provisioning Oracle RAC database with File System on a new cluster, see Provisioning Oracle Real Application Clusters Database with File System on a New Cluster.
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, and setting up of Oracle Software Library.	- To learn about the prerequisites for provisioning Oracle Grid Infrastructure and Oracle RAC database and configuring ASM and Database, see Prerequisites for Provisioning Grid Infrastructure with Oracle RAC Database. - To learn about the prerequisites for provisioning Oracle RAC database with File System on an existing cluster, see Prerequisites for Provisioning Oracle RAC Database with File System on an Existing Cluster. - To learn about the prerequisites for provisioning Oracle RAC database with File System on a new cluster, see Prerequisites for Provisioning Oracle RAC Database with File System on a New Cluster.
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle Grid Infrastructure and Oracle RAC.	- To provision Oracle Grid Infrastructure and Oracle RAC database and configure ASM and Database, follow the steps explained in Procedure for Provisioning Grid Infrastructure with Oracle RAC Database. - To provision Oracle RAC database with File System on an existing cluster, follow the steps explained in Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster. - To provision Oracle RAC database with File System on a new cluster, follow the steps explained in Procedure for Provisioning Oracle RAC Database with File System on a New Cluster.

Oracle Real Application Clusters Database Topology

Oracle Enterprise Manager Cloud Control enables standardized gold image-based deployments of Oracle RAC databases with provisioning profiles, input lock down in designer role, and associating compliance standards with databases. [Figure 7-1](#) shows a typical Oracle RAC database topology that you can provision using Cloud Control.

Figure 7-1 Oracle RAC Database Topology



The topology shows a N-node setup using Grid Infrastructure, clustered ASM, and policy-managed Oracle RAC database. An ASM disk array is shared thorough the cluster setup. The Grid Infrastructure uses an ASM diskgroup named QUORUM for Oracle Cluster Registry (OCR) and Voting Disk (Heartbeat). The Oracle RAC database uses another diskgroup named DATA. This stores database datafiles. The nodes are multi-homed such that a high speed internal network between nodes facilitates cluster operation, and a public network is used for external connectivity. The networks are public, private, and storage network between nodes and the ASM disk array.

Oracle Flex Clusters and Flex ASM

Oracle Flex Clusters

Oracle Grid Infrastructure installed in an Oracle Flex Cluster configuration is a scalable, dynamic, robust network of nodes. Oracle Flex Clusters provide a platform for a variety of applications, including Oracle Real Application Clusters (Oracle RAC) databases with large numbers of nodes. Oracle Flex Clusters also provide a platform for other service deployments that require coordination and automation for high availability.

All nodes in an Oracle Flex Cluster belong to a single Oracle Grid Infrastructure cluster. This architecture centralizes policy decisions for deployment of resources based on application needs, to account for various service levels, loads, failure responses, and recovery.

Oracle Flex Clusters contain two types of nodes arranged in a hub and spoke architecture: Hub Nodes and Leaf Nodes. The number of Hub Nodes in an Oracle Flex Cluster can be as many as 64. The number of Leaf Nodes can be many more. Hub Nodes and Leaf Nodes can host different types of applications.

Hub Nodes are similar to Oracle Grid Infrastructure nodes in an Oracle Clusterware standard Cluster configuration: they are tightly connected, and have direct access to shared storage. In an Oracle Flex Cluster configuration, shared storage can be provisioned to Leaf Nodes independent of the Oracle Grid Infrastructure.

Leaf Nodes are different from standard Oracle Grid Infrastructure nodes, in that they do not require direct access to shared storage, but instead request data through Hub Nodes. Hub Nodes can run in an Oracle Flex Cluster configuration without having any Leaf Nodes as cluster member nodes, but Leaf Nodes must be members of a cluster that includes at least one Hub Node.

Oracle Flex ASM

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. With this deployment, larger clusters of Oracle ASM instances can support more database clients while reducing the Oracle ASM footprint for the overall system.

For more information on Oracle Flex Clusters and Flex ASM see the links below:

- [Oracle Database documentation](#)
- <http://www.oracle.com/technetwork/articles/database/flexasm-flexcluster-benefits-odb12c-2177371.html>

Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management

This section describes how you can provision Grid Infrastructure with Oracle Real Application Clusters (Oracle RAC) Database and configure Database with Oracle Automatic Storage Management (Oracle ASM).

In particular, this section covers the following:

- [Prerequisites for Provisioning Grid Infrastructure with Oracle RAC Database](#)
- [Procedure for Provisioning Grid Infrastructure with Oracle RAC Database](#)

Prerequisites for Provisioning Grid Infrastructure with Oracle RAC Database

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Grid Infrastructure with Oracle RAC Database

To provision the grid infrastructure with Oracle RAC database, and to configure the database with Oracle ASM, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle RAC Database Deployment Procedure and click **Launch**. The Oracle RAC Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy and configure Grid Infrastructure, select **Deploy software only** to provision Oracle RAC databases.
- To deploy database software select **Deploy and create a RAC database** which creates a new database software and configures it after installing the Oracle RAC database.

Click the **Lock** icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link, and then click **Next**.
5. In the Specify OS Users page, specify the operating system users and groups required to provision the database.

 Note:

Choose the credentials to be used to provision the database. To avoid using credentials, refer [Using No Root Credentials for Provisioning Oracle Real Application Clusters \(Oracle RAC\) Databases](#).

For Database User and ASM User, select the Normal User and Privileged User to perform the tasks as the OS user.

Click **Next**.

6. For database versions 11.2 and above, in the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - ASM Database Administrator (ASMDBA)
 - ASM Instance Operator (ASMOPER)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)
 - ASM Instance Administrator (OSASM)

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a Configured status.

7. Click on the Deploy Software link.
8. In the Select 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. As a designer, you can click the **Lock** icon to lock these fields. These fields will then not be available for editing in the operator role.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.

**Note:**

For Windows operating systems, if the Oracle Grid Infrastructure or Oracle Database component selected is of version 12.1 or higher, you can install all services as a named Oracle service user with limited privileges. This will enhance security for database services.

In the Windows Security option section, you can configure the option for an existing user and specify the User Name and Password. Select **Decline Security** option if you want all the services to be installed and configured as an administrative user.

In the Destination location, specify the following:

- **Oracle Base for Grid Infrastructure**, a location on the destination host where the diagnostic and administrative logs, and other logs associated with the Grid Infrastructure can be stored.
- **Grid Infrastructure Home**, 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. Optionally, select **Shared Grid Infrastructure home** to enable Grid Infrastructure Oracle Home on shared locations. Ensure that the directory path you provide meets the requirements described in [Requirements for Grid Infrastructure Software Location Path](#).
- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database. Optionally, select **Shared Database Oracle home** to enable Database Oracle Home on shared locations.

For Grid Infrastructure, Oracle Base is /u01/app/user and Oracle Home is %GI_ORACLE_BASE%/.../grid. Oracle Base for Grid Infrastructure is referred to as %GI_ORACLE_BASE% and that for Database is referred to as %ORACLE_BASE%. You can use %ORACLE_BASE% and %GI_ORACLE_BASE% to specify the relative paths which will be interpolated to their respective values.

In the Additional Parameters section, specify the **Working 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. For **Installer Parameters**, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle Grid Infrastructure and database. 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.

If required, you can also specify OCFS devices in the Installer Parameters field in the following format, separating devices with commas:

Device Number:Partition Number: Drive letter: [DATA | SOFTWARE]

For example, `–ocfs_devices=1:1:E:DATA,1:2:F:DATA`

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a Configured status.

9. Click on the Configure Grid Infrastructure link.
10. In the Select Storage page, select the storage type for database as **Automatic Storage Management** or **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.

You can choose either Automatic Storage Management or File System for Oracle RAC Database. The storage type for Grid Infrastructure is set to Automatic Storage Management with Flex redundancy.

As a designer, you can click the **Lock** icon to lock these fields. These fields will then not be available for editing in the operator role.

Click **Next**.

11. In the Configure Cluster page, in the Basic Settings section, specify the **Cluster Name**, **SCAN Name**, and **SCAN Port**. 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.

Provision Oracle RAC Database: Configure Cluster

Basic Settings

Cluster Name	crs1	Check Availability
SCAN Name	crs1-scan	Check Availability
SCAN Port	1521	Check Availability

GNS Settings

☒ Configure Grid Naming Service (GNS)

☐ Configure nodes Virtual IPs as assigned by DHCP

Check Availability

GNS Sub Domain

GNS VIP Address

Hostname	Role	Virtual Hostname
racdbnode01...	HUB	racdbnode01-1521

Network

Add Delete... >>		
Interface Name	Subnet	Usage
eth0	10.10.10.0	Public

In the GNS Settings section, select **Configure Grid Naming Service** and specify the **GNS Sub Domain** and **GNS VIP Address** if you want virtual host names outside the cluster to have dynamically assigned names.

In the Network 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, or ASM, or ASM & Private, or Do Not Use.

Click **Add** to add an interface and specify the **Interface Name** and **Interface Subnet** and click **OK**. Select the Usage as **Public**, **Private**, **ASM**, **ASM & Private**, or **Do Not Use** if you do not want to use the interface.

If you have chosen storage type as Automatic Storage Management for either or both Grid Infrastructure and Oracle RAC Database, in the Automatic Storage Management (ASM) section, select from the ASM Disk Groups that have been discov-

ered by Cloud Control and are displayed in the table. Click **Add** to add an ASM Disk Group. In the Add/Edit Disk Group dialog box, specify the **Disk Group Name**, **Disk List**, and specify the redundancy as **Normal**, **High**, **External**, **Flex**, or **Extended**. If required, set the Advanced Disk Group Attributes and click **OK**.

Select the OCR/Voting Disk to store cluster registry and voting disk files, and specify the ASM credentials for ASMSNMP and SYS users.

Configure Automatic Storage Management (ASM) ⓘ

+ Add Delete... >>

Disk Group Name	Disk List	Redundancy
No ASM Disk Groups specified. To add ASM Disk Group(s) use 'Add...' option.		

OCR/Voting Disk

Select the Disk Group to store the cluster registry and voting disk files.

OCR/Voting Disk

Specify Password

Provide password for ASMSNMP and SYS users

ASM Password

Confirm Password

Click **Next**. You will come back to the Configure page. If you have configured the storage options for the Grid Infrastructure and database, the Configure Grid Infrastructure task will have a completed status.

12. Click on the Create Databases link.

13. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Identification and Placement page, select the Oracle RAC database configuration type, whether Policy Managed or Admin Managed.

For admin-managed database, select nodes on which you want to create the cluster database. You must specify the node selected as the reference node in the Database Version and Type page.

For policy-managed database, select the server pools to be used for creating the database, from the list of existing server pools, or choose to create a new server pool. Policy-managed databases can be created for database versions 11.2 and

higher. For database versions lower than 11.2, you will need to select nodes to create the Oracle RAC database.

You can also choose to create a CDB and populate it with PDBs.

Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP. You can choose to specify the same or different passwords for each of these user accounts.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

15. In the Storage Locations page, select the same storage type you specified for Oracle Database in the Select Storage page.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored. These locations must be on shared storage such as cluster file system location or ASM diskgroups.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify upto nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

16. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

17. In the Additional Configuration Options page, select custom scripts from the Software Library or your local disk. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click the **Lock** icon to lock the field. Click **Next**.

18. Review the information you have provided and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a Configured status. Click **Next**.

19. Click the Compliance Standards link.

20. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.

21. In the Configure page, click **Next**.

22. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.

23. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

24. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.

25. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.

26. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

27. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly provisioned databases appear as Cloud Control targets.

 Note:

If the Deployment Procedure fails, then review log files described in [Reviewing Log Files](#).

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. These requirements are suggested as the GI configuration changes the ownership of the GI home and all parent directories to be owned by root.

- 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.

Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster

This section describes how you can provision Oracle Real Application Clusters (Oracle RAC) database with a file system on an existing cluster. In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle RAC Database with File System on an Existing Cluster](#)
- [Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster](#)

Prerequisites for Provisioning Oracle RAC Database with File System on an Existing Cluster

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster

To provision Oracle RAC databases with file system on an existing cluster, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.

2. In the Database Procedures page, select the Provision Oracle RAC Database Deployment Procedure and click **Launch**. The Oracle RAC Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select **Use existing Grid Infrastructure** and select the **Cluster Name**.
- To deploy database software select **Deploy software and create a RAC database** which creates a new database and configures it after installing the Oracle RAC database.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system users and groups required to provision the database.

 Note:

To use no root credentials, see [Using No Root Credentials for Provisioning Oracle Real Application Clusters \(Oracle RAC\) Databases](#).

For Database User, select the Normal User and Privileged User to be added to the OS group.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.
 - Inventory Group (OINSTALL)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a Configured status.

7. Click on the Deploy Software link.
8. In the Select Software Locations page, specify the locations where the software binaries of Oracle RAC database can be placed, that is, the \$ORACLE_HOME location. As a designer, you can click on the Lock icon to lock these fields. These fields will then not be available for editing in the operator role.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.



Note:

For Windows operating systems, if the Oracle Grid Infrastructure or Oracle Database component selected is of version 12.1 or higher, you can install all services as a named Oracle service user with limited privileges. This will enhance security for database services.

In the Windows Security option section, you can configure the option for an existing user and specify the User Name and Password. Select **Decline Security** option if you want all the services to be installed and configured as an administrative user.

In the Destination location, specify the following:

- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database. Select **Shared Database Oracle home** to enable Database Oracle Home on shared locations.

In the Additional Parameters section, specify the **Working 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. For **Installer Parameters**, specify any additional Oracle Universal Installer (OUI) parameters you want to run while provisioning Oracle RAC database. 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.

You can also specify OCFS devices in the Installer Parameters field in the following format, separating devices with commas:

Device Number:Partition Number: Drive letter: [DATA | SOFTWARE]

For example, `-ocfs_devices=1:1:E:DATA,1:2:F:DATA`

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a Configured status.

9. Click on the Create Databases link.
10. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

11. In the Identification and Placement page, select the Oracle RAC database configuration type, whether Policy Managed or Admin Managed.

For Admin-managed database, select nodes on which you want to create the cluster database.

For policy-managed database, select the server pools to be used for creating the database, from the list of existing server pools, or choose to create a new server pool. Policy-managed databases can be created for database versions 11.2 and higher. For database versions lower than 11.2, you will need to select nodes to create the Oracle RAC database.

Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP. You can choose to specify the same or different passwords for each of these user accounts.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

12. In the Storage Locations page, select the storage type for Oracle RAC Database as **File System**.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored. These locations must be on shared storage such as cluster file system location or ASM diskgroups.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify upto nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

13. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Additional Configuration Options page, select custom scripts from the Software Library or your local disk. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click the **Lock** icon to lock the field. Click **Next**.

15. Review the information you have provided and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a Configured status. Click **Next**.

16. Click the Compliance Standards link.

17. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.

18. In the Configure page, click **Next**.

19. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.

20. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.

21. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.

22. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
23. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
24. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly provisioned databases appear as Cloud Control targets.

**Note:**

If the Deployment Procedure fails, then review log files described in [Reviewing Log Files](#).

Provisioning Oracle Real Application Clusters Database with File System on a New Cluster

This section describes how to provision Oracle Real Application Clusters (Oracle RAC) Database with file system on a new cluster. In particular, this section covers the following:

- [Prerequisites for Provisioning Oracle RAC Database with File System on an Existing Cluster](#)
- [Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster](#)

Prerequisites for Provisioning Oracle RAC Database with File System on a New Cluster

Before running the Deployment Procedure, meet the prerequisites listed in [Setting Up Database Provisioning](#).

Procedure for Provisioning Oracle RAC Database with File System on a New Cluster

To provision Oracle RAC databases on a new cluster, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the **Provision Oracle RAC Database** Deployment Procedure and click **Launch**. The Oracle RAC Database provisioning wizard is launched.

3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select **Deploy software and configure** to provision Oracle RAC databases.
- To deploy database software select **Deploy and create a RAC One Node database** which creates a new database and configures it after installing the Oracle RAC database.

Click on the Lock icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.
5. In the Specify OS Users page, specify the operating system users and groups required to provision the database.

For Database User and ASM User, select the Normal User and Privileged User to be added to the OS group.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that the groups corresponding to the following roles already exist on the hosts you select for provisioning.



Note:

To use no root credentials, refer to [Using No Root Credentials for Provisioning Oracle Real Application Clusters \(Oracle RAC\) Databases](#).

- Inventory Group (OINSTALL)
- ASM Database Administrator (ASMDBA)
- ASM Instance Operator (ASMOPER)
- Database Administrator (OSDBA)
- Database Operator (OSOPER)
- ASM Instance Administrator (OSASM)

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a Configured status.

7. Click on the Deploy Software link.
8. In the Select 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. As a designer, you can click on the Lock icon to

lock these fields. These fields will then not be available for editing in the operator role.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.

 Note:

For Windows operating systems, if the Oracle Grid Infrastructure or Oracle Database component selected is of version 12.1 or higher, you can install all services as a named Oracle service user with limited privileges. This will enhance security for database services.

In the Windows Security option section, you can configure the option for an existing user and specify the User Name and Password. Select **Decline Security** option if you want all the services to be installed and configured as an administrative user.

In the Destination location, specify the following:

- **Oracle Base for Grid Infrastructure**, a location on the destination host where the diagnostic and administrative logs, and other logs associated with the Grid Infrastructure can be stored.
- **Grid Infrastructure Home**, 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. Select **Shared Grid Infrastructure home** to enable Grid Infrastructure Oracle Home on shared locations. Ensure that the directory path you provide meets the requirements described in [Requirements for Grid Infrastructure Software Location Path](#).
- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database. Select **Shared Database Oracle home** to enable Database Oracle Home on shared locations.

For Grid Infrastructure, Oracle Base is `/u01/app/user` and Oracle Home is `%ORACLE_BASE/.../grid`. You can use `%ORACLE_BASE%` and `%GI_ORACLE_BASE%` to specify the relative paths which will be interpolated to their respective values.

In the Additional Parameters section, specify the **Working 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. For **Installer Parameters**, 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.

You can also specify OCFS devices in the Installer Parameters field in the following format, separating devices with commas:

Device Number:Partition Number: Drive letter: [DATA | SOFTWARE]

For example, `–ocfs_devices=1:1:E:DATA,1:2:F:DATA`

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a Configured status.

9. Click on the Configure Grid Infrastructure link.
10. In the Select Storage page, select the storage type for Grid Infrastructure as **Automatic Storage Management** and database as **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.

As a designer, you can click on the Lock icon to lock these fields. These fields will then not be available for editing in the operator role.

Click **Next**.

11. In the Configure GI page, in the Basic Settings section, specify the **Cluster Name**, **SCAN Name**, and **SCAN Port**. 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.

If you have chosen storage type as Automatic Storage Management for Grid Infrastructure, in the ASM Storage section, select from the ASM Disk Groups that have been discovered by Cloud Control and are displayed in the table. Click **Add** to add an ASM Disk Group. In the Add/Edit Disk Group dialog box, specify the **Disk**

Group Name, **Disk List**, and specify the redundancy as **Normal**, **High**, or **External**. Click **OK**. Select the OCR/Voting Disk to store cluster registry and voting disk files, and specify the ASM credentials for ASMSNMP and SYS users.

Click **Next**. You will come back to the Configure page. If you have configured the storage options for the Grid Infrastructure and database, the Configure Grid Infrastructure task will have a completed status.

12. Click on the Create Databases link.

13. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle home. The default location is `ORACLE_HOME/assistants/dbca/templates`.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

14. In the Identification and Placement page, select the Oracle RAC database configuration type, whether Policy Managed or Admin Managed.

For admin-managed database, select nodes on which you want to create the cluster database.

For policy-managed database, select the server pools to be used for creating the database, from the list of existing server pools, or choose to create a new server pool. Policy-managed databases can be created for database versions 11.2 and higher. For database versions lower than 11.2, you will need to select nodes to create the Oracle RAC database.

Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP. You can choose to specify the same or different passwords for each of these user accounts.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

15. In the Storage Locations page, select the storage type for Oracle RAC Database as **File System**.

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored. These locations must be on shared storage such as cluster file system location or ASM diskgroups.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify upto five locations.

In the Recovery Files Location section, select **Use same storage type as database files location** to use the same storage type for recovery files as database files. Select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

Select **Enable Archiving** to enable archive logging. Click **Specify Archive Log Locations** and specify upto nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

16. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with datafiles in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click on the Lock icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

17. In the Additional Configuration Options page, select custom scripts from the Software Library or your local disk. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options.

Click on the Lock icon to lock the field. Click **Next**.

18. Review the information you have provided and click **Next**. You will come back to the Configure page. If you have configured the database, the Create Databases task will have a Configured status. Click **Next**.
19. Click the Compliance Standards link.
20. In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.
21. In the Configure page, click **Next**.
22. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.
23. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed.

Click **Next**.
24. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
25. In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
26. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.
27. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly provisioned databases appear as Cloud Control targets.

 Note:

If the Deployment Procedure fails, then review log files described in [Reviewing Log Files](#).

Using No Root Credentials for Provisioning Oracle Real Application Clusters (Oracle RAC) Databases

No root credential is supported for provisioning Oracle RAC databases. To use this feature, do the following:

1. On the Specify OS users page, select **Override Preferred Credentials**. On the Specify OS users dialogue box that appears, create the normal name credential, and then set Run Privilege to **None**. Click **OK**.
2. Select the new normal name credential for both Normal user and Priveleged user.
3. Click **Submit**.

When the database provisioning process reaches the step which requires root credentials, the process will stop. You will need to run the command line manually. To do this, set the environment to \$AGENT_HOME, and then run the command line copy from the Instructions field for the following three steps:

- Execute fixups manually
 - Execute Root scripts manually (for CRS install phase)
 - Execute Root scripts manually (for RAC database install phase)
4. After the command line is run manually using root user for each step, click **Confirm**. The database provisioning process then continues till it completes.

8

Provisioning Oracle Real Application Clusters One (Oracle RAC One) Node Databases

This chapter explains how you can provision Oracle Real Application Clusters One (Oracle RAC One) node databases using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with Provisioning Oracle RAC One Node Databases](#)
- [Deployment Procedures for Provisioning Oracle RAC One Node Databases](#)
- [Provisioning Oracle RAC One Node Databases](#)

Getting Started with Provisioning Oracle RAC One Node Databases

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

Table 8-1 Getting Started with Provisioning Oracle RAC One Node Databases

Step	Description	Reference Links
Step 1	Understanding the Deployment Procedures To provision Oracle RAC One node databases, you will need to run two deployment procedures.	• To learn about the deployment procedures to run for provisioning Oracle RAC One node databases, see Deployment Procedures for Provisioning Oracle RAC One Node Databases.
Step 2	Understanding the Usecase This section lists the use case to provision Oracle RAC One node databases.	• To understand the usecase for provisioning Oracle RAC One node databases, see Provisioning Oracle RAC One Node Databases.
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, and setting up of Oracle Software Library.	• To learn about the prerequisites for provisioning Oracle RAC One node databases, see Prerequisites for Provisioning Oracle RAC One Node Databases.

Table 8-1 (Cont.) Getting Started with Provisioning Oracle RAC One Node Databases

Step	Description	Reference Links
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle RAC One node databases.	To provision Oracle RAC One node databases, follow the steps explained in Procedure for Provisioning Oracle RAC One Node Databases.

Deployment Procedures for Provisioning Oracle RAC One Node Databases

To provision Oracle RAC one database using Cloud Control, use the following Deployment Procedures:

- *Provision Oracle RAC Database + Create Oracle Database*

Use the Provision Oracle RAC Database deployment procedure to provision Oracle RAC database software and then run the Create Oracle Database deployment procedure to create Oracle RAC One databases.

Provisioning Oracle RAC One Node Databases

This section describes how you can provision Oracle RAC one Node databases.

This section covers the following:

- [Prerequisites for Provisioning Oracle RAC One Node Databases](#)
- [Procedure for Provisioning Oracle RAC One Node Databases](#)

Prerequisites for Provisioning Oracle RAC One Node Databases

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

Following are the infrastructure-related prerequisites to be met by the administrator who creates the infrastructure for provisioning deployment procedures:

- Ensure that you meet the infrastructure requirements described in [Setting Up Your Infrastructure](#).
- 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.
- Discover and monitor the destination hosts in Cloud Control. For this purpose, you need the latest version of Oracle Management Agent (Management Agent) on the destination hosts. For more information refer to the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*. Ensure that the agents are installed in the same location on all hosts.

- Set up the Oracle Software Library (Software Library). Ensure that the installation media, database templates, or provisioning entities are available in the Software Library. For information about creating them, see [Setting Up Database Provisioning](#). Alternatively, use a provisioning profile to store the database template. For information about creating a database provisioning profile, see [Creating Database Provisioning Profiles](#).
- Store the operating system credentials of the destination hosts as preferred credentials in Oracle Management Repository (Management Repository) or use Named Credentials.

If you are using SUDO, PowerBroker, see [Setting Up Credentials](#) for information on setting up these authentication utilities.

- The user configuring the deployment procedure will need to be a member of the groups specified below. 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.

The Oracle Database user (typically *oracle*) must be a member of the following groups:

- Inventory Group (OINSTALL) as the primary group
- Database Administrator (OSDBA)
- Database Operator (OSOPER)

The Grid Infrastructure user (typically *grid*) must be a member of the following groups:

- Inventory Group (OINSTALL) as the primary group
- ASM Database Administrator (ASMDBA)
- ASM Instance Operator (ASMOPER)
- ASM Instance Administrator (OSASM)

The Oracle RAC Database user must be a member of the following group:

- ASM Database Administrator (ASMDBA)

- Ensure that you use an operating system user that has write permission on the following locations:
 - 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 you have Operator -Any Target privileges in Cloud Control.

Prerequisites for Operators

Following are the deployment procedure-related prerequisites to be met by the operator who runs the provisioning deployment procedures:

- Ensure that as an operator, you have permissions to view credentials (set and locked by the designer), view targets, submit jobs, and launch deployment procedures.
- 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.
 - Inventory Group (OINSTALL)
 - ASM Database Administrator (ASMDBA)
 - ASM Instance Operator (ASMOPER)
 - Database Administrator (OSDBA)
 - Database Operator (OSOPER)
 - ASM Instance Administrator (OSASM)
- Ensure that you have Operator -Any Target privileges in Cloud Control.

Procedure for Provisioning Oracle RAC One Node Databases

To provision Oracle RAC One node database, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle RAC Database Deployment Procedure and click **Launch**. The Oracle RAC Database provisioning wizard is launched.
3. In the Select Hosts page, if you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then, select the profile with previously saved configuration parameters.

In the Select destination hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

In the Select Tasks to Perform section, select the platform, the version for the process, and the components you want to provision:

- To deploy Grid Infrastructure, select **Deploy software and configure** to provision Oracle RAC databases.
- To deploy database software select **Deploy and create a RAC One Node database** which creates a new database and configures it after installing the Oracle RAC database.

Click the **Lock** icon against the fields that you do not want to be edited in the operator role. For more information about the lock down feature in deployment procedures, see [Introduction to Database Provisioning](#).

Click **Next**.

4. In the Configure page, click on the Setup Hosts link.

5. In the Specify OS Users page, specify the operating system users and groups required to provision the database.

For Database User and ASM User, select the Normal User and Privileged User to be added to the OS group.

Click **Next**.

6. In the Specify OS Groups page, specify the OS Groups to use for operating system authentication. Ensure that these groups already exist on the hosts you select for provisioning.

- Inventory Group (OINSTALL)
- ASM Database Administrator (ASMDBA)
- ASM Instance Operator (ASMOPER)
- Database Administrator (OSDBA)
- Database Operator (OSOPER)
- ASM Instance Administrator (OSASM)

Click **Next**. You will come back to the Configure page. If you have configured the destination hosts, the Setup Hosts task will have a Configured status.

7. Click on the Deploy Software link.

8. In the Select 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. As a designer, you can click the **Lock** icon to lock these fields. These fields will then not be available for editing in the operator role.

In the Source section, select the Software Library location for the **Grid Infrastructure** and **Oracle Database** binaries.

In the Destination location, specify the following:

- **Oracle Base for Grid Infrastructure**, a location on the destination host where the diagnostic and administrative logs, and other logs associated with the Grid Infrastructure can be stored.
- **Grid Infrastructure Home**, 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. Select **Shared Grid Infrastructure home** to enable Grid Infrastructure Oracle Home on shared locations. Ensure that the directory path you provide meets the requirements described in [Requirements for Grid Infrastructure Software Location Path](#).
- **Oracle Base for Database**, 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.
- **Database Oracle Home**, a location on the destination host where the database software can be provisioned. This is the Oracle home directory for the database. Select **Shared Database Oracle home** to enable Database Oracle Home on shared locations.

In the Additional Parameters section, specify the **Working 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. For **Installer Param-**

eters, 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.

Click **Next**. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a Configured status.

9. Click on the Configure Grid Infrastructure link.
10. In the Select Storage page, select the storage type for database as **Automatic Storage Management** or **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.

You can choose either Automatic Storage Management or File System for Oracle RAC Database. The storage type for Grid Infrastructure is set to Automatic Storage Management with Flex redundancy.

As a designer, you can click the **Lock** icon to lock these fields. These fields will then not be available for editing in the operator role.

Click **Next**.

11. In the Configure Cluster page, in the Basic Settings section, specify the **Cluster Name**, **SCAN Name**, and **SCAN Port**. 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.

In the GNS Settings section, select **Configure Grid Naming Service** and specify the **GNS Sub System** and **GNS VIP Address** if you want virtual host names outside the cluster to have dynamically assigned names.

In the GI Network 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.

Click **Add** to add an interface and specify the **Interface Name** and **Interface Subnet** and click **OK**. Select the Usage as **Public**, **Private**, or **Do Not Use** if you do not want to use the interface.

If you have chosen storage type as Automatic Storage Management for either or both Grid Infrastructure and Oracle RAC Database, in the ASM Storage section, select from the ASM Disk Groups that have been discovered by Cloud Control and are displayed in the table. Click **Add** to add an ASM Disk Group. In the Add/Edit Disk Group dialog box, specify the **Disk Group Name**, **Disk List**, and specify the redundancy as **Normal**, **High**, **External**, **Flex**, or **Extended**. Click **OK**. Select the OCR/Voting Disk to store cluster registry and voting disk files, and specify the ASM credentials for ASMSNMP and SYS users.

If you have chosen storage type as File System for Grid Infrastructure or Oracle RAC database, in the File System Storage 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.

As a designer, you can click the **Lock** icon to lock these fields. These fields will then not be available for editing in the operator role.

Click **Next**. You will come back to the Configure page. If you have configured the storage options for the Grid Infrastructure and database, the Configure Grid Infrastructure task will have a completed status. Click **Next**.

12. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the deployment, if any. Click **Next**.
13. In the Schedule page, 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. Click **Next**.
14. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
15. In the Database Procedures page, select the Create Oracle Database Deployment Procedure and click **Launch**. The Create Oracle Database provisioning wizard is launched.
16. In the Database Version and Type page, select the database **Version** and select **Oracle RAC One Node Database**.

In the Cluster section, select the cluster and Oracle Home provisioned earlier. Select a reference host to perform validations to use as reference for creating the database on the cluster.

Select **Cluster Credentials** or add new. Click the plus icon to add new credentials and specify **User Name**, **Password**, and **Run Privileges** and save the credentials.

Click **Next**.

17. In the Database Template page, choose the database template location. The location can be Software Library or Oracle home. The template selected must be compatible with the selected Oracle home version.

If you have selected **Software Library**, click on the search icon and select the template from the Software Library. Specify **Temporary Storage Location on Managed Host(s)**. This location must be present on the reference node that you selected earlier.

Click **Show Template Details** to view details of the selected template. You can view initialization parameters, table spaces, data files, redo log groups, common options, and other details of the template.

If you have selected **Oracle Home**, select the template from the Oracle home. The default location is ORACLE_HOME/assistants/dbca/templates.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

18. In the Identification and Placement page, select nodes on which you want to create the cluster database. Specify **Global Database Name** and **SID** prefix. Specify the **Database Credentials** for SYS, SYSTEM, and DBSNMP. Select the type of Oracle RAC database, whether Policy Managed or Admin Managed. Specify the **Service Name**.

 Note:

Database Service Name is used by applications to connect to the Oracle RAC One Node database and to facilitate online relocation.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

19. In the Storage Locations page, select the storage type, whether File System or Automatic Storage Management (ASM).

In the Database Files Location section, specify the location where data files, temporary files, redo logs, and control files will be stored.

- Select **Use Database File Locations from Template** to select defaults from the template used.
- Select **Use Common Location for All Database Files** to specify a different location.

If you select **Use Oracle Managed Files (OMF)**, in the Multiplex Redo Logs and Control Files section, you can specify locations to store duplicate copies of redo logs and control files. Multiplexing provides greater fault-tolerance. You can specify up to five locations.

In the Recovery Files Location section, select **Use Flash Recovery Area** and specify the location for recovery-related files and Fast Recovery Area Size.

In the Archive Log Settings section, select **Enable Archiving** to enable archive logging. In the Specify Archive Log Locations, you can specify up to data files nine archive log locations. If the log location is not specified, the logs will be saved in the default location.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

20. In the Initialization Parameters page, select the memory management type as **Automatic Memory Management** or **Automatic Shared Memory Management**. Select **Specify Memory Settings as Percentage of Available Memory** to specify memory settings as percentage of available physical memory. For Automatic Shared Memory management, specify **Total SGA** and **Total PGA**. For Automatic Memory Management, specify **Total Memory for Oracle**.

In the Database sizing section, specify the **Block Size** and number of **Processes**. If you have selected a database template with data files in the Database Template page, you cannot edit the Block Size.

Specify the Host CPU Count. The maximum CPU count that can be specified is equal to the number of CPUs present on the host.

In the Character Sets section, select the default character set. The default character set is based on the locale and operating system.

Select a national character set. The default is **AL16UTF16**.

In the Database Connection Mode section, select the dedicated server mode. For shared server mode, specify the number of shared servers.

Click the **Lock** icon to lock the fields you have configured. These fields will not be available for editing in the operator role.

Click **Next**.

21. In the Additional Configuration Options page, select custom scripts from the Software Library. If you have selected a Structure Only database template in the Database Template page, you can also view and edit database options. Click the **Lock** icon to lock the field. Click **Next**.
22. In the Schedule page, specify a Deployment Procedure Instance Name and a schedule for the deployment. If you want to run the procedure immediately, then retain the default selection, that is Immediately. If you want to run the procedure later, then select Later and provide time zone, start date, and start time details. Click **Next**.
23. In the Review page, 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. Click **Save** to save the deployment procedure for future deployment.
24. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the Status link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Provisioning Oracle Real Application Clusters for 12c and 18c

This chapter explains how you can provision Oracle Real Application Clusters (Oracle RAC) for 12c and 18c Release 1 using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with Provisioning Oracle Real Application Clusters for 12c and 18c](#)
- [Core Components Deployed When Provisioning Oracle RAC](#)
- [Cloning a Running Oracle Real Application Clusters](#)
- [Provisioning Oracle Real Application Clusters Using Gold Image](#)
- [Provisioning Oracle Real Application Clusters Using Archived Software Binaries](#)

Getting Started with Provisioning Oracle Real Application Clusters for 12c and 18c

This section helps you get started with this chapter by providing an overview of the steps involved in provisioning Oracle RAC for 12c and 18c Release 1. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision Oracle RAC for 12c and 18c Release 1. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 9-1 Getting Started with Provisioning Oracle RAC

Step	Description	Reference Links
Step 1	Understanding the Components Provisioned Understand the core components provisioned.	To learn about the core components that are provisioned, see Core Components Deployed When Provisioning Oracle RAC .
Step 2	Selecting the Use Case This chapter covers a few use cases for provisioning Oracle RAC. Select the use case that best matches your requirements.	• To learn about cloning an existing Oracle RAC, see Cloning a Running Oracle Real Application Clusters. • To learn about provisioning Oracle RAC using a gold image, see Provisioning Oracle Real Application Clusters Using Gold Image. • To learn about provisioning Oracle RAC using the software binaries from an installation medium, see Provisioning Oracle Real Application Clusters Using Archived Software Binaries.

Table 9-1 (Cont.) Getting Started with Provisioning Oracle RAC

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, and setting up of Oracle Software Library.	- To learn about prerequisites for cloning an existing Oracle RAC, see Prerequisites for Cloning a Running Oracle Real Application Clusters. - To learn about the prerequisites for provisioning Oracle RAC using a gold image, see Prerequisites for Provisioning Oracle Real Application Clusters Using Gold Image. - To learn about the prerequisites for provisioning Oracle RAC using the software binaries from an installation medium, see Prerequisites for Provisioning Oracle Real Application Clusters Using Archived Software Binaries.
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle RAC.	- To clone an existing Oracle RAC, follow the steps explained in Procedure for Cloning a Running Oracle Real Application Clusters. - To provision Oracle RAC using a gold image, follow the steps explained in Procedure for Provisioning Oracle Real Application Clusters Using Gold Image. - To provision Oracle RAC using the software binaries from an installation medium, follow the steps explained in Procedure for Provisioning Oracle Real Application Clusters Using Archived Software Binaries.

Core Components Deployed When Provisioning Oracle RAC

When you provision Oracle RAC, essentially, the Deployment Procedures deploy the following core components:

- Oracle Clusterware
- Oracle RAC Database
- Optionally, Automatic Storage Management (ASM)

You can deploy ASM either in the same Oracle home as the one for Oracle RAC Database, or in a completely different Oracle home (recommended).



Note:

When you run the Deployment Procedures to provision Oracle RAC on a shared file system, the software binaries are installed in the shared location, but the configuration happens on all nodes. To configure new nodes, run the *One Click Extend Cluster Database* procedure to extend the Oracle RAC stack to other nodes.

Cloning a Running Oracle Real Application Clusters

This section describes how you can clone an existing Oracle RAC installation that is running on a host monitored by Cloud Control.

This section covers the following:

- [Prerequisites for Cloning a Running Oracle Real Application Clusters](#)
- [Procedure for Cloning a Running Oracle Real Application Clusters](#)

Prerequisites for Cloning a Running Oracle Real Application Clusters

Before running the Deployment Procedure, meet the following prerequisites.

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- If you want to clone Oracle RAC 11g Release 1 (11.1.0.6) on Solaris platforms, then apply patch# 6486988 on the Oracle home that needs to be cloned.
- Ensure that the target hosts have the necessary hardware and software required for Oracle RAC. The hardware requirements include setting up of the following:
 - Private Network: The network interface cards must be installed on each node and connected to each other.
 - Shared Storage Between Nodes: The shared storage is required for OCR, Voting disks and the data files.
- Ensure that the Virtual IPs are set up in the DNS. If you choose to set up the Virtual IPs locally, then the IP addresses can be specified using the Deployment Procedure, and the procedure will set them up for you.
- 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. For information about creating templates, see [Creating Database Templates](#).
- 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*.

For more information, see Oracle Clusterware Installation Guide available at:

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

- Ensure that the User IDs for operating system users and the Group IDs for operating system groups are identical on all nodes of the cluster.

Prerequisites for Operators

- Ensure that you do NOT use an NIS-based operating system user.
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure and its commands on the target hosts. 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 information about customization, see [Customizing Deployment Procedures](#).

- 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 Cloud Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then from the Host menu, click **Configuration** and then click **Compare**.

- While selecting the source, remember to remove `sqlnet.ora` from the list of files mentioned in **Files to Exclude**.
- Ensure that the umask value on the target host is 022.

Procedure for Cloning a Running Oracle Real Application Clusters

To clone an existing Oracle RAC installation, follow these steps:

1. From the Enterprise menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select one of the following, and click **Launch**.
 - a. To run the Deployment Procedure on UNIX platforms, select **Provision Oracle Clusterware / RAC for UNIX and RDBMS versions 12c/18c**.
 - b. To run the Deployment Procedure on Microsoft Windows platforms, select **Provision Oracle Clusterware / RAC for Windows and RDBMS versions 12c/18c**.

Cloud Control displays the Select Source page of the Deployment Procedure.

3. On the Select Source page, do the following:
 - a. In the Select Source section, select **Select from Existing Installations**. Then click the torch icon for **Reference Host** and select the host on which the existing Oracle RAC installation is running. Once you select the reference host, the application automatically displays the working directory and the details of the selected Oracle Clusterware and Oracle Database.

If you want to save the selected Oracle Clusterware and Oracle Database as gold images in the Software Library, then click **Save to Software Library**. Oracle Clusterware is saved as a *Clusterware Clone* component type and Oracle Database is stored as a *Database Clone* component type, respectively.

 Note:

- Maintain different locations as working directories in case a shared disk is used between the source host and the destination host.
- `sqlnet.ora` hardcodes Oracle base from source install. If you do not remove the file, Oracle tools and utilities will use an incorrect Oracle base and an error message stating that the current location is not writable is displayed.

- b. Click **Next**.
 4. On the Select Hosts page, do the following:
 - a. In the Hosts to Include in Cluster section, click **Add** and select the target hosts that should form the cluster. To see more details about the selected hosts, click **Show Options**.

 Note:

When you click **Add**, the Select Target pop-up window appears. On this page, by default, the **Show Suitable Hosts** option is selected and the table lists only those hosts that are best suited for provisioning. If you do not find the host you want to add, then select **Show All Hosts** to view a complete list of hosts.

By default, Private Host Name and Virtual Host Name are automatically prefilled with values. Edit them and specify values that match with your environment. Optionally, you can also specify their IP addresses.

 Note:

If the prefilled, default values of Private Host Name and Virtual Host Name are incorrect, then see the workaround described in [Troubleshooting Issues](#).

If you already have these details stored in the cluster configuration file, then click **Import From File** to select that cluster configuration file. This file typically contains information about the new hosts to be added. To understand how a cluster configuration file looks, see the sample file shown in [Sample Cluster Configuration File](#).

To configure the private and public network interfaces, click **Select Interfaces**. By default, the interfaces that have the same name and subnet for the selected target hosts are displayed. However, you can also choose to view all the interfaces for the selected target hosts. You can either select one of the existing interfaces or specify a completely new one if the interface you want to use does not exist.

- b. In the Network Interface Configuration section, review the details of the private and public interfaces.
 - c. Click **Next**.
5. On the Credentials/Schedule page, do the following:
- a. In the Reference Host Credentials section, retain the default selection, that is, **Use Preferred Credentials**.

 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 [Customizing Deployment Procedures](#).

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*).

- b. In the Target Host(s) Credentials section, provide the credentials as described in Step 6 (a).

 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*.

- c. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
 - d. Click **Next**.
6. On the Configure Cluster page, do the following:

- a. In the Cluster Name and Location section, review the default name and location details provided for Oracle Clusterware and Oracle RAC Database. While Oracle recommends you to retain the default values, you can always edit them to provide custom values.

For security purposes, the clusterware configuration sets the ownership of Oracle Clusterware home and all its parent directories to be owned by *root*. Hence, Oracle recommends you to install Oracle Clusterware outside the Oracle base of the Oracle RAC home.

The default cluster name you see here is based on the host cluster name you provided in the Agent Deploy application in Cloud Control, while deploying Management Agents on a cluster. The scratch location you see here is a temporary location on the target host where temporary files are placed before provisioning and configuring Oracle RAC.

For **Additional Parameters**, specify any additional parameters you want to run while installing Oracle Clusterware. 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`. Ensure that the parameters are separated by white space.

 Note:

- If you do not see a default cluster name in the **Cluster Name** field, then you might have selected nodes that are not master nodes of the cluster. In this case, manually specify a cluster name, but ensure that the name you specify is the same host cluster name you provided in the Agent Deploy application in Cloud Control, while deploying Management Agents on that cluster.

- b. In the Database Details section, retain the default selection for creating a starter database.

 Note:

If the database creation steps are disabled in the Deployment Procedure, then you will not see this section.

If you want to create a general-purpose database, then leave all the fields in this section blank. Otherwise, provide the required details as described in this step.

If you have a custom response file that already has the options enabled, then select **Use response file to create database**, and specify the full path to a location where the file is available. The file may be available on the target host, in a shared location accessible from the target host, in the Software Library, or in a location where an existing database is running.

 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 target host or a shared location, any template file can be used.

If you do not have a custom response file, then select **Do not use response file**, and provide the global database name, the credentials, and the additional parameters you want to run while creating the starter database.

 Note:

Ensure that the database name you specify is in the format database_name.database_domain. It must have 1 to 8 alphanumeric characters. For example, dbname.example.com. Also note that the credentials you provide are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

If you want to use the structure of an existing database and have a custom template to structure the new database, then in **Template File for Database**, specify the full path to a location where the template file is available. The file may be available on the target host or on a shared location accessible from the target host.

 Note:

If you do not store the response files and templates in a central location, you can always customize the Deployment Procedure to add another step that copies the response file or template to the target host before invoking the configuration tools to create the database.

- c. In the Backup and Recovery Details section, retain the default selection, that is, **Do not Enable Automated Backups** if you do not want to have backups taken.

Alternatively, if you want to enable automated backups, select **Enable Automated Backups**, specify the full path to a directory location from where the backed-up files can be recovered, and provide the operating system credentials for running the backup job. Note that recovery location is the same location as the backup location because this is where the files are backed up and also recovered from.

- d. In the ASM Instance Details section (appears only if you had selected to deploy ASM), retain the default selection, that is, **Create ASM Instance**, and specify the credentials, additional ASM parameters to be used, and the ASM disk string to be used.

 Note:

If you are provisioning Oracle Database 10g and Oracle ASM 10g, then ensure that you specify the same password for database as well as ASM.

If you have a custom response file that already has the options enabled, then select **Use response file to create ASM database**, and specify the full path to a location where the file is available. The file may be available on the target host or on a shared location accessible from the target hosts.

If you do not want to use a response file, then select **Do not use response file**.

- e. Click **Next**.
7. On the Storage page, do the following:
 - a. In the Shared Storage Configuration section, provide details about the storage devices and click **Next**. Specify the partition name and the mount location, and select the mount format and a storage device for storing data. While partition name is the path to the location where the device is installed, mount location is the mount point that represents the partition location.

While configuring the storage device, at a minimum, you must have a partition for at least OCR, Voting Disk, and data files. You cannot designate the same storage device to multiple partitions.

Oracle recommends designating the OCR and the OCR Mirror devices to different partitions. Similarly, Oracle recommends designating the Voting Disk, Voting Disk1, and Voting Disk2 to different partitions.

Before clicking **Next**, do the following:

 - If you want to clear the data on selected raw devices before creating and configuring the cluster, then select **Clear raw devices**.
 - If you have configured only for a few storage devices, then select **Do not provision storage** for others that you do not want to provision.
 - Specify the ASM disk string to be used.
 - b. In the Options section, select the ASM redundancy mode. The default is None, which requires 7 GB of space. While Normal requires 16 GB of space, High requires 32 GB.
8. (Optional) On the Advanced Configuration page, do the following:

 Note:

If the configuration steps are disabled in the Deployment Procedure, then you will not see this page.

- a. In the Bonding Interface (Private Interconnect) section, select **Configure Bonding Interface** if you want to configure the bonding interface. To bind the interfaces, specify details as described in [Table 9-2](#).

- The default mode is *append*. You can however select *edit* to modify, and *replace* to replace the current `sysctl.conf` file.

- On the Review page, review the details you have provided for provisioning Oracle RAC, and click **Submit**. If the details you provided seem to be missing on this page, then see the workaround described in [Troubleshooting Issues](#).
- After the Deployment Procedure ends successfully, instrument the database to collect configuration information.

This section describes how you can provision a gold image of Oracle RAC.

Ensure that you use a gold image that was created using the Oracle home directory of a RAC database. You cannot use a gold image that was created using the Oracle home directory of a standalone database.

- Prerequisites for Provisioning Oracle Real Application Clusters Using Gold Image
- Procedure for Provisioning Oracle Real Application Clusters Using Gold Image

Before running the Deployment Procedure, meet the following prerequisites.

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- Ensure that you create gold images of existing Oracle RAC Database and Oracle Grid Infrastructure.

- Ensure that the target hosts have the necessary hardware and software required for Oracle RAC. The hardware requirements include setting up of the following:

- Private Network: The network interface cards must be installed on each node and connected to each other.
- Shared Storage Between Nodes: The shared storage is required for OCR, Voting disks and the data files.
- Ensure that the Virtual IPs are set up in the DNS. If you choose to set up the Virtual IPs locally, then the IP addresses can be specified using the Deployment Procedure, and the procedure will set them up for you.
- 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 [Creating Database Templates](#).

- 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*.

For more information, see Oracle Clusterware Installation Guide available at:

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

- Ensure that the User IDs for operating system users and the Group IDs for operating system groups are identical on all nodes of the cluster.

Prerequisites for Operators

- Ensure that you do NOT use an NIS-based operating system user.
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure and its commands on the target hosts. 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 information about customization, see [Customizing Deployment Procedures](#).
- While selecting the source, remember to remove *sqlnet.ora* from the list of files mentioned in **Files to Exclude**.
- Ensure that the umask value on the target host is 022.

Procedure for Provisioning Oracle Real Application Clusters Using Gold Image

To provision a gold image of an Oracle RAC installation, follow these steps:

1. From the Enterprise menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select one of the following, and click **Launch**.
 - a. To run the Deployment Procedure on UNIX platforms, Select **Provision Oracle Clusterware / RAC for UNIX and RDBMS versions 12c/18c**.
 - b. To run the Deployment Procedure on Microsoft Windows platforms, select **Provision Oracle Clusterware / RAC for Windows and RDBMS versions 12c/18c**.

Cloud Control displays the Select Source page of the Deployment Procedure.

3. On the Select Source page, do the following:
 - a. In the Select Source section, select **Select from Software Library**.
 - b. In the Source for Clusterware section, click the torch icon and select the generic component that has the gold image of Oracle Clusterware. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

 Note:

If you do not see the required component in the Software Library, then follow the workaround described in [Troubleshooting Issues](#).

- c. In the Source for RAC section, click the torch icon and select the generic component that has the gold image of Oracle Database. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

 Note:

If you do not see the required component in the Software Library, then follow the workaround described in [Troubleshooting Issues](#).

- d. (Optional) In the Source for ASM section, do one of the following:

If you do not want to deploy ASM, then retain the default selection, that is, **Do not Provision ASM**.

If you want to deploy ASM in the same Oracle home as the Oracle RAC, then select **Use the same source as the RAC home**. Alternatively, if you can select **Choose a component** and upload an ASM component from the Software Library.
 - e. Click **Next**.

4. On the Select Hosts page, do the following:
 - a. In the Hosts to Include in Cluster section, click **Add** and select the target hosts that should form the cluster. To see more details about the selected hosts, click **Show Options**.

 Note:

When you click **Add**, the Select Target pop-up window appears. On this page, by default, the **Show Suitable Hosts** option is selected and the table lists only those hosts that are best suited for provisioning. If you do not find the host you want to add, then select **Show All Hosts** to view a complete list of hosts.

By default, Private Host Name and Virtual Host Name are automatically prefilled with values. Edit them and specify values that match with your environment. Optionally, you can also specify their IP addresses.

If the prefilled, default values of Private Host Name and Virtual Host Name are incorrect, then see the workaround described in [Troubleshooting Issues](#).

If you already have these details stored in the cluster configuration file, then click **Import From File** to select that cluster configuration file. This file typically contains information about the new hosts to be added. To understand how a cluster configuration file looks, see the sample file shown in [Sample Cluster Configuration File](#).

To configure the private and public network interfaces, click **Select Interfaces**. By default, the interfaces that have the same name and subnet for the selected target hosts are displayed. However, you can also choose to view all the interfaces for the selected target hosts. You can either select one of the existing interfaces or specify a completely new one if the interface you want to use does not exist.

- b. In the Network Interface Configuration section, review the details of the private and public interfaces.
 - c. Click **Next**.
5. On the Credentials/Schedule page, do the following:
 - a. In the Target Host(s) Credentials section, retain the default selection, that is, **Use Preferred Credentials**.

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 [Customizing Deployment Procedures](#) .

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*.

- b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
 - c. Click **Next**.
6. On the Configure Cluster page, do the following:
- a. In the Cluster Name and Location section, review the default name and location details provided for Oracle Clusterware and Oracle RAC Database. While Oracle recommends you to retain the default values, you can always edit them to provide custom values.

For security purposes, the clusterware configuration sets the ownership of Oracle Clusterware home and all its parent directories to be owned by root. Hence, Oracle recommends you to install Oracle Clusterware outside the Oracle base of the Oracle RAC home.

The default cluster name you see here is based on the host cluster name you provided in the Agent Deploy application in Cloud Control, while deploying Management Agents on a cluster. The scratch location you see here is a temporary location on the target host where temporary files are placed before provisioning and configuring Oracle RAC.

For **Additional Parameters**, specify any additional parameters you want to run while installing Oracle Clusterware. 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. Ensure that the parameters are separated by white space.

- b. In the Database Details section, retain the default selection for creating a starter database.

 Note:

If the database creation steps are disabled in the Deployment Procedure, then you will not see this section.

If you want to create a general-purpose database, then leave all the fields in this section blank. Otherwise, provide the required details as described in this step.

If you have a custom response file that already has the options enabled, then select **Use response file to create database**, and specify the full path to a location where the file is available. The file may be available on the target host, in a shared location accessible from the target host, in the Software Library, or in a location where an existing database is running.

 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 target host or a shared location, any template file can be used.

If you do not have a custom response file, then select **Do not use response file**, and provide the global database name, the credentials, and the additional parameters you want to run while creating the starter database.

 Note:

Ensure that the database name you specify is in the format database_name.database_domain. It must have 1 to 8 alphanumeric characters. For example, dbname.example.com. Also note that the credentials you provide are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

If you want to use the structure of an existing database and have a custom template to structure the new database, then in **Template File for Database**, specify the full path to a location where the template file is available. The file may be available on the target host or on a shared location accessible from the target host.

 Note:

If you do not store the response files and templates in a central location, you can always customize the Deployment Procedure to add another step that copies the response file or template to the target host before invoking the configuration tools to create the database.

- c. In the Backup and Recovery Details section, retain the default selection, that is, **Do not Enable Automated Backups** if you do not want to have backups taken.

Alternatively, if you want to enable automated backups, select **Enable Automated Backups**, specify the full path to a directory location from where the backed-up files can be recovered, and provide the operating system credentials for running the backup job. Note that recovery location is the same loca-

tion as the backup location because this is where the files are backed up and also recovered from.

- d. In the ASM Instance Details section (appears only if you had selected to deploy ASM), retain the default selection, that is, **Create ASM Instance**, and specify the credentials, additional ASM parameters to be used, and the ASM disk string to be used.

 Note:

If you are provisioning Oracle Database 10g and Oracle ASM 10g, then ensure that you specify the same password for database as well as ASM.

If you have a custom response file that already has the options enabled, then select **Use response file to create ASM database**, and specify the full path to a location where the file is available. The file may be available on the target host or on a shared location accessible from the target hosts.

If you do not want to use a response file, then select **Do not use response file**.

- e. Click **Next**.
7. On the Storage page, do the following:
 - a. In the Shared Storage Configuration section, provide details about the storage devices and click **Next**. Specify the partition name and the mount location, and select the mount format and a storage device for storing data. While partition name is the path to the location where the device is installed, mount location is the mount point that represents the partition location.

While configuring the storage device, at a minimum, you must have a partition for at least OCR, Voting Disk, and data files. You cannot designate the same storage device to multiple partitions.

Oracle recommends designating the OCR and the OCR Mirror devices to different partitions. Similarly, Oracle recommends designating the Voting Disk, Voting Disk1, and Voting Disk2 to different partitions.

Before clicking **Next**, do the following:

 - If you want to clear the data on selected raw devices before creating and configuring the cluster, then select **Clear raw devices**.
 - If you have configured only for a few storage devices, then select **Do not provision storage** for others that you do not want to provision.
 - Specify the ASM disk string to be used.
 - b. In the Options section, select the ASM redundancy mode. The default is None, which requires 7 GB of space. While Normal requires 16 GB of space, High requires 32 GB.
 8. (Optional) On the Configuration page, do the following:

**Note:**

If the configuration steps are disabled in the Deployment Procedure, then you will not see this page.

- a. In the Bonding Interface (Private Interconnect) section, select **Configure Bonding Interface** if you want to configure the bonding interface. To bind the interfaces, specify details as described in [Table 9-2](#).
 - b. In the Sysctl File Configuration section, select **Configure Sysctl** file if you want to configure the sysctl.conf file. Specify the mode of editing the system configuration file and the location of the reference system configuration file used for modifying the kernel parameters.

The default mode is *append*. You can however select *edit* to modify, and *replace* to replace the current sysctl.conf file.

Ensure that the reference file you specify is available in a shared location accessible by the Oracle Management Service.
9. On the Review page, review the details you have provided for provisioning Oracle RAC, and click **Submit**. If the details you provided seem to be missing on this page, then see the workaround described in [Troubleshooting Issues](#).
 10. After the Deployment Procedure ends successfully, instrument the database to collect configuration information.

Provisioning Oracle Real Application Clusters Using Archived Software Binaries

This section describes how you can provision Oracle RAC that is identical to the one available on the installation medium.

This section covers the following:

- [Prerequisites for Provisioning Oracle Real Application Clusters Using Archived Software Binaries](#)
- [Procedure for Provisioning Oracle Real Application Clusters Using Archived Software Binaries](#)

Prerequisites for Provisioning Oracle Real Application Clusters Using Archived Software Binaries

Before running the Deployment Procedure, meet the following prerequisites.

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Provisioning Oracle Real Application Clusters for 12c and 18c](#).
- Ensure that you upload the software binaries of Oracle RAC Database and Oracle Grid Infrastructure to the Software Library.

- Ensure that the target hosts have the necessary hardware and software required for Oracle RAC. The hardware requirements include setting up of the following:
 - Private Network: The network interface cards must be installed on each node and connected to each other.
 - Shared Storage Between Nodes: The shared storage is required for OCR, Voting disks and the data files.
- Ensure that the Virtual IPs are set up in the DNS. If you choose to set up the Virtual IPs locally, then the IP addresses can be specified using the Deployment Procedure, and the procedure will set them up for you.
- 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 [Creating Database Templates](#).

- 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*.

For more information, see Oracle Clusterware Installation Guide available at:

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

- Ensure that the User IDs for operating system users and the Group IDs for operating system groups are identical on all nodes of the cluster.

Prerequisites for Operators

- Ensure that you do NOT use an NIS-based operating system user.
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure and its commands on the target hosts. 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 information about customization, see [Customizing Deployment Procedures](#).
- Ensure that the umask value on the target host is 022.

Procedure for Provisioning Oracle Real Application Clusters Using Archived Software Binaries

To provision a fresh Oracle RAC installation, follow these steps:

1. From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**.
2. In the Database Provisioning page, select one of the following, and click **Launch**.
 - a. To run the Deployment Procedure on UNIX platforms, select **Provision Oracle Clusterware / RAC for UNIX and RDBMS versions 12c/18c**.
 - b. To run the Deployment Procedure on Microsoft Windows platforms, select **Provision Oracle Clusterware / RAC for Windows and RDBMS versions 12c/18c**.

Cloud Control displays the Select Source page of the Deployment Procedure.

3. On the Select Source page, do the following:
 - a. In the Select Source section, select **Select from Software Library**.
 - b. In the Source for Clusterware section, click the torch icon and select the generic component that has the software binaries of Oracle Clusterware. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

 Note:

If you do not see the required component in the Software Library, then follow the workaround described in [Troubleshooting Issues](#).

- c. In the Source for RAC section, click the torch icon and select the generic component that has the software binaries of Oracle Database. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

 Note:

If you do not see the required component in the Software Library, then follow the workaround described in [Troubleshooting Issues](#).

- d. (Optional) In the Source for ASM section, do one of the following:

If you do not want to deploy ASM, then retain the default selection, that is, **Do not Provision ASM**.

If you want to deploy ASM in the same Oracle home as the Oracle RAC, then select **Use the same source as the RAC home**. Alternatively, if you can select **Choose a component** and upload an ASM component from the Software Library.
 - e. Click **Next**.

4. On the Select Hosts page, do the following:
 - a. In the Hosts to Include in Cluster section, click **Add** and select the target hosts that should form the cluster. To see more details about the selected hosts, click **Show Options**.

 Note:

When you click **Add**, the Select Target pop-up window appears. On this page, by default, the **Show Suitable Hosts** option is selected and the table lists only those hosts that are best suited for provisioning. If you do not find the host you want to add, then select **Show All Hosts** to view a complete list of hosts.

By default, Private Host Name and Virtual Host Name are automatically prefilled with values. Edit them and specify values that match with your environment. Optionally, you can also specify their IP addresses.

Note:

If the prefilled, default values of Private Host Name and Virtual Host Name are incorrect, then see the workaround described in [Troubleshooting Issues](#).

If you already have these details stored in a cluster configuration file, then click **Import From File** to select that cluster configuration file. This file typically contains information about the new hosts to be added. To understand how a cluster configuration file looks, see the sample file shown in [Sample Cluster Configuration File](#).

To configure the private and public network interfaces, click **Select Interfaces**. By default, the interfaces that have the same name and subnet for the selected target hosts are displayed. However, you can also choose to view all the interfaces for the selected target hosts. You can either select one of the existing interfaces or specify a completely new one if the interface you want to use does not exist.

- b. In the Network Interface Configuration section, review the details of the private and public interfaces.
 - c. Click **Next**.
5. On the Credentials/Schedule page, do the following:
 - a. In the Target Host(s) Credentials section, retain the default selection, that is, **Use Preferred Credentials**.

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 [Customizing Deployment Procedures](#).

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*.

- b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
- c. Click **Next**.
6. On the Configure Cluster page, do the following:
 - a. In the Cluster Name and Location section, review the default name and location details provided for Oracle Clusterware and Oracle RAC Database. While Oracle recommends you to retain the default values, you can always edit them to provide custom values.

For security purposes, the clusterware configuration sets the ownership of Oracle Clusterware home and all its parent directories to be owned by root. Hence, Oracle recommends you to install Oracle Clusterware outside the Oracle base of the Oracle RAC home.

The default cluster name you see here is based on the host cluster name you provided in the Agent Deploy application in Cloud Control, while deploying Management Agents on a cluster. The scratch location you see here is a temporary location on the target host where temporary files are placed before provisioning and configuring Oracle RAC.

For **Additional Parameters**, specify any additional parameters you want to run while installing Oracle Clusterware. 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. Ensure that the parameters are separated by white space.

- b. In the Database Details section, retain the default selection for creating a starter database.

 Note:

If the database creation steps are disabled in the Deployment Procedure, then you will not see this section.

If you want to create a general-purpose database, then leave all the fields in this section blank. Otherwise, provide the required details as described in this step.

If you have a custom response file that already has the options enabled, then select **Use response file to create database**, and specify the full path to a location where the file is available. The file may be available on the target host, in a shared location accessible from the target host, in the Software Library, or in a location where an existing database is running.

 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 target host or a shared location, any template file can be used.

If you do not have a custom response file, then select **Do not use response file**, and provide the global database name, the credentials, and the additional parameters you want to run while creating the starter database.

 Note:

Ensure that the database name you specify is in the format database_name.database_domain. It must have 1 to 8 alphanumeric characters. For example, dbname.example.com. Also note that the credentials you provide are used for SYS, SYSTEM, SYSMAN, and DBSNMP accounts.

If you want to use the structure of an existing database and have a custom template to structure the new database, then in **Template File for Database**, specify the full path to a location where the template file is available. The file may be available on the target host or on a shared location accessible from the target host.

 Note:

If you do not store the response files and templates in a central location, you can always customize the Deployment Procedure to add another step that copies the response file or template to the target host before invoking the configuration tools to create the database.

- c. In the Backup and Recovery Details section, retain the default selection, that is, **Do not Enable Automated Backups** if you do not want to have backups taken.

Alternatively, if you want to enable automated backups, select **Enable Automated Backups**, specify the full path to a directory location from where the backed-up files can be recovered, and provide the operating system credentials for running the backup job. Note that recovery location is the same loca-

tion as the backup location because this is where the files are backed up and also recovered from.

- d. In the ASM Instance Details section (appears only if you had selected to deploy ASM), retain the default selection, that is, **Create ASM Instance**, and specify the credentials, additional ASM parameters to be used, and the ASM disk string to be used.



Note:

If you are provisioning Oracle Database 10g and Oracle ASM 10g, then ensure that you specify the same password for database as well as ASM.

If you have a custom response file that already has the options enabled, then select **Use response file to create ASM database**, and specify the full path to a location where the file is available. The file may be available on the target host or on a shared location accessible from the target hosts.

If you do not want to use a response file, then select **Do not use response file**.

- e. Click **Next**.

7. On the Storage page, do the following:

- a. In the Shared Storage Configuration section, provide details about the storage devices and click **Next**. Specify the partition name and the mount location, and select the mount format and a storage device for storing data. While partition name is the path to the location where the device is installed, mount location is the mount point that represents the partition location.

While configuring the storage device, at a minimum, you must have a partition for at least OCR, Voting Disk, and data files. You cannot designate the same storage device to multiple partitions.

Oracle recommends designating the OCR and the OCR Mirror devices to different partitions. Similarly, Oracle recommends designating the Voting Disk, Voting Disk1, and Voting Disk2 to different partitions.

Before clicking **Next**, do the following:

- If you want to clear the data on selected raw devices before creating and configuring the cluster, then select **Clear raw devices**.
- If you have configured only for a few storage devices, then select **Do not provision storage** for others that you do not want to provision.
- Specify the ASM disk string to be used.

- b. In the Options section, select the ASM redundancy mode. The default is None, which requires 7 GB of space. While Normal requires 16 GB of space, High requires 32 GB.

8. (Optional) On the Configuration page, do the following:

 Note:

If the configuration steps are disabled in the Deployment Procedure, then you will not see this page.

- a. In the Bonding Interface (Private Interconnect) section, select **Configure Bonding Interface** if you want to configure the bonding interface. To bind the interfaces, specify details as described in [Table 9-2](#).
- b. In the Sysctl File Configuration section, select **Configure Sysctl** file if you want to configure the `sysctl.conf` file. Specify the mode of editing the system configuration file and the location of the reference system configuration file used for modifying the kernel parameters.

The default mode is *append*. You can however select *edit* to modify, and *replace* to replace the current `sysctl.conf` file.

Ensure that the reference file you specify is available in a shared location accessible by the Oracle Management Service.

9. On the Review page, review the details you have provided for provisioning Oracle RAC, and click **Submit**. If the details you provided seem to be missing on this page, then see the workaround described in [Troubleshooting Issues](#).
10. After the Deployment Procedure ends successfully, instrument the database to collect configuration information.

Table 9-2 Configuration Page - Element Description

Element	Description
Bonding Device Name	Specify the name of the bond to be created. For example, <code>bond0</code> .
Subnet Mask	Specify the subnet mask for the IP address. For example, <code>255.255.255.0</code> .
Default Gateway	Specify the default gateway for the bonding device. For example, <code>10.1.2.3</code> .
DNS Servers	Specify the Domain Name Server (DNS) list for the bonding device. For multiple DNS servers, the values should be comma-separated. Default values are picked up from the <code>/etc/resolv.conf</code> file. Entries provided here will be appended.
Slave Devices List	Specify the list of slave devices for the bonding device. For multiple slave devices, the values should be comma-separated. For example, <code>eth1,eth2,eth3</code> .
Bonding Mode	Specifies one of four policies allowed for the bonding module. Acceptable values for this parameter are: 0 (Balance-rr)— Sets a round-robin policy for fault tolerance and load balancing. Transmissions are received and sent out sequentially on each bonded slave interface beginning with the first one available. 1 (Active-backup)— Sets an active-backup policy for fault tolerance. Transmissions are received and sent out through the first available bonded slave interface. Another bonded slave interface is only used if the active bonded slave interface fails. 2 (Balance-xor)— Sets an XOR (exclusive-or) policy for fault tolerance and load balancing. Using this method, the interface matches up the incoming request's MAC address with the MAC address for one of the slave NICs. Once this link is established, transmissions are sent out sequentially beginning with the first available interface. 3 (Broadcast)— Sets a round-robin policy for fault tolerance and load balancing. Transmissions are sent out sequentially on each bonded slave interface beginning with the first one available.

Table 9-2 (Cont.) Configuration Page - Element Description

Element	Description
Domain Name	Specify the domain name for the assigned host name. For example, example.com.
Primary Slave Device	Specify the interface name, such as eth0, of the primary device. The primary device is the first of the bonding interfaces to be used and is not abandoned unless it fails. This setting is particularly useful when one NIC in the bonding interface is faster and, therefore, able to handle a bigger load. This setting is only valid when the bonding interface is in active-back-up mode.
ARP Interval	Specify (in milliseconds) how often ARP monitoring occurs. If using this setting while in mode 0 or 2 (the two load-balancing modes) the network switch must be configured to distribute packets evenly across the NICs. The value is set to 0 by default, which disables it.
MII Interval	Specify (in milliseconds) how often MII link monitoring occurs. This is useful if high availability is required because MII is used to verify that the NIC is active to verify that the driver for a particular NIC supports the MII tool. If using a bonded interface for high availability, the module for each NIC must support MII. Setting the value to 0 (the default), turns this feature off. When configuring this setting, a good starting point for this parameter is 100.
MII Interval Down Delay	Specify (in milliseconds) how long to wait after link failure before disabling the link. The value must be a multiple of the value specified in the miimon parameter. The value is set to 0 by default, which disables it.
MII Interval Up Delay	Specify (in milliseconds) how long to wait before enabling a link. The value must be a multiple of the value specified in the miimon parameter. The value is set to 0 by default, which disables it.
NTP Server	Specify the NTP server for the assigned host name. For example, 1.2.3.4.

Sample Cluster Configuration File

The following shows the contents of a typical cluster configuration file:

```
# Cluster Configuration file
```

```
# Node information
```

# Public Node Name	Private Node Name	Private IP (Optional)	Virtual Host Name	Virtual IP (Optional)
node1.example.com	node1-priv.example.com	-	node1-vip.example.com	-
node2.example.com	node2-priv.example.com	10.2.109.103	node2-vip.example.com	134.2.109.103

Provisioning Oracle Real Application Clusters (Oracle RAC) Databases Using No Root Credentials

No root credential is supported for provisioning Oracle RAC databases. To use this feature, do the following:

1. On the Specify OS users page, select **Override Preferred Credentials**. On the Specify OS users dialogue box that appears, create the normal name credential, and then set Run Privilege to **None**. Click **OK**.
2. Select the new normal name credential for both Normal user and Privileged user.
3. Click **Submit**.

When the database provisioning process reaches the step which requires root credentials, the process will stop. You will need to run the command line manually. To do this, set the environment to \$AGENT_HOME, and then run the command line copy from the Instructions field for the following three steps:

- Execute fixups manually
 - Execute Root scripts manually (for CRS install phase)
 - Execute Root scripts manually (for RAC database install phase)
4. Once the command line is run manually using root user for each step, click **Confirm**. The database provisioning process then continues till it completes.

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 using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with Extending Oracle Real Application Clusters](#)
- [Extending Oracle Real Application Clusters](#)

Getting Started with Extending Oracle Real Application Clusters

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 10-1 Getting Started with Extending Oracle RAC

Step	Description	Reference Links
Step 1	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, and setting up of Oracle Software Library.	To learn about the prerequisites for extending Oracle RAC, see Prerequisites for Extending Oracle Real Application Clusters .
Step 2	Running the Deployment Procedure Run the Deployment Procedure to successfully extend an existing Oracle RAC.	To extend Oracle RAC, follow the steps explained in Procedure for Extending Oracle Real Application Clusters .

Extending Oracle Real Application Clusters

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 for Extending Oracle Real Application Clusters](#)
- [Procedure for Extending Oracle Real Application Clusters](#)

Prerequisites for Extending Oracle Real Application Clusters

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- 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.

See [Customizing Deployment Procedures](#).

- 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*.

Procedure for Extending Oracle Real Application Clusters

To extend an existing Oracle RAC, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select **Extend Oracle Real Application Clusters** and click **Launch**.

Cloud Control displays the Extend Real Application Clusters page.

3. 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 do not already exist.

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 Cloud Control. In Cloud Control, click the **Targets** menu, and then click **All Targets**. 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, click **Monitoring Configuration**. On the Configure Target page, specify the clusterware home location and click **Update**.
- Configure Cloud Control settings to display them. In Cloud 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.

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.
 - 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.

 Note:

Ensure that you select nodes that are monitored by Oracle Management Agents 12c Release 1 (12.1.0.1) 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 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 the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#).

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 (*dbainstall*).

 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, see *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**.
- 4. 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 as described in [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.

Deleting or Scaling Down Oracle Real Application Clusters

This chapter describes how you can delete or scale down Oracle Real Application Clusters (Oracle RAC). In particular, this chapter covers the following:

- [Getting Started with Deleting or Scaling Down Oracle Real Application Clusters](#)
- [Deleting the Core Components of Oracle Real Application Clusters](#)
- [Deleting the Entire Oracle RAC](#)
- [Scaling Down Oracle RAC by Deleting Some of Its Nodes](#)

Getting Started with Deleting or Scaling Down Oracle Real Application Clusters

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 11-1 Getting Started with Deleting or Scaling Down an Existing Oracle RAC

Step	Description	Reference Links
Step 1	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 requirements.	• To learn about deleting the entire Oracle RAC stack, see 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 Scaling Down Oracle RAC by Deleting Some of Its Nodes.
Step 2	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, and setting up of Oracle Software Library.	• To learn about the prerequisites for deleting an entire Oracle RAC stack, see Prerequisites for Deleting the Entire Oracle RAC. • To learn about the prerequisites for scaling down an existing Oracle RAC stack by deleting one or more of its own nodes, see Prerequisites for Scaling Down Oracle RAC by Deleting Some of Its Nodes.

Table 11-1 (Cont.) Getting Started with Deleting or Scaling Down an Existing Oracle RAC

Step	Description	Reference Links
Step 3	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 Procedure for Deleting the Entire Oracle RAC. - To scale down an existing Oracle RAC stack by deleting one or more of its own nodes, follow the steps explained in Procedure for Scaling Down Oracle RAC by Deleting Some of Its Nodes.

Deleting the Core Components of Oracle Real Application Clusters

Using the *Delete/Scale down Oracle Real Application Clusters* 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.

This Deployment Procedure enables you to descale 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 [Supported Use Cases and Targets Provisioned Using Database Provisioning Procedures](#), 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 Cloud Control.

Deleting the Entire Oracle RAC

This section describes how you can delete the entire Oracle RAC. In particular, this section covers the following:

- [Prerequisites for Deleting the Entire Oracle RAC](#)
- [Procedure for Deleting the Entire Oracle RAC](#)

Prerequisites for Deleting the Entire Oracle RAC

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- Ensure that *oinstall* and *dba* groups are available.
- Ensure that operating system users such as *oracle* and *grid* 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 [Customizing Deployment Procedures](#).

- REMOVE_ANY_TARGET Enterprise Manager privilege

Procedure for Deleting the Entire Oracle RAC

To delete the entire Oracle RAC, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning page, select **Delete/Scale down Oracle Real Application Clusters**, and click **Launch**.

Cloud Control displays the Delete/Scale down Oracle Real Application Clusters page.

3. 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 al-

so 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 the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#).

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 (*dbaloinstall*).

- e. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
 - f. Click **Review**.
4. 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 [Reviewing Log Files](#).

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 Cloud Control. In particular, this section covers the following:

- [Prerequisites for Scaling Down Oracle RAC by Deleting Some of Its Nodes](#)
- [Procedure for Scaling Down Oracle RAC by Deleting Some of Its Nodes](#)

Prerequisites for Scaling Down Oracle RAC by Deleting Some of Its Nodes

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- 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 [Customizing Deployment Procedures](#).

Procedure for Scaling Down Oracle RAC by Deleting Some of Its Nodes

To scale down Oracle RAC by deleting one or more nodes that are part of it, follow these steps:

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

2. In the Database Provisioning page, select **Delete/Scale down Oracle Real Application Clusters**, and click **Launch**.

Cloud Control displays the Delete/Scale down Oracle Real Application Clusters page.

3. 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 Cloud 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 destination hosts have the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#).

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 (*dbaloinstall*).

- e. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
 - f. Click **Review**.
4. 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 [Reviewing Log Files](#).

Provisioning Oracle Database Replay Client

This chapter explains how you can provision Oracle Database Replay Client using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with Provisioning Oracle Database Replay Client](#)
- [Cloning a Running Oracle Database Replay Client](#)
- [Provisioning an Oracle Database Replay Client Using Gold Image](#)
- [Provisioning an Oracle Database Replay Client Using Installation Binaries](#)

Getting Started with Provisioning Oracle Database Replay Client

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 12-1 Getting Started with Provisioning Oracle Database Replay Client

Step	Description	Reference Links
Step 1	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 requirements.	• To learn about cloning an existing Oracle Database Replay Client, see Cloning a Running Oracle Database Replay Client. • To learn about provisioning Oracle Database Replay Client using a gold image, see Provisioning an Oracle Database Replay Client Using Gold Image. • To learn about provisioning a standalone Oracle Database Replay Client, see Provisioning an Oracle Database Replay Client Using Installation Binaries.
Step 2	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, and setting up of Oracle Software Library.	• To learn about the prerequisites for cloning an existing Oracle Database Replay Client, see Prerequisites for Cloning a Running Oracle Database Replay Client. • To learn about the prerequisites for provisioning Oracle Database Replay Client using a gold image, see Prerequisites for Provisioning an Oracle Database Replay Client Using Gold Image. • To learn about the prerequisites for provisioning a standalone Oracle Database Replay Client, see Prerequisites for Provisioning an Oracle Database Replay Client Using Installation Binaries.

Table 12-1 (Cont.) Getting Started with Provisioning Oracle Database Replay Client

Step	Description	Reference Links
Step 3	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 Procedure for Cloning a Running Oracle Database Replay Client. - To provision Oracle Database Replay Client using a gold image, follow the steps explained in Procedure for Provisioning an Oracle Database Replay Client Using Gold Image. - To provision a standalone Oracle Database Replay Client, follow the steps explained in Procedure for Provisioning an Oracle Database Replay Client Using Installation Binaries.

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 Cloud 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 for Cloning a Running Oracle Database Replay Client](#)
- [Procedure for Cloning a Running Oracle Database Replay Client](#)

Prerequisites for Cloning a Running Oracle Database Replay Client

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- 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 Cloud Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then from the Host menu, click **Configuration** and then click **Compare**.

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 [Customizing Deployment Procedures](#).

- 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*.

Procedure for Cloning a Running Oracle Database Replay Client

To clone an existing instance of Oracle Database Replay Client, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle Database Client Deployment Procedure and click **Launch**. The Oracle Database Client provisioning wizard is launched.
3. 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. Select **Override Preferred Credentials** to specify other credentials.

- 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 hosts. 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 the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#) .

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 host name 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**.
4. 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 [Reviewing Log Files](#).

Provisioning an 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 for Provisioning an Oracle Database Replay Client Using Gold Image](#)
- [Procedure for Provisioning an Oracle Database Replay Client Using Gold Image](#)

Prerequisites for Provisioning an Oracle Database Replay Client Using Gold Image

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- Ensure that the gold image is available either in the Software Library or in a shared, 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 [Customizing Deployment Procedures](#).

- 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.

- 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*.

Procedure for Provisioning an Oracle Database Replay Client Using Gold Image

To provision a gold image of Oracle Database Replay Client from the software library, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle Database Client Deployment Procedure and click **Launch**. The Oracle Database Client provisioning wizard is launched.

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

Cloud 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 [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 provisioning.

- 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 hosts. 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 destination hosts have the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#).

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 host name 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 [Reviewing Log Files](#).

Provisioning an 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 Client 11g Release 1 (11.1.0.6) and Oracle Database Client 11g Release 2, for test database versions 10.2.0.4, 10.2.0.5, 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 for Provisioning an Oracle Database Replay Client Using Installation Binaries](#)
- [Procedure for Provisioning an Oracle Database Replay Client Using Installation Binaries](#)

Prerequisites for Provisioning an Oracle Database Replay Client Using Installation Binaries

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).

 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 Cloud Control Deployment Procedure.

 Note:

If you want to create a component for the software binaries of Oracle Database Replay Client, do not save the shiphome or component to the Components folder in Software Library. Create a new folder in Software Library and then save the component.

- Ensure that the installation binaries are downloaded, and archived and uploaded as a component in the Software Library.
- 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 Cloud Control, click **Targets** and then **Hosts**. On the Hosts page, click the name of the source host to access its Home page, and then from the Host menu, click **Configuration** and then click **Compare**.

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 [Customizing Deployment Procedures](#).

- 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*.

Procedure for Provisioning an Oracle Database Replay Client Using Installation Binaries

To provision a fresh Oracle Database Replay Client, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Procedures page, select the Provision Oracle Database Client Deployment Procedure and click **Launch**. The Oracle Database Replay Client provisioning wizard is launched.

Cloud Control displays the Select Source and Destination page of the Deployment Procedure.

3. 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, Cloud 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 [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 provisioning.

- 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 hosts. 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 the same credentials, which may be different from the preferred credentials, then you can override the preferred credentials with the 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 [Customizing Deployment Procedures](#).

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 host name 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**.
- 4. 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 [Reviewing Log Files](#).

Provisioning Oracle Standby Databases

This chapter contains the following sections:

- [Overview of Creating a Standby Database](#)
- [Creating a New Physical Standby Database \(single-instance only\)](#)
- [Creating a New Logical Standby Database \(single-instance only\)](#)
- [Managing an Existing Standby Database with Data Guard Broker](#)
- [Creating a Primary Database Backup Only](#)

Overview of Creating a Standby Database

Creating a standby database is the first thing you must do before you can manage and monitor the databases. Enterprise Manager Cloud Control provides the Add Standby Database wizard to create a broker configuration that includes a primary database and one or more standby databases.

You can create a physical or a logical standby database. Physical standby databases are physically identical to the primary database, mounted or open read-only when in recovery mode, and support all datatypes and DDL. Logical standby databases are physically different from the primary database, open read-write only when in recovery mode, can be used for data protection and reporting, and support only some databases and DDL.

The standby database creation process performs the following steps:

- Performs an online backup (or optionally uses an existing backup) of the primary database control file, datafiles, and archived redo log files
- Transfers the backup pieces from the primary host to the standby host
- Creates other needed files (e.g., initialization, password) on the standby host
- Restores the control file, datafiles, and archived redo log files to the specified locations on the standby host
- Adds online redo log files and other files to the standby database as needed
- Configures the recovered database as a physical or logical standby database

Creating a New Physical Standby Database (single-instance only)

To create a new physical standby database (single-instance only), follow these steps:



Note:

New physical databases will be created as single-instance databases. The Enterprise Manager Convert to Cluster Database function can be used to convert the standby database to a cluster database after it is created.

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, you see a list of databases. Select the primary database for which you want to create a new physical standby database.
3. On the primary database home page, click **Availability** and then select **Add Standby Database**.
4. On the Database Login page, enter your credentials. Click **Login**.



Note:

You need to connect to the primary database using SYSDBA credentials, if you are not yet connected.

If you log in as a user with SYSDBA privileges, you will have access to all Data Guard functionality, including all monitoring and management features. If you log in as a non-SYSDBA user, you will have access to monitoring functions only; features such as standby creation, switchover, and failover will not be available.

5. On the Add Standby Database page, select **Create a new physical standby database**. Click **Continue**.



Note:

If you choose to create a new physical or logical standby database, Data Guard checks the following when you click Continue:

- Server parameter file (SPFILE) -- Data Guard requires that all databases in a configuration use a server parameter file (SPFILE). If the wizard encounters a primary database that does not use an SPFILE, the wizard stops and returns a message asking you to create one. You can create one with a non-default name. Data Guard only requires that the primary database uses an SPFILE.
- NOARCHIVELOG mode -- Regardless of what method you choose to add the standby database, the primary database must be in ARCHIVELOG mode. If the primary database is in NOARCHIVELOG mode, you will be asked to exit the wizard and put the primary database into ARCHIVELOG mode.

6. The Add Standby Database wizard opens. It takes you through the following steps:
 - Determine the backup type.
 - Set up the backup options.

- Select the Oracle home in which to create the standby database.
- Set up the location for standby database files.
- Provide standby database configuration parameters.
- Review the information before clicking Finish.

Step 1: Determine the backup type

Enterprise Manager uses Oracle Recovery Manager (RMAN) to create a single-instance standby database from a new or existing backup of the primary database. You can select one of two backup operations to use for the standby database creation:

- Perform a live backup of the primary database using RMAN to copy database files, or by copying the database files via staging areas.
- Use an existing RMAN backup or an existing backup from a previous standby database creation.

Click **Next**.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 2: Set up the backup options

For an online backup using RMAN, enter the Degree of Parallelism (the number of parallel channels used by RMAN to copy the database files). The default number is 2.

For an online backup by copying the database files via staging areas, provide the staging area location. You can also choose if you want to retain the directory or delete it after the standby database has been created.

For an offline backup using RMAN, provide the RMAN backup location and the staging area location.

Enter the primary host credentials. You can use existing credentials or create new credentials. If you create new credentials you can save the credentials to use for another database standby creation later. To do this, check the **Set As Preferred Credentials** box.

Click **Next**.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 3: Select the Oracle home in which to create the standby database

The standby database can be created in any Oracle home that was discovered by Oracle Enterprise Manager. Only Oracle homes on hosts that match the operating system of the primary host are shown. You must select a discovered Oracle home and provide a unique instance name for the standby database. Standby host credentials are required to continue.

Step 4: Set up the location for standby database files

Since the primary and standby databases are the same host, the standby database files are placed into an Oracle Optimal Flexible Architecture (OFA) directory structure. Click **Customize** to modify individual file names.

In the Listener Configuration section, specify the name and port of the listener that will be used for the standby database. If a new name and port are specified that are not in use by an existing listener, a new listener using the specified port will be created.

Click **Next**.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 5: Provide standby database configuration parameters

Standby database configuration parameters must be set. These parameters include the database name, database unique name, target name, and standby archive location. The standby archive location can be a regular directory or a flash recovery area. The default values are based on corresponding primary database settings.

After you verify that the parameters are correct, click **Next**.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 6: Review the information before clicking Finish

The Add Standby Database wizard allows one last review of the data you input for the configuration and standby database. Click **Finish** when you are certain all of the information is correct.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

By clicking Standby Database Storage, you can see additional information about all the standby database file locations.

Once you click **Finish**, the standby database creation process runs as an Oracle Enterprise Manager job. You can cancel the standby creation at any point before the job submission.

After the job is submitted, you will be returned to the Data Guard Overview page. In the Status column of the Standby Databases table, you will see Creation in progress listed. If you click that link, you can monitor the progress of the standby database creation.



Note:

To add additional standby databases after the initial creation of the configuration, click **Add Standby Database** to run the Add Standby Database wizard again.

Creating a New Logical Standby Database (single-instance only)

To create a new physical standby database (single-instance only), follow these steps:

 Note:

New logical standby databases will be created as single-instance databases. The Enterprise Manager Convert to Cluster Database function can be used to convert the standby database to a cluster database after it is created.

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, you see a list of databases. Select the primary database for which you want to create a new physical standby database.
3. On the primary database home page, click **Availability** and then select **Add Standby Database**.
4. On the Database Login page, enter your credentials. Click **Login**.

 Note:

You need to connect to the primary database using SYSDBA credentials, if you are not yet connected.

If you log in as a user with SYSDBA privileges, you will have access to all Data Guard functionality, including all monitoring and management features. If you log in as a non-SYSDBA user, you will have access to monitoring functions only; features such as standby creation, switchover, and failover will not be available.

5. On the Add Standby Database page, select **Create a new physical standby database**. Click **Continue**.

 Note:

If you choose to create a new physical or logical standby database, Data Guard checks the following when you click Continue:

- Server parameter file (SPFILE) -- Data Guard requires that all databases in a configuration use a server parameter file (SPFILE). If the wizard encounters a primary database that does not use an SPFILE, the wizard stops and returns a message asking you to create one. You can create one with a non-default name. Data Guard only requires that the primary database uses an SPFILE.
- NOARCHIVELOG mode -- Regardless of what method you choose to add the standby database, the primary database must be in ARCHIVELOG mode. If the primary database is in NOARCHIVELOG mode, you will be asked to exit the wizard and put the primary database into ARCHIVELOG mode.

6. On the database page, in the Standby Databases section, click **Add Standby Database**.
7. The following steps assume a broker configuration already exists with one primary database and one physical standby database, and creates a new logical standby database. It shows how the wizard takes you through additional steps to select the Oracle home for the database and to copy datafiles to the standby database.

The Add Standby Database wizard takes you through the following steps:

- a. Determine the backup type.
- b. Set up the backup options.
- c. Select the Oracle home in which to create the standby database.
- d. Set up the location for standby database files.
- e. Provide standby database configuration parameters.
- f. Review the information before clicking Finish.

Step 1: Determine the backup type

Enterprise Manager uses Oracle Recovery Manager (RMAN) to create a single-instance standby database from a new or existing backup of the primary database. You can select one of two backup operations to use for the standby database creation:

- Perform a live backup of the primary database
- Use an existing backup of the primary database

You can click Cancel to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 2: Set up the backup options

A working directory is needed to store the primary database backup files. It can optionally be retained and used to create additional standby databases in the future. Specify a location on the primary host in which the working directory can be created.

Primary host credentials are required for this step. Enter the credentials of the owner of the primary database Oracle server installation. These credentials can be saved by checking the box marked Save as Preferred Credential.

You can click Cancel to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 3: Select the Oracle home in which to create the standby database

The standby database can be created in any Oracle home that was discovered by Oracle Enterprise Manager. Only Oracle homes on hosts that match the operating system of the primary host are shown. You must select a discovered Oracle home and provide a unique instance name for the standby database. Standby host credentials are required to continue.

Step 4: Set up the location for standby database files

Part of the create broker configuration process involves making the datafiles for the primary database available to the standby host. You have the option of customizing the location for the standby database files. Standby host credentials are required to continue. The following list describes your options:

- Specify the backup file access method
Choose the method by which you want to make the primary database backup files accessible to the standby host. The two options are:
 - Transfer files from the primary host working directory to a standby host working directory
 - Directly access the primary host working directory location from the standby host using a network path name
- Specify the standby database file location
Choose the locations for the standby database files. You have two options:
 - Convert to Oracle OFA (Optimal Flexible Architecture)
 - Keep file names and locations the same as the primary database
- Specify the network configuration file location
Data Guard will add configuration information for the standby database to the network configuration files (listener.ora and tnsnames.ora) in the specified directory on the standby host.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 5: Provide standby database configuration parameters

Standby database configuration parameters must be set. These parameters include the database name, database unique name, target name, and standby archive location. The standby archive location can be a regular directory or a flash recovery area. The default values are based on corresponding primary database settings.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

Step 6: Review the information before clicking Finish

The Add Standby Database wizard allows one last review of the data you input for the configuration and standby database. Click **Finish** when you are certain all of the information is correct.

You can click **Cancel** to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

By clicking Standby Database Storage, you can see additional information about all the standby database file locations.

Once you click **Finish**, the standby database creation process runs as an Oracle Enterprise Manager job. You can cancel the standby creation at any point before the job submission.

After the job is submitted, you will be returned to the Data Guard Overview page. In the Status column of the Standby Databases table, you will see Creation in progress listed. If you click that link, you can monitor the progress of the standby database creation.



Note:

To add additional standby databases after the initial creation of the configuration, click **Add Standby Database** to run the Add Standby Database wizard again.

Managing an Existing Standby Database with Data Guard Broker

To add an existing standby database, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, you see a list of databases. Select the database you want to manage an existing standby database.
3. On the primary database home page, click **Availability** and then select **Add Standby Database**.
4. On the Database Login page, enter your credentials. Click **Login**.



Note:

You need to connect to the primary database using SYSDBA credentials, if you are not yet connected.

If you log in as a user with SYSDBA privileges, you will have access to all Data Guard functionality, including all monitoring and management features. If you log in as a non-SYSDBA user, you will have access to monitoring functions only; features such as standby creation, switchover, and failover will not be available.

5. In the Add Standby Database page, select **Manage an Existing Standby Database with Data Guard Broker**.
6. Select an existing standby database that you want to be managed by the Data Guard broker. The database you choose must have been created from the primary database and must be configured to function as a standby database.

All discovered databases in your environment (both RAC and non-RAC databases) will be shown in the list.

Click **Next**.



Note:

You can click **Cancel** at any time to terminate the current process and begin again at the introductory page of the Add Standby Database wizard.

7. Enter the log in details for the database. You can select **Named** or **New** credentials. For new credentials, create a unique credential. You can set it to Preferred Credential if you want to use it again.

Click **Next**.

8. (optional) Change the Standby Archive Location setting of the existing standby cluster database. Click **Next**.
9. Review the data for the configuration and standby database. Click **Finish**

Creating a Primary Database Backup Only

An additional option is to only create a primary database backup without creating a standby database. This backup can be used for future standby database creations by re-running the Add Standby Database wizard and choosing to create a standby database from an existing backup. The existing backup can be used from either the primary or standby host. Consider this option if you are not able to use the file transfer mechanism provided by Enterprise Manager to transfer the backup files from the primary to the standby hosts, and instead wish to transfer and stage the backup files on the standby host using your own mechanism.

To create a primary database backup only, follow these steps:

1. From the **Targets** menu, select **Databases**.

2. On the Databases page, you see a list of databases. Select the database you want to manage an existing standby database.
3. On the primary database home page, click **Availability** and then select **Add Standby Database**.
4. On the Database Login page, enter your credentials. Click **Login**.

 Note:

You need to connect to the primary database using SYSDBA credentials, if you are not yet connected.

If you log in as a user with SYSDBA privileges, you will have access to all Data Guard functionality, including all monitoring and management features. If you log in as a non-SYSDBA user, you will have access to monitoring functions only; features such as standby creation, switchover, and failover will not be available.

5. In the Add Standby Database page, select **Create a Primary Backup Only**. Click **Continue**.
6. On the Backup Options page, specify a location on the primary host where a directory can be created to store the primary database backup files. Click **Next**.
7. On the Schedule page, specify a name, description, and start time for the backup job.

You can choose to start the backup immediately or at a later time. If you want to start at a later time, set the time and date for when the backup should start.

Click **Next**.
8. Review the data for the configuration and standby database. Click **Finish**

Managing Pluggable Databases Using Enterprise Manager

This chapter explains how you can manage pluggable databases (PDBs) using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started With Managing Pluggable Databases Using Enterprise Manager](#)
- [Overview of Managing Pluggable Databases Using Enterprise Manager](#)
- [Provisioning Pluggable Databases Using Enterprise Manager](#)
- [Relocating Pluggable Databases Using Enterprise Manager](#)
- [Removing Pluggable Databases Using Enterprise Manager](#)
- [Viewing Pluggable Database Job Details Using Enterprise Manager](#)
- [Administering Pluggable Databases Using Enterprise Manager](#)

Getting Started With Managing Pluggable Databases Using Enterprise Manager

This section helps you get started with this chapter by providing an overview of the steps involved in creating a new pluggable database (PDB), cloning a PDB, migrating a non-container database (CDB) as a PDB, unplugging a PDB, and deleting PDBs. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully perform these tasks using Cloud Control. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 14-1 Getting Started with Managing Pluggable Databases

Step	Description	Reference Links
Step 1	Obtaining an Overview Obtain a conceptual overview of PDBs.	To obtain a conceptual overview of PDBs, see Overview of Managing Pluggable Databases Using Enterprise Manager .
Step 2	Selecting the Use Case Among the following use cases, select the one that best matches your requirement: • Creating a new PDB • Plugging in an unplugged PDB • Cloning a PDB • Migrating a non-CDB as a PDB • Unplugging and dropping a PDB • Deleting PDBs	

Table 14-1 (Cont.) Getting Started with Managing Pluggable Databases

Step	Description	Reference Links
Step 3	Meeting the Prerequisites Meet the prerequisites for the selected use case.	- To meet the prerequisites for creating a new PDB, see Prerequisites for Creating a New Pluggable Database. - To meet the prerequisites for plugging in an unplugged PDB, see Prerequisites for Plugging In an Unplugged Pluggable Database. - To meet the prerequisites for cloning a PDB, see Prerequisites for Cloning a Pluggable Database. - To meet the prerequisites for migrating a non-CDB as a PDB, see Prerequisites for Migrating a Non-CDB as a Pluggable Database. - To meet the prerequisites for unplugging and dropping a PDB, see Prerequisites for Unplugging and Dropping a Pluggable Database. - To meet the prerequisites for deleting PDBs, see Prerequisites for Deleting Pluggable Databases.
Step 4	Following the Procedure Follow the procedure for the selected use case.	- To create a new PDB, see Creating a New Pluggable Database. - To plug in an unplugged PDB, see Plugging In an Unplugged Pluggable Database. - To clone a PDB, see Cloning a Pluggable Database. - To migrate a non-CDB as a PDB, see Migrating a Non-CDB as a Pluggable Database. - To unplug and drop a PDB, see Unplugging and Dropping a Pluggable Database. - To delete PDBs, see Deleting Pluggable Databases.

Overview of Managing Pluggable Databases Using Enterprise Manager

An Oracle Database can contain a portable collection of schemas, schema objects, and nonschema objects, that appear to an Oracle Net client as a separate database. This self-contained collection is called a pluggable database (PDB). A multitenant container database (CDB) is a database that includes one or more PDBs. Oracle Database 12c Release 1 (12.1) and later versions allow you to create many PDBs within a single CDB. Applications that connect to databases view PDBs and earlier versions of Oracle Database (earlier than 12.1) in the same manner.

Cloud Control enables administrators to manage the entire PDB lifecycle, including provisioning CDBs, provisioning PDBs (from the seed or from an unplugged PDB), cloning existing PDBs, migrating non-CDBs as PDBs, unplugging PDBs, and deleting PDBs.


 Note:

To manage the PDB lifecycle using Cloud Control, you must have the 12.1.0.3 Enterprise Manager for Oracle Database plug-in, or a later version, deployed. To delete PDBs using Cloud Control, you must have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed.

For information on how to deploy a plug-in and upgrade an existing plug-in, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Figure 14-1 provides a graphical overview of how you can manage the PDB lifecycle in Cloud Control.

Figure 14-1 Managing Pluggable Databases



For more information about PDBs and CDBs, see Managing Plug-Ins in the *Oracle Database Administrator's Guide*.

Provisioning Pluggable Databases Using Enterprise Manager

You can provision PDBs by creating a new PDB within a CDB, by cloning an existing PDB, or by migrating existing non-CDBs to a CDB as PDBs. You can also use unplugged PDBs for provisioning, by plugging them into a CDB.

This section provides information about provisioning a PDB using the Cloud Control console. In particular, it contains the following:

- [Creating a New Pluggable Database Using Enterprise Manager](#)
- [Plugging In an Unplugged Pluggable Database Using Enterprise Manager](#)
- [Cloning a Pluggable Database Using Enterprise Manager](#)
- [Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager](#)



Note:

You can also provision PDBs using EM CLI. For information on how to do so, see [Provisioning Pluggable Databases](#).

Creating a New Pluggable Database Using Enterprise Manager

This section provides information about creating a new PDB using Cloud Control. In particular, it contains the following:

- [Prerequisites for Creating a New Pluggable Database](#)
- [Creating a New Pluggable Database](#)

Prerequisites for Creating a New Pluggable Database

Before creating a new PDB using Cloud Control, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The CDB within which you want to create a PDB must exist, and must be a Cloud Control target.



Note:

For information on how to create a new CDB, see [Creating Databases](#).

- The CDB (within which you want to create a PDB) must not be in read-only, upgrade, or downgrade mode.
- The target host user must be the owner of the Oracle home that the CDB (within which you want to create the PDB) belongs to.

Creating a New Pluggable Database

To create a new PDB in a CDB using Cloud Control, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.



Note:

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, under the Container Database section, select the CDB within which you want to create new PDBs.




Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section, select **Create New Pluggable Databases**.
4. Click **Launch**.

Figure 14-2 displays the Provision Pluggable Database Console.

Figure 14-2 Provision Pluggable Database Console



Note:


You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using *sysdba* user account credentials.

5. In the Creation Options page of the Create Pluggable Database Wizard, in the Pluggable Database Creation Options section, select **Create a New PDB**.
6. In the Container Database Host Credentials section, select or specify the target CDB Oracle home owner host credentials. If you have already registered the cre-

dentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-3 displays the Creation Options page.

Figure 14-3 Creating a New Pluggable Database: Creation Options Page



7. Click **Next**.
8. In the Identification page, enter a unique name for the PDB you are creating.
If you prefer to create more than one PDB in this procedure, then select **Create Multiple Copies**, and set the number of PDBs that you want to create. Note that you can create a maximum of 252 PDBs within a CDB.

 Note:

If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name `accountsPDB`, then the PDBs are created with the names `accountsPDB1`, `accountsPDB2`, `accountsPDB3`, `accountsPDB4`, and `accountsPDB5`.

9. In the PDB Administrator section, enter the credentials of the admin user account you need to create for administering the PDB.

Figure 14-4 displays the Identification page.

Figure 14-4 Creating a New Pluggable Database: Identification Page

cdbam (Container Database)

Creation Options **Identification** Storage Schedule Review

Create Pluggable Database : Identification

PDB Name

A PDB name uniquely identifies a PDB in a CDB. The PDB name is also used as a service name and it is recommended to be unique across all CDBs on a host or cluster.

* PDB Name

☐ Create Multiple Copies

Number of Copies

NOTE For multiple copies, PDB name is generated by appending sequence number (<PDB Name>#)

PDB Administrator

A PDB administrator is a local user with privileges to administer a PDB.

☒ Create PDB Administrator

Username

Password

Confirm Password



Note:

If you choose to create multiple PDBs, then an admin user account is created for each PDB that you create, with the same set of the specified credentials.

10. Click **Next**.
11. In the Storage page, in the PDB Datafile Locations section, select the type of location where you want to store the datafiles.
 - If the target CDB (*CDB in which you are creating the PDB*) is enabled with Oracle Managed Files and if you want to use the same, then select **Use Oracle Managed Files (OMF)**.
 - If you want to enter a custom location, then select **Use Common Location for PDB Datafiles**. Select the storage type and the location where the datafiles can be stored.
12. In the Temporary Working Directory section, enter a location where the temporary files generated during the PDB creation process can be stored.
13. In the Post-Creation Scripts section, select a custom SQL script you want to run as part of this procedure, once the PDB is created.

Figure 14-5 displays the Storage page.

Figure 14-5 Creating a New Pluggable Database: Storage Page

cdbam (Container Database)

Creation Options Identification **Storage** Schedule Review

Create Pluggable Database : Storage

PDB Datafile Locations
Select the storage locations for the PDB(s) to be created.

☐ Use Oracle Managed Files (OMF)
☒ Use Common Location for PDB Datafiles

Storage Type: File System

Location: /scratch/user01/app/user01/oradata/cdbam/prov_pdb

Temporary Working Directory
Specify the location to store temporary files generated during PDB creation.

Temporary Location: /tmp

Post-Creation Scripts
Specify a custom SQL script to be executed after the PDB creation. Optionally, you may select the components from Software Library that contain the custom scripts.

☒ Select from Software Library

SQL Script:

14. Click **Next**.

15. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. Otherwise, select **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, and set the permissible delay time.

Figure 14-6 displays the Schedule page.

Figure 14-6 Creating a New Pluggable Database: Schedule Page

cdbam (Container Database)

Creation Options Identification Storage **Schedule** Review

Create Pluggable Database : Schedule

Deployment Instance: CreatePluggableDatabase_1381225684

Schedule

Start: ☒ Immediately ☐ Later

Grace Period: ☐ Do not run if it cannot start within 1 hours of the scheduled start time

16. Click **Next**.

17. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 14-7 displays the Review page.

Figure 14-7 Creating a New Pluggable Database: Review Page

cdbam (Container Database)

Creation Options Identification Storage Schedule **Review**

Create Pluggable Database : Review

General

Container Database Name	cdbam
Pluggable Database	prov_pdb
PDB Administrator	PDBADMIN
Pluggable Database Creation Options	Seed PDB
Host Name	abc.example.com
Host Credentials	NC_HOST_USER01
Create as Clone	No
Lock All Existing PDB Users	No

Storage

Use Oracle Managed Files (OMF)	No
Location	/scratch/user01/app/user01/oradata/cdbam/prov
Storage Type	File System

18. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

Figure 14-8 displays the Procedure Activity page.

Figure 14-8 Creating a New Pluggable Database: Procedure Activity Page

Procedure Activity: CreatePluggableDatabase_1381224688
 Elapsed Time: 3 minutes, 22 seconds

Select	Name	Type	Status
☒	Initialization	Computational	✓
☐	Snapshot Preparation	Computational	✓
☐	Source PDB Snapshot Creation	Dynamic Proc	✓
☐	Post Snapshot Creation	Computational	✓
☐	Clone Data Preparation	Computational	✓
☐	Target Storage Data Preparation	Computational	✓
☐	Pre PDB Clone Steps Execution	Rolling	✓
☒	Pluggable Databases Creation	Parallel	✓
☐	Post PDB Clone Steps Execution	Computational	✓

Initialization
 Type: Computational
 Elapsed Time: 20 seconds
 Start Date: Oct 8, 2013 2:31:43 AM PDT
 Completed Date: Oct 8, 2013 2:32:03 AM PDT

Step: Evaluate expression (Succeeded)
 Start Date: Oct 8, 2013 2:31:43 AM PDT
 Completed Date: Oct 8, 2013 2:32:03 AM PDT

Pluggable Database Identification validation succeeded

Success:
 Pluggable database name check succeeded.
 Pluggable database admin user check succeeded.

Pluggable Database Storage validation succeeded

Success:
 /oradbocfs/oradata/ is shared across the cluster nodes.

Step: Initialization has been executed successfully.

When you create a new PDB, the Enterprise Manager job system creates a Create Pluggable Database job. For information about viewing the details of this job, see [Viewing Create Pluggable Database Job Details](#).

Plugging In an Unplugged Pluggable Database Using Enterprise Manager

This section provides information about plugging in an unplugged PDB into a CDB, using Cloud Control. In particular, it contains the following:

- [Prerequisites for Plugging In an Unplugged Pluggable Database](#)
- [Plugging In an Unplugged Pluggable Database](#)

Prerequisites for Plugging In an Unplugged Pluggable Database

Before plugging in an unplugged PDB using Cloud Control, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The target CDB (the CDB within which you want to plug in the unplugged PDB) must exist, and must be a Cloud Control target.

Note:

For information on how to create a new CDB, see [Creating Databases](#).

- The target CDB must not be in read-only, upgrade, or downgrade mode.

- The XML file that describes the unplugged PDB, and the other files associated with the unplugged PDB, such as the datafiles and the wallet file, must exist and must be readable.
- The target host user must be the owner of the Oracle home that the CDB (within which you want to plug in the unplugged PDB) belongs to.
- The platforms of the source CDB host (the host on which the CDB that previously contained the unplugged PDB is installed) and the target CDB host (the host on which the target CDB is installed) must have the same endianness, and must have compatible database options installed.
- The source CDB (the CDB that previously contained the unplugged PDB) and the target CDB must have compatible character sets and national character sets. Every character in the source CDB character set must be available in the target CDB character set, and the code point value of every character available in the source CDB character set must be the same in the target CDB character set.

Plugging In an Unplugged Pluggable Database

To plug in an unplugged PDB to a CDB using Cloud Control, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.



Note:

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, under the Container Database section, select the CDB to which you want to add the unplugged PDBs.



Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section, select **Create New Pluggable Databases**.
4. Click **Launch**.

Figure 14-9 displays the Provision Pluggable Database Console.

Figure 14-9 Provision Pluggable Database Console

Container Database
Specify a CDB target on which the PDB operation needs to be performed. You can create a CDB using the Create Oracle Database deployment procedure.

Container Database 

PDB Operations

☐ Migrate Existing Databases
Migrates non-CDBs as new Pluggable Databases

☒ Create New Pluggable Databases
Creates PDBs from sources such as the seed Pluggable Database, unplugged Pluggable Databases, or by cloning an existing Pluggable Database

☐ Unplug Pluggable Databases
Unplugs and drops the Pluggable Database after retaining the datafiles and PDB template which can later be used for plugging back the PDB

☐ Delete Pluggable Databases
Drops the Pluggable Database along with the datafiles

 Note:

You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using *sysdba* user account credentials.

5. In the Creation Options page of the Create Pluggable Database Wizard, in the Pluggable Database Creation Options section, select **Plug an Unplugged PDB**.
6. In the Container Database Host Credentials section, select or specify the target CDB Oracle home owner host credentials. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-10 displays the Creation Options page.

Figure 14-10 Plugging In an Unplugged Pluggable Database: Creation Options Page

The screenshot displays the 'Create Pluggable Database : Creation Options' page. At the top, there's a breadcrumb trail: 'Creation Options' > 'Identification' > 'Storage' > 'Schedule' > 'Review'. The 'Creation Options' tab is active. Under 'Pluggable Database Creation Options', three radio buttons are present: 'Create a New PDB', 'Plug an Unplugged PDB' (which is selected and highlighted with a red box), and 'Clone an Existing PDB'. Below these, there are two more radio buttons: 'Full Clone' and 'Snap Clone'. A 'Source PDB' text field with a search icon is also visible. The 'Container Database Host Credentials' section follows, with a prompt to 'Specify credentials for host "abc.example.com"'. It features three radio buttons: 'Credential' (selected), 'Preferred', 'Named', and 'New'. Below 'Preferred', there's a 'Preferred Credential Name' dropdown menu set to 'Database Host Credentials' and a 'Credential Details' section stating 'Default preferred credentials are not set.' A red box highlights the 'Preferred' radio button and the dropdown menu. At the bottom, a blue note icon precedes the text: 'NOTE Host credentials are required to perform validations and initiate PDB creation on the CDB host or cluster.'

7. Click **Next**.
8. In the Identification page, enter a unique name for the PDB you are plugging in.
Select **Create As Clone** to ensure that Oracle Database generates a unique PDB DBID, GUID, and other identifiers expected for the new PDB.
If you prefer to create more than one PDB in this procedure, then select **Create Multiple Copies**, and set the number of PDBs that you want to create. Note that you can create a maximum of 252 PDBs within a CDB.

Note:

If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name `accountsPDB`, then the PDBs are created with the names `accountsPDB1`, `accountsPDB2`, `accountsPDB3`, `accountsPDB4`, and `accountsPDB5`.

9. In the PDB Administrator section, do one of the following to administer the PDB:
 - If you prefer to use the admin user account that was created as part of the source PDB that you are plugging in, then deselect **Create PDB Administrator**.

- If you want to create a brand new admin user account for the PDB you are plugging in, then select **Create PDB Administrator**, and enter the desired credentials.


 Note:

If you choose to create multiple PDBs, then an admin user account is created for each PDB that you create, with the same set of the specified credentials.

To lock and expire all the users in the newly created PDB, (except the newly created Admin), select **Lock All Existing PDB Users**.

Figure 14-11 displays the Identification page.

Figure 14-11 Plugging In an Unplugged Pluggable Database: Identification Page



The screenshot shows the 'Create Pluggable Database : Identification' page. The 'PDB Name' section includes a text field with 'prov_pdb', checkboxes for 'Create as Clone' and 'Create Multiple Copies', and a 'Number of Copies' spinner set to 2. A note states: 'NOTE For multiple copies, PDB name is generated by appending sequence number (<PDB Name>#)'. The 'PDB Administrator' section includes a checkbox for 'Create PDB Administrator', fields for 'Username' (PDBADMIN), 'Password', and 'Confirm Password', and a checkbox for 'Lock All Existing PDB Users'. A red box highlights the 'PDB Name' and 'PDB Administrator' sections.

10. In the PDB Template Location section, select the location where the source PDB's template is available, and then select the type of PDB template.
 - If the PDB template is available on your CDB host (*CDB to which you are plugging in the unplugged PDB*), then select **Target Host File System**.
 - If the PDB template is a single archive file—a TAR file with datafiles and metadata XML file included in it, then select **Create the PDB from PDB Archive**, then select the PDB template.
 - If the PDB template is a PDB file set—a separate DFB file with all the datafiles and a separate metadata XML file, then select **Create the PDB using PDB File Set**, then select the DBF and XML files.
 - If you want to plug in a PDB using the PDB metadata XML file and the existing datafiles, then select **Create PDB using Metadata file**.

- If the PDB template is available in Oracle Software Library (Software Library), then select **Software Library**, then select the component in the Software Library that contains the PDB template.

Figure 14-12 displays the PDB Template Location section of the Identification page.

Figure 14-12 Plugging In an Unplugged Pluggable Database: PDB Template Location Section

11. Click **Next**.

12. In the Storage page, do one of the following:

- In the previous page, if you chose to create the PDB from a pluggable database archive (single TAR file) or using a pluggable database file set (DFB file and an XML file), then select the type of location where you want to store the target datafiles for the PDB you are plugging in.
 - If the target CDB (*CDB to which you are plugging in the unplugged PDB*) is enabled with Oracle Managed Files and if you want to use the same, then select **Use Oracle Managed Files (OMF)**.
 - If you want to enter a common custom location, then select **Use Common Location for PDB datafiles**. Select the storage type and the location where the datafiles can be stored.

- In the previous page, if you chose to create the PDB using a pluggable database template (XML file only), then do the following:

In the PDB Datafile Locations section, validate the locations mapped for the datafiles. If they are incorrect, correct the paths. Alternatively, if you have a single location where the datafiles are all available, then enter the absolute path in the **Set Common Source File Mapping Location** field, and click **Set**.

You can choose to store the target datafiles for the PDB you are plugging in, in the same location as the source datafiles. However, if you want the target datafiles to be stored in a different location, then select **Copy Datafiles**, and select the type of location:

- If the target CDB (*CDB to which you are plugging in the unplugged PDB*) is enabled with Oracle Managed Files and if you want to use the same, then select **Use Oracle Managed Files (OMF)**.

- If you want to enter a common custom location, then select **Use Common Location for Pluggable Database Files**. Select the storage type and the location where the datafiles can be stored.
 - If you prefer to use different custom locations for different datafiles, then select **Customized Location**, and enter the custom location paths.
13. In the Temporary Working Directory section, enter a location where the temporary files generated during the PDB creation process can be stored.
 14. In the Post-Creation Scripts section, select a custom SQL script you want to run as part of this procedure, once the PDB is plugged in.

If the script is available in the Software Library, select **Select from Software Library**, then select the component that contains the custom script.

Figure 14-13 displays the Storage page.

Figure 14-13 Plugging In an Unplugged Pluggable Database: Storage Page

The screenshot shows the 'Create Pluggable Database : Storage' page for a container database named 'cdbam'. The page has a breadcrumb trail: 'Creation Options' > 'Identification' > 'Storage' > 'Schedule' > 'Review'. The 'Storage' page is divided into three main sections:

- PDB Datafile Locations:** This section asks to 'Select the storage locations for the PDB(s) to be created.' It has three radio buttons: 'Use Oracle Managed Files (OMF)', 'Use Common Location for PDB Datafiles' (which is selected), and 'Use PDB File Locations Same as Source'. Below the radio buttons, there is a 'Storage Type' dropdown menu set to 'File System' and a 'Location' text field containing '/scratch/user01/app/user01/oradata/cdbam/prov_pdb'.
- Temporary Working Directory:** This section asks to 'Specify the location to store temporary files generated during PDB creation.' It has a 'Temporary Location' text field containing '/tmp'.
- Post-Creation Scripts:** This section asks to 'Specify a custom SQL script to be executed after the PDB creation. Optionally, you may select the components from Software Library that contain the custom scripts.' It has a checkbox labeled 'Select from Software Library' which is checked. Below it is an 'SQL Script' text field and a 'Reset' button.

15. Click **Next**.
16. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. Otherwise, select **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, then set the permissible delay time.

Figure 14-14 displays the Schedule page.

Figure 14-14 Plugging In an Unplugged Pluggable Database: Schedule Page

The screenshot shows the 'Schedule' page for creating a pluggable database. At the top, there is a breadcrumb trail: 'Creation Options' > 'Identification' > 'Storage' > 'Schedule' > 'Review'. The main title is 'Create Pluggable Database : Schedule'. Below this, the 'Deployment Instance' is set to 'CreatePluggableDatabase_1381301659'. The 'Schedule' section has a 'Start' dropdown with 'Immediately' selected and 'Later' as an option. A time zone selector shows '(GMT-08:00) Los Angeles - Pacific Time (PT)'. The 'Grace Period' is set to '1 hours' with a checkbox for 'Do not run if it cannot start within'.

17. Click **Next**.

18. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 14-15 displays the Review page.

Figure 14-15 Plugging In an Unplugged Pluggable Database: Review Page

The screenshot shows the 'Review' page for creating a pluggable database. At the top, there is a breadcrumb trail: 'Creation Options' > 'Identification' > 'Storage' > 'Schedule' > 'Review'. The main title is 'Create Pluggable Database : Review'. The 'General' section displays the following details:

Container Database Name	cdbam
Pluggable Database	prov_pdb
Pluggable Database Creation Options	PDB Archive
PDB Archive	/scratch/user01/app/user01/product/12.1.0/db025/assistants/dbca/templates/PROV_PDB.tar.gz (File System)
Host Name	abc.example.com
Host Credentials	USER01
Create as Clone	No
Lock All Existing PDB Users	No

The 'Storage' section displays the following details:

Use Oracle Managed Files (OMF)	No
Location	/scratch/user01/app/user01/oradata/cdbam/prov
Storage Type	File System

19. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

Figure 14-16 displays the Procedure Activity page.

Figure 14-16 Plugging In an Unplugged Pluggable Database: Procedure Activity Page

Procedure Activity: CreatePluggableDatabase_1381224688

Elapsed Time: 3 minutes, 22 seconds

Procedure Steps

Select	Name	Type	Status
☒	Initialization	Computational	✓
☐	Snapshot Preparation	Computational	✓
☐	Source PDB Snapshot Creation	Dynamic Proc	✓
☐	Post Snapshot Creation	Computational	✓
☐	Clone Data Preparation	Computational	✓
☐	Target Storage Data Preparation	Computational	✓
☐	Pre PDB Clone Steps Execution	Rolling	✓
☒	Pluggable Databases Creation	Parallel	✓
☐	Post PDB Clone Steps Execution	Computational	✓

Initialization x Pluggable Databases Creation x

✓ **Initialization**

Type Computational Start Date Oct 8, 2013 2:31:43 AM PDT
Elapsed Time 20 seconds Completed Date Oct 8, 2013 2:32:03 AM PDT

Step: Evaluate expression (Succeeded)

Start Date Oct 8, 2013 2:31:43 AM PDT
Completed Date Oct 8, 2013 2:32:03 AM PDT

Pluggable Database Identification validation succeeded

Success:
Pluggable database name check succeeded.
Pluggable database admin user check succeeded.

Pluggable Database Storage validation succeeded

Success:
/oradbocfs/oradata/ is shared across the cluster nodes.

Step: Initialization has been executed successfully.

When you plug in an unplugged PDB, the Enterprise Manager job system creates a Create Pluggable Database job, see [Viewing Create Pluggable Database Job Details](#).

Cloning a Pluggable Database Using Enterprise Manager

You can clone a PDB using either the Full Clone method, or the Snap Clone method. This section provides information about cloning a PDB using these methods, in Cloud Control. In particular, it contains the following:

- [Prerequisites for Cloning a Pluggable Database](#)
- [Cloning a Pluggable Database](#)

Prerequisites for Cloning a Pluggable Database

To clone a PDB using Cloud Control, you must meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The source PDB (the PDB that you want to clone) must exist, and must be a Cloud Control target.

Note:

For information on how to create a new PDB, see [Creating a New Pluggable Database Using Enterprise Manager](#).

- The source PDB must be open.
- The target CDB (the CDB into which you want to plug in the cloned PDB) must exist, and must be a Cloud Control target.

**Note:**

For information on how to create a new CDB, see [Creating Databases](#).

- The target CDB must not be in read-only, upgrade, or downgrade mode.
- The target host user must be the owner of the Oracle home that the target CDB belongs to.

To clone a PDB using the Snap Clone method, you must meet the following additional prerequisites:

- The 12.1.0.5 Enterprise Manager for Oracle Database plug-in must be downloaded and deployed. Also, the 12.1.0.3 SMF plug-in or higher must be downloaded and deployed.
- The PDB that you want to clone must reside on a registered storage server. This storage server must be synchronized.

For information on how to register a storage server and synchronize storage servers, see *Oracle Enterprise Manager Cloud Administration Guide*.

- All the datafiles of the PDB that you want to clone must reside on the storage volumes of the storage server, and not on the local disk.
- Metric collections must be run on the source CDB (the CDB containing the PDB that you want to clone), the source CDB host, and the PDB that you want to clone.
- The Snap Clone feature must be enabled for the PDB that you want to clone.

For information on how to enable the Snap Clone feature, see *Oracle Enterprise Manager Cloud Administration Guide*.

Cloning a Pluggable Database

To clone an existing PDB using either the Full Clone or the Snap Clone method, follow these steps:

**Note:**

If you use the Full Clone method to clone a PDB, you can clone the PDB only to the source CDB (the CDB containing the PDB that you are cloning).

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.

**Note:**

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, in the CDB section, select the CDB to which you want to add the cloned PDB.


 Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section, select **Create New Pluggable Databases**.
4. Click **Launch**.

Figure 14-17 displays the Provision Pluggable Database Console.

Figure 14-17 Provision Pluggable Database Console



 Note:

You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using *sysdba* user account credentials.

5. In the Creation Options page of the Create Pluggable Database Wizard, in the PDB Creation Options section, select **Clone an Existing PDB**.

To clone a PDB using the traditional method of cloning the PDB datafiles, select **Full Clone**. Use this method if you want to clone a PDB for long term usage. This method is ideal for load testing, when you plan to make significant data updates to the PDB clone. However, this method takes a longer period of time, and a clone that is created using this method occupies a fairly large amount of space, as compared to the Snap Clone method.

To clone a PDB using the Storage Management Framework (SMF) Snap Clone feature, select **Snap Clone**. Use this method if you want to clone a PDB for short term purposes. This method is ideal for functional testing, as the cloning process is quick, and a PDB clone that is created using this method occupies very little space. However, this method is not suitable if you plan to make significant data updates to the PDB clone.

For **Source PDB**, select the PDB that you want to clone.

Figure 14-18 displays the PDB Creation Options section of the Creation Options page.

Figure 14-18 Cloning a Pluggable Database: Specifying Creation Options

6. In the CDB Host Credentials section, select or specify the target CDB Oracle Home owner host credentials. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-19 displays the CDB Host Credentials section of the Creation Options page.

Figure 14-19 Cloning a Pluggable Database: Specifying CDB Host Credentials

7. Click **Next**.
8. In the Identification page, enter a unique name for the PDB you are cloning.
If you prefer to create more than one PDB in this procedure, then select **Create Multiple Copies**, and set the number of PDBs you want to create. Note that you can create a maximum of 252 PDBs.

 Note:


If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all the cloned PDBs, and the suffix is a numeric value that indicates the count of PDBs.

For example, if you create five PDBs with the name `accountsPDB`, then the PDBs are created with the names `accountsPDB1`, `accountsPDB2`, `accountsPDB3`, `accountsPDB4`, and `accountsPDB5`.

9. In the PDB Administrator section, do one of the following to administer the PDB:
 - If you prefer to use the admin user account that was created as part of the source PDB that you are cloning, then deselect **Create PDB Administrator**.
 - If you want to create a brand new admin user account for the PDB you are cloning, then select **Create PDB Administrator**, and enter the desired credentials.

Figure 14-20 displays the Identification page.

Figure 14-20 Cloning a Pluggable Database: Identification Page



The screenshot shows the 'Create Pluggable Database: Identification' page. At the top, there's a breadcrumb trail: 'Creation Options' > 'Identification' > 'Storage' > 'Schedule' > 'Review'. Below this is the title 'Create Pluggable Database : Identification'. The main section is 'PDB Name', which includes a text field for 'PDB Name' containing 'prov_pdb', a checkbox for 'Create Multiple Copies' (unchecked), and a 'Number of Copies' spinner set to 2. A note below states: 'NOTE For multiple copies, PDB name is generated by appending sequence number (<PDB Name>#)'. The next section is 'PDB Administrator', which has a checkbox for 'Create PDB Administrator' (checked). Below this are three input fields: 'Username' (containing 'PDBADMIN'), 'Password', and 'Confirm Password'. Red boxes are drawn around the 'PDB Name' section and the 'PDB Administrator' section.

 Note:

If you choose to create multiple PDBs, then an admin user account is created for each PDB that you create, with the same set of the specified credentials.

10. In the Source CDB Login Credentials section, select or specify the login credentials of the source CDB. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

The credentials are used to bring the source PDB to read-only mode before the cloning operation begins, and to restore it to the original state after the cloning operation ends.

If you chose the Snap Clone method (on the Source page of the Create Pluggable Database Wizard) to clone the PDB, specify the host credentials for the source CDB.



Note:

If you are cloning the source PDB to the source CDB itself, then the Source CDB Login Credentials section is not displayed, that is, you do not need to provide the source CDB login credentials or the source CDB host credentials.

If you are cloning the source PDB to a CDB different from the source CDB, and this CDB resides on the source CDB host, then you must provide the source CDB login credentials. You do not need to provide the source CDB host credentials.

If you are cloning the source PDB to a CDB different from the source CDB, and this CDB resides on a host different from the source CDB host, then you must provide the source CDB login credentials and the source CDB host credentials.

11. Click **Next**.

12. In the Storage page, specify the storage information.

If you chose the Full Clone method to clone the PDB, select the type of location where you want to store the PDB datafiles in the following manner:

- If the source CDB is enabled with Oracle Managed Files and if you want to use the same, then select **Use Oracle Managed Files (OMF)**.
- If you want to enter a custom location, then select **Use Common Location for PDB Datafiles**. Select the storage type and the location where the datafiles can be stored.

Figure 14-21 displays the Storage page for the Full Clone method.

Figure 14-21 Cloning a Pluggable Database (Full Clone): Storage Page

The screenshot displays the 'Create Pluggable Database : Storage' page for a container database named 'sicdb'. The page is part of a wizard with steps: Creation Options, Identification, Storage (selected), Schedule, and Review. A message states 'All validations succeeded. See details.' The 'PDB Datafile Locations' section is highlighted with a red box. It contains two radio buttons: 'Use Oracle Managed Files (OMF)' and 'Use Common Location for PDB Datafiles' (selected). Below these, the 'Storage Type' is set to 'File System' and the 'Location' is '/scratch/aime/app/aime/oradata/sicdb/prov_pdb'.


If you chose the Snap Clone method to clone the PDB, do the following:

- In the PDB Datafile Locations section, specify a value for **Mount Point Prefix**, that is, the mount location for the storage volumes. You can choose to specify the same prefix for all the volumes, or a different prefix for each volume. Also, specify a value for **Writable Space**, that is, the space that you want to allocate for writing the changes made to the PDB clone. You can choose to specify the same writable space value for all the volumes, or a different value for each volume.
- In the Privileged Host Credentials section, select or specify the credentials of the *root* user. These credentials are used for mounting the cloned volumes on the destination host.

If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-22 displays the Storage page for the Snap Clone method.

Figure 14-22 Cloning a Pluggable Database (Snap Clone): Storage Page



sidb (Container Database)

Creation Options Identification **Storage** Schedule Review

Create Pluggable Database : Storage

PDB Datafile Locations
Specify mount points and writable space for the done volumes. The writable space required for the done volumes depends on the amount of updates planned on the cloned PDB(s).

Source PDB Size 886 MB

Volume	Mount Point Prefix	Writable Space (GB)
abc:dev1:vol_1		0.5

Privileged Host Credentials
Specify privileged host credentials for "abc.example.com"

Credential ☒ Preferred ☐ Named ☐ New

Preferred Credential Name Normal Host Credentials

Credential Details Default preferred credentials are not set.

☒ **NOTE** Privileged host credentials are required to mount the volumes at the specified locations.


13. In the Temporary Working Directory section, enter a location where the temporary files generated during the PDB creation process can be stored.
14. In the Post-Creation Scripts section, select a custom SQL script you want to run as part of this procedure, once the PDB is cloned.
15. Click **Next**.
16. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. Otherwise, select **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, and set the permissible delay time.

Figure 14-23 displays the Schedule page.

Figure 14-23 Cloning a Pluggable Database: Schedule Page



sidb (Container Database)

Creation Options Identification Storage **Schedule** Review

Create Pluggable Database : Schedule

Deployment Instance CreatePluggableDatabase_1381312636

Schedule

Start ☒ Immediately ☐ Later (GMT-08:00) Los Angeles - Pacific Time (PT)

Grace Period ☐ Do not run if it cannot start within 1 hours of the scheduled start time

17. Click **Next**.
18. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 14-24 displays the Review page.

Figure 14-24 Cloning a Pluggable Database: Review Page

Create Pluggable Database : Review

General

Container Database Name: sicdb

Pluggable Database: prov_pdb

Pluggable Database Creation Options: Source PDB (Snap Clone)

Host Name: abc.example.com

Host Credentials: AIME_NORMAL

Source PDB: sicdb_PROV_PDB_MAST

Storage

Volume	Mount Point Prefix	Writable Space (GB)
abc:dev_1:vol_1	/abc	0.5

19. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

Figure 14-25 displays the Procedure Activity page.

Figure 14-25 Cloning a Pluggable Database: Procedure Activity Page

Procedure Activity: CreatePluggableDatabase_1381224688

Elapsed Time: 3 minutes, 22 seconds

Procedure Steps

Select	Name	Type	Status
☒	Initialization	Computational	Success
☐	Snapshot Preparation	Computational	Success
☐	Source PDB Snapshot Creation	Dynamic Proc	Success
☐	Post Snapshot Creation	Computational	Success
☐	Clone Data Preparation	Computational	Success
☐	Target Storage Data Preparation	Computational	Success
☐	Pre PDB Clone Steps Execution	Rolling	Success
☒	Pluggable Databases Creation	Parallel	Success
☐	Post PDB Clone Steps Execution	Computational	Success

Initialization

Type: Computational

Start Date: Oct 8, 2013 2:31:43 AM PDT

Elapsed Time: 20 seconds

Completed Date: Oct 8, 2013 2:32:03 AM PDT

Step: Evaluate expression (Succeeded)

Start Date: Oct 8, 2013 2:31:43 AM PDT

Completed Date: Oct 8, 2013 2:32:03 AM PDT

Pluggable Database Identification validation succeeded

Success:
Pluggable database name check succeeded.
Pluggable database admin user check succeeded.

Pluggable Database Storage validation succeeded

Success:
/oradbocfs/oradata/ is shared across the cluster nodes.

Step: Initialization has been executed successfully.

When you clone a PDB, the Enterprise Manager job system creates a Create Pluggable Database job, see [Viewing Create Pluggable Database Job Details](#).

**Note:**

The pluggable database cloning procedure contains a prerequisite step that checks the plug compatibility which is disabled by default. If disabled, this step will succeed irrespective of any violations. If enabled, this step will fail if there are any plug compatibility violations found.

To enable this step, run the following command:

```
emctl set property -name oracle.sysman.db.pdb.prereq_enabled -  
sysman_pwd <sysman password> -value true
```

To disable this step, run the following command:

```
emctl set property -name oracle.sysman.db.pdb.prereq_enabled -  
sysman_pwd <sysman password> -value false
```

Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager

This section provides information about migrating a non-CDB as a PDB using Cloud Control. In particular, it contains the following:

- [Prerequisites for Migrating a Non-CDB as a Pluggable Database](#)
- [Migrating a Non-CDB as a Pluggable Database](#)

Prerequisites for Migrating a Non-CDB as a Pluggable Database

Before migrating a non-CDB as a PDB using Cloud Control, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The target CDB (the CDB to which you want to migrate a non-CDB as a PDB) must exist, and must be a Cloud Control target.

**Note:**

For information on how to create a new CDB, see [Creating Databases](#).

- The target CDB must not be in read-only, upgrade, or downgrade mode.
- The non-CDB that you want to migrate and the target CDB must be running in ARCHIVELOG mode.

For information on setting the archiving mode of a database, see *Oracle Database Administrator's Guide*.

- The database administrators of the database you want to migrate, and the target CDB must have SYSDBA privileges.
- The target host user must be the owner of the Oracle home that the target CDB belongs to.

Migrating a Non-CDB as a Pluggable Database

To migrate a non-CDB as a PDB using Cloud Control, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.

Note:

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, in the CDB section, select the CDB to which you want to migrate a non-CDB as a PDB.

Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section of the Provision Pluggable Database page, select the **Migrate Existing Databases** option and click **Launch**.

Figure 14-26 displays the Provision Pluggable Database Console.

Figure 14-26 Provision Pluggable Database Console

Container Database
Specify a CDB target on which the PDB operation needs to be performed. You can create a CDB using the Create Oracle Database deployment procedure.

Container Database

PDB Operations

☒ **Migrate Existing Databases**
Migrates non-CDBs as new Pluggable Databases

☐ **Create New Pluggable Databases**
Creates PDBs from sources such as the seed Pluggable Database, unplugged Pluggable Databases, or by cloning an existing Pluggable Database

☐ **Unplug Pluggable Databases**
Unplugs and drops the Pluggable Database after retaining the datafiles and PDB template which can later be used for plugging back the PDB

☐ **Delete Pluggable Databases**
Drops the Pluggable Database along with the datafiles

4. On the Database Login page, select the Credential Name from the drop-down list. Click **Login**.

- On the Migrate Non-CDBs launch page, select a data migration method, that is, **Export/Import** or **Plug as a PDB**. If you select **Plug as a PDB**, ensure that the non-CDB that you want to migrate is open, and is in read-only mode.

Enter the appropriate credentials for the Oracle Home Credential section.

Click **Next**.

Figure 14-27 displays the Method page.

Figure 14-27 Migrating a Non-CDB: Method Page

- On the Database page, select a Non-CDB to be migrated. You can select more than one. Click **Add**. In the database pane, provide the appropriate credential, properties, export, import, and datafile location information. Click **Next**.

Figure 14-28 displays how to select the non-CDB you want to migrate.

Figure 14-28 Migrating a Non-CDB: Selecting a Non-CDB

Figure 14-29 displays how to specify the database and database host credentials.

Figure 14-29 Migrating a Non-CDB: Specifying Credentials

Source Information

Database Credentials

Enter the credentials for the database: uatdb01.example.com.

Credential: ☒ Preferred ☐ Named ☐ New

Preferred Credential Name: SYSDBA Database Credentials

Attribute	Value
Username	sys
Password	*****
Role	sysdba

[More Details](#)

Database Host Credentials

Enter the database host credentials for the host: abcdefgh.example.com with write permission on the export directory.

Credential: ☒ Preferred ☐ Named ☐ New

Preferred Credential Name: Database Host Credentials

Attribute	Value
Username	oracle
Password	*****

[More Details](#)

Figure 14-30 displays how to specify the stage location for migration.

Figure 14-30 Migrating a Non-CDB: Specifying Stage Location

Export Options

Specify a location on the non-CDB host that will be used to temporarily stage the migration metadata and/or datafiles.

* Export: /tmp

Figure 14-31 displays how to specify the PDB administrator details and datafile location.

Figure 14-31 Migrating a Non-CDB: Specifying PDB Properties and Datafile Location

Destination Information

PDB Properties

* PDB Name: uatdb01

* PDB Administrator: pdbadmin

* Password: *****

* Confirm Password: *****

Datafile Location

Specify a location on the CDB host where datafiles for the newly created PDB are placed.

* Datafile Location: /app/oracle/product/oradata/devtd

☒ Create Subdirectory

Figure 14-32 displays how to specify the import directory.

Figure 14-32 Migrating a Non-CDB: Specifying Import Directory

Import Directory

Specify a location on the CDB host that will be used to temporarily stage the migration metadata and/or datafiles.

This directory will be cleaned up upon successful migration.

* Import Directory: /product/admin/devtdb01/dpdump/

- On the Schedule page, enter the appropriate job and scheduling details. Click **Next**.

Figure 14-33 displays the Schedule page.

Figure 14-33 Migrating a Non-CDB: Schedule Page

Migrate Non-CDBs : Schedule

Job Details

* Job Name: Migrate_Into_devdb01_29-01-2013_015134_790

Job:

Description:

Schedule

Start: ☒ Immediately ☐ Later (GMT-06:00) Chicago - Central Time (CT)

Grace Period: ☐ Do not run if it cannot start within 1 hours of the scheduled start time

Duration: ☒ Indefinitely ☐ For 1 hours ☐ Until

8. On the Review page, review all details entered. If there are no changes required, click **Submit**.

Figure 14-34 displays the Review page.

Figure 14-34 Migrating a Non-CDB: Review Page

Migrate Non-CDBs : Review

General

CDB database

Data Migration Method Plug as a PDB

Job Name Migrate_Into_devdb01_29-01-2013_015134_790

Job Description

Databases

The following databases will be migrated to the CDB.

Database	Host	Export Directory	Import Directory	PDB Name	PDB Administrator	Datafile Location
uatdb01.example.com	abcdefghi.example.com	/tmp/exp_uatdb01_29-01-2013_015134_790	/app/oracle/product/admin/devdb01/dpdu_uatdb01	uatdb01	pdadmin	/app/oracle/prod

Number of Databases to be migrated: 1

Relocating Pluggable Databases Using Enterprise Manager

This section provides information about relocating a Pluggable Database (PDB). By following the steps for relocation, a Pluggable Database is moved from one Container Database (source) to another (destination). While cloning a Pluggable Database to a different Container Database (CDB) creates a copy of the PDB while the original still exists in the parent CDB, relocation completely moves the PDB from one CDB to another.

Note:

The relocation of Pluggable Databases is applicable only to Database versions 12.2 and above.

In particular, this section contains the following:

- [Prerequisites for Relocating Pluggable Databases](#)
- [Relocating Pluggable Databases](#)

Prerequisites for Relocating Pluggable Databases

Before relocating a PDB of version 12.2 or above using Cloud Control, ensure that you meet the following prerequisite:

- Both the source and destination Container Databases should be in Archive Log mode and have Local Undo configured.

Relocating Pluggable Databases

To relocate a Pluggable Database from one Container Database to another follow the steps below:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source PDB and right-click the name of the PDB that you want to relocate.
4. From the **Oracle Database** menu, and select **Relocate**.
5. On the Source and Destination: Relocate Pluggable Database page, do the following:
 - a. Specify the SYSDBA credentials for the source CDB. You can choose to use the preferred credentials, use a saved set of named credentials, or specify a new set of credentials.
 - b. In the Container Database Destination section, select the destination CDB and specify a name, and a display name for the PDB.
 - c. In the Credentials section, specify the SYSDBA credentials of the CDB and the CDB host credentials.
 - d. If you have an existing database link which connects from the destination CDB to the source CDB, then select the existing database link.

The privileges required by the existing database link are 'CREATE SESSION', 'SYSOPER', and 'CREATE PLUGGABLE DATABASE'.

 Note:

This step is optional. If the existing database link is not selected, then a new database link is created automatically.

- e. Select either of the following Relocation modes:
 - Maximum Availability ensures there is no need to update the Connect Descriptor soon after the PDB relocation. You may do it at a later time, if required.
 - Maximum Performance requires you to update the Connect Descriptor after the PDB relocation.
6. Click **Relocate**.

**Note:**

- To edit other parameters of the PDB you can click **Advanced** and edit the parameters which are similar to the PDB clone wizard. For more information, see [Creating a Full Clone Pluggable Database Using the Clone Wizard](#).
- If you have chosen Maximum Availability as the relocation mode, you will see the PDB as a managed target in the source CDB. However, the PDB status is always shown as down there. To eliminate the PDB listing from the source CDB, update the Connect Descriptor manually and then run the following command from the source CDB:

```
emcli complete_post_pdb_relocation -pdb_target_name="source  
pluggable database target name" -input_file="path of the  
input file containing properties"
```

The properties that need to be mentioned in the input file are:

- SRC_CDB_CRED = SYSDBA source container database credentials (format - CRED_NAME:OWNER).
- SRC_HOST_CRED = Source container database host credentials (format - CRED_NAME:OWNER).

Command example: `emcli complete_post_pdb_relocation -pdb_target_name="PRODCDB_GEN" -input_file=data:/u01/post_relocate.props`

Removing Pluggable Databases Using Enterprise Manager

This section provides information related to unplugging PDBs and deleting PDBs. In particular, it contains the following:

- [Unplugging and Dropping a Pluggable Database Using Enterprise Manager](#)
- [Deleting Pluggable Databases Using Enterprise Manager](#)

Unplugging and Dropping a Pluggable Database Using Enterprise Manager

This section provides information about unplugging and dropping a PDB using Cloud Control. In particular, it contains the following:

- [Prerequisites for Unplugging and Dropping a Pluggable Database](#)
- [Unplugging and Dropping a Pluggable Database](#)



Note:

As an alternative to using the method described in this section, you can use EM CLI to unplug and drop PDBs. For more information, see [Unplugging and Dropping a Pluggable Database](#).

Prerequisites for Unplugging and Dropping a Pluggable Database

Before unplugging and dropping a PDB using Cloud Control, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The PDB that you want to unplug and drop must have been opened at least once.
- The target host user must be the owner of the Oracle home that the CDB (containing the PDB that you want to unplug and drop) belongs to.

Unplugging and Dropping a Pluggable Database

To unplug a PDB from its CDB using Cloud Control, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.



Note:

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, in the CDB section, select the CDB from which you want to unplug the PDBs.



Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section, select **Unplug Pluggable Database**.
4. Click **Launch**.

[Figure 14-35](#) displays the Provision Pluggable Database Console.

Figure 14-35 Provision Pluggable Database Console

Container Database
Specify a CDB target on which the PDB operation needs to be performed. You can create a CDB using the Create Oracle Database deployment procedure.

Container Database:

PDB Operations

- ☐ Migrate Existing Databases
Migrates non-CDBs as new Pluggable Databases
- ☐ Create New Pluggable Databases
Creates PDBs from sources such as the seed Pluggable Database, unplugged Pluggable Databases, or by cloning an existing Pluggable Database
- ☒ Unplug Pluggable Databases
Unplugs and drops the Pluggable Database after retaining the datafiles and PDB template which can later be used for plugging back the PDB
- ☐ Delete Pluggable Databases
Drops the Pluggable Database along with the datafiles



Note:

You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using *sysdba* user account credentials.

5. In the Select PDB page of the Unplug Pluggable Database Wizard, in the Select Pluggable Database section, select the PDB you want to unplug. Note that the PDB once unplugged will be stopped and dropped.
6. In the CDB Host Credentials section, select or specify the target CDB Oracle Home owner host credentials. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-36 displays the Select PDB page.

Figure 14-36 Unplugging a Pluggable Database: Select PDB Page

cdbam (Container Database)

Select PDB Destination Schedule Review

Unplug Pluggable Database : Select PDB

Select Pluggable Database

Select the PDB you want to unplug. Selected PDB will be unplugged and dropped from the CDB.

Pluggable Database:

Container Database Host Credentials

Specify host credentials. Host credentials are required to perform validations and initiate PDB unplug on the CDB host or cluster.

Credential: ☒ Preferred ☐ Named ☐ New

Preferred Credential Name:

Credential Details: Default preferred credentials are not set.

7. In the Destination page, select the type of PDB template you want to generate for unplugging the PDB, and the location where you want to store it. The PDB template consists of all datafiles as well as the metadata XML file.
 - If you want to store the PDB template on your CDB host (*CDB from where you are unplugging the PDB*), then select **Target Host File System**.
 - If you want to generate a single archive file—a TAR file with the datafiles and the metadata XML file included in it, then select **Generate PDB Archive**. Select a location where the archive file can be created.

 Note:

Oracle recommends you to select this option if the source and target CDBs are using file system for storage. This option is not supported for PDBs using ASM as storage.

- If you want to generate an archive file set—a separate DFB file with all the datafiles and a separate metadata XML file, then select **Generate PDB File Set**. Select the locations where the DFB and XML files can be created.

 Note:

Oracle recommends you to select this option if the source and target CDBs are using ASM for storage.

- If you want to generate only a metadata XML file, leaving the datafiles in their current location, then select **Generate PDB Metadata File**. Select a location where the metadata XML file can be created.
- If you want to store the PDB template in Oracle Software Library (Software Library), then select **Software Library**.
 - If you want to generate a single archive file—a TAR file with the datafiles and the metadata XML file included in it, then select **Generate PDB Archive**. If you want to generate an archive file set—a separate DFB file with all the datafiles and a separate metadata XML file, then select **Generate PDB File Set**. If you want to generate only a metadata XML file, leaving the datafiles in their current location, then select **Generate PDB Metadata File**.

- Enter a unique PDB template name.

The template is created in the default location that has the following format:

Database Configuration/<db_release>/<platform>/Database Templates

For example,

Database Configuration/12.1.0.0.2/unix/Database Templates

- Enter a temporary location where the archive can be created by Enterprise Manager before it is uploaded to the Software Library.

Figure 14-37 displays the Destination page.

Figure 14-37 Unplugging a Pluggable Database: Destination Page

The screenshot shows the 'Unplug Pluggable Database : Destination' page. At the top, there's a breadcrumb trail: 'Select PDB' > 'Destination' > 'Schedule' > 'Review'. Below this, the title is 'Unplug Pluggable Database : Destination'. Under 'PDB Template Location', there's a note: 'Unplug operation generates a PDB Template, which can be a PDB archive, PDB file set or PDB Metadata file. You can choose file system or Software Library to store PDB template. Using Software Library allows plugging in PDBs from a central location.' There are two radio buttons: 'Target Host File System' (selected) and 'Software Library'. Below this, there are four sections for generating different PDB components, each with a radio button and a text field for the location:

- Generate PDB Archive:** PDB archive is a compressed TAR file which consists of PDB XML metadata file and all datafiles that belong to PDB. Not supported for PDB using ASM as storage. PDB Archive Location: /scratch/user01/app/user01/product/12.1.0/ab025/assistant/db
- Generate PDB File Set:** PDB file set consists of PDB XML metadata file and RMAN backup of PDB. Recommended choice for PDBs using ASM as storage. PDB Metadata File: /scratch/user01/app/user01/product/12.1.0/ab025/assistant/db
- Generate PDB Metadata File:** Use this option to generate PDB metadata file and leave the PDB datafiles in the current location. Recommended for large PDB where packaging datafiles as TAR or RMAN backup is not efficient. PDB Metadata File: /scratch/user01/app/user01/product/12.1.0/ab025/assistant/db

8. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. Otherwise, select **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, and set the permissible delay time.

Figure 14-38 displays the Schedule page.

Figure 14-38 Unplugging a Pluggable Database: Schedule Page


The screenshot shows the 'Unplug Pluggable Database : Schedule' page. At the top, there's a breadcrumb trail: 'Select PDB' > 'Destination' > 'Schedule' > 'Review'. Below this, the title is 'Unplug Pluggable Database : Schedule'. There's a 'Deployment Instance' text field with the value 'UnplugPluggableDatabase_1381308454'. Below this, there's a 'Schedule' section. It has a 'Start' section with two radio buttons: 'Immediately' (selected) and 'Later'. To the right of these is a time zone dropdown showing '(GMT-08:00) Los Angeles - Pacific Time (PT)'. Below the 'Start' section is a 'Grace Period' section with a checkbox 'Do not run if it cannot start within' followed by a spinner set to '1' and a dropdown for 'hours', and then 'of the scheduled start time'.

9. Click **Next**.
10. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 14-39 displays the Review page.

Figure 14-39 Unplugging a Pluggable Database: Review Page



11. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

When you unplug and drop a PDB, the Enterprise Manager job system creates an Unplug Pluggable Database job. For information about viewing the details of this job, see [Viewing Unplug Pluggable Database Job Details](#).

Deleting Pluggable Databases Using Enterprise Manager

This section provides information about permanently deleting PDBs from a CDB, using Cloud Control. In particular, it contains the following:

- [Prerequisites for Deleting Pluggable Databases](#)
- [Deleting Pluggable Databases](#)

Prerequisites for Deleting Pluggable Databases

Before permanently deleting a set of PDBs from a CDB using Cloud Control, ensure that you meet the following prerequisites:

- The 12.1.0.5 Enterprise Manager for Oracle Database plug-in must be downloaded and deployed.
For information on how to download and deploy a plug-in, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
- Oracle Software Library (Software Library) must be set up in Cloud Control.
For information on how to set up Software Library in Cloud Control, see [Setting Up Oracle Software Library](#).
- The PDBs that you want to delete must have been opened at least once.
- The target host user must be the owner of the Oracle home that the CDB (containing the PDBs that you want to delete) belongs to.

Deleting Pluggable Databases

To permanently delete a set of PDBs from a CDB using Cloud Control, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.



Note:

You can also access the Provision Pluggable Database Console from the home page of the CDB. To do so, in the CDB's home page, from the **Oracle Database** menu, select **Provisioning**, then select **Provision Pluggable Database**.

2. In the Provision Pluggable Database Console, in the CDB section, select the CDB from which you want to delete the PDBs.




Note:

Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

3. In the PDB Operations section, select **Delete Pluggable Databases**.
4. Click **Launch**.

Figure 14-40 displays the Provision Pluggable Database Console.

Figure 14-40 Provision Pluggable Database Console



Note:

You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using **sysdba** user account credentials.

5. In the Select PDBs page of the Delete Pluggable Databases Wizard, click **Add**. Select the PDBs that you want to delete, then click **Select**.

 Note:


If you choose to delete a PDB that was created using the Snap Clone method, the PDB mount points on the CDB host are cleaned up. The corresponding storage volumes on the storage server are also deleted. This action is irreversible.

6. In the CDB Host Credentials section, select or specify the target CDB Oracle Home owner host credentials. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

If one (or more) of the PDBs that you selected for deletion is the Snap Clone of another PDB, you must also provide the privileged host credentials, that is, the credentials of the *root* user. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 14-41 displays the CDB Host Credentials section of the Select PDBs page.

Figure 14-41 Deleting Pluggable Databases: Specifying Credentials



7. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

If you want to run the procedure immediately, then retain the default selection, that is, **Immediately**. Otherwise, select **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, and set the permissible delay time.

Figure 14-42 displays the Schedule page.

Figure 14-42 Deleting Pluggable Databases: Schedule Page

cdbam (Container Database)

Select PDBs **Schedule** Review

Delete Pluggable Databases : Schedule

Deployment Instance: DeletePluggableDatabase_1381309209

Schedule

Start: ☒ Immediately ☐ Later (GMT-08:00) Los Angeles - Pacific Time (PT)

Grace Period: ☐ Do not run if it cannot start within 1 hours of the scheduled start time

8. Click **Next**.
9. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 14-43 displays the Review page.

Figure 14-43 Deleting Pluggable Databases: Review Page

cdbam (Container Database)

Select PDBs Schedule **Review**

Delete Pluggable Databases : Review

General

Container Database Name: cdbam

Host Name: abc.example.com

PDB for Delete: USER01_MAST


Delete Pluggable Databases

Target Name	PDB Name	Datafiles
cdbam_USER01_MAST	USER01_MAST	Show

10. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

Figure 14-44 displays the Procedure Activity page.

Figure 14-44 Deleting Pluggable Databases: Procedure Activity Page



When you delete a PDB, the Enterprise Manager job system creates a Delete Pluggable Database job. For information about viewing the details of this job, see [Viewing Delete Pluggable Database Job Details](#).

Viewing Pluggable Database Job Details Using Enterprise Manager

This section provides information about viewing the details of the jobs that are created by the Enterprise Manager job system when you create a PDB, unplug a PDB, or delete a PDB. It contains the following:

- [Viewing Create Pluggable Database Job Details](#)
- [Viewing Unplug Pluggable Database Job Details](#)
- [Viewing Delete Pluggable Database Job Details](#)

Viewing Create Pluggable Database Job Details

To view the details of a create PDB job, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. Click the deployment procedure that contains the required create PDB job.

● CreatePluggableDatabase_1381260390	Succeeded	Pluggable Databases Creation	Database Creation	SYSMAN
● SMF_SRL_1381253068-70_sicnas301	Succeeded	Snapshot Reclaimable Size	Storage Management Framework	TESTSUPERADMIN
● SMF_MDU_1381154586-426_sicnas301	Succeeded	Upload Storage Metadata	Storage Management Framework	TESTSUPERADMIN
● SMF_SRL_1381242026608_sicnas301	Succeeded	Snapshot Reclaimable Size	Storage Management Framework	TESTSUPERADMIN

3. Expand the deployment procedure steps. Select the PDB creation job.

Procedure Activity > CreatePluggableDatabase_1381260390

Procedure Activity: CreatePluggableDatabase_1381260390

> Elapsed Time: 2 minutes, 5 seconds

Procedure Steps

View Show All Steps

Select	Name	Type	Status
☐	Initialization	Computation	
☐	Snapshot Preparation	Computation	
☐	Source PDB Snapshot Creation	Dynamic Proc	
☐	Post Snapshot Creation	Computation	
☐	Clone Data Preparation	Computation	
☐	Target Storage Data Preparation	Computation	
☐	Pre PDB Clone Steps Execution	Rolling	
☐	Pluggable Databases Creation	Parallel	
☐	racdb	Cluster Data	
☒	Pluggable Database Creation	Job	
☐	PDB Storage Details Updation	Dynamic Proc	
☐	Post PDB Clone Steps Execution	Computation	

4. Click **Job Summary**.

Pluggable Database Creation x

Pluggable Database Creation Job Summary

Type Job Start Date Oct 8, 2013 12:34:58 PM PDT
Elapsed Time 1 minutes, 46 seconds Completed Date Oct 8, 2013 12:36:44 PM PDT

Step: Prepare Configuration Data (Succeeded)

Start Date Oct 8, 2013 12:34:58 PM PDT
Completed Date Oct 8, 2013 12:36:44 PM PDT
Targets racdb

Successfully created the temporary staging location(s).
[./scratch/aime/app/aime/e82fa81d4e0348fe040f00aceb157e8/incident]
Config data preparation complete.

5. To view a summary of the job details, click **Summary**.

In the Prepare Configuration Data step, the system prepares for PDB creation.

In the Check Prerequisites step, the system checks the prerequisites for PDB creation.

In the Verify and Prepare step, the system runs tasks prior to PDB creation.

In the Perform Configuration step, the PDB creation is performed. For details of the performed tasks and their status, refer to the remote log files present on the host.

In the Post Configuration step, Enterprise Manager is updated with the newly created PDB details, and the custom scripts are run.


6. To view a visual representation of the create PDB job progress, click **Results**.

In the Configuration Progress section, you can view the completion percentage of the job, and a list of pending, currently running, and completed job steps. You can also view errors, warnings, and logs. The tail of the log for the currently running job step is displayed.


Viewing Unplug Pluggable Database Job Details

To view the details of an unplug PDB job, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. Click the deployment procedure that contains the required unplug PDB job.



3. Expand the deployment procedure steps. Select the unplug PDB job.



4. Click **Job Summary**.
5. To view a summary of the job details, click **Summary**.

In the Prepare Configuration Data step, the system prepares for unplugging a PDB.

In the Check Prerequisites step, the system checks the prerequisites for unplugging a PDB.

In the Verify and Prepare step, the system runs tasks prior to unplugging the PDB.

In the Perform Configuration step, the PDB unplugging is performed. For details of the performed tasks and their status, refer to the remote log files present on the host.

In the Post Configuration step, Enterprise Manager is updated with the unplugged PDB details.

- To view a visual representation of the unplug PDB job progress, click **Results**.

In the Configuration Progress section, you can view the completion percentage of the job, and a list of pending, currently running, and completed job steps. You can also view errors, warnings, and logs. The tail of the log for the currently running job step is displayed.

Viewing Delete Pluggable Database Job Details

To view the details of a delete PDB job, follow these steps:

- From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
- Click the deployment procedure that contains the required delete PDB job.

☒ DeletePluggableDatabase_1381261461	Succeeded	Delete Pluggable Database Deployment Procedure	Database Creation
☐ CreatePluggableDatabase_1381260390	Succeeded	Pluggable Databases Creation	Database Creation
☐ SMF_SRI_1381253066750_slonas301	Succeeded	Snapshot Reclaimable Size	Storage Management Framework
☐ SMF_MDU_1381154586426_slonas301	Succeeded	Upload Storage Metadata	Storage Management Framework

- Expand the deployment procedure steps. Select the delete PDB job.

Provisioning

Procedure Activity > DeletePluggableDatabase_1381261461

Procedure Activity: DeletePluggableDatabase_1381261461

> Elapsed Time: 50 seconds

Procedure Steps

View Show All Steps

Select	Name	Type	Status
☐	Initialization Step	Computation	
☐	Snap Clone PDB Delete Preparation step	Computation	
☒	▼ Delete Pluggable Databases	Parallel	
☐	▼ racdb	Cluster Data	
☒	Delete Pluggable Database Job Step	Job	
☐	Updates storage details for the PDB	Dynamic Proc	
☐	Cleanup Operations	Computation	

- Click **Job Summary**.

Delete Pluggable Databases > Delete Pluggable Database Job Step

Delete Pluggable Database Job Step

Elapsed Time: 35 seconds

Type: Job Start Date: Oct 8, 2013 12:44:55 PM PDT
Completed Date: Oct 8, 2013 12:45:30 PM PDT

Step: Prepare Configuration Data (Succeeded)

Start Date: Oct 8, 2013 12:44:55 PM PDT
Completed Date: Oct 8, 2013 12:45:30 PM PDT
Targets: racdb

[Job Summary](#)

5. To view a summary of the job details, click **Summary**.

In the Prepare Configuration Data step, the system prepares for deleting the PDBs.

In the Verify and Prepare step, the system runs tasks prior to deleting the PDBs.

In the Perform Configuration step, the PDB deletion is performed. For details of the performed tasks and their status, refer to the remote log files present on the host.

In the Post Configuration step, Enterprise Manager is updated with the deleted PDB details.

6. To view a visual representation of the delete PDB job progress, click **Results**.

In the Configuration Progress section, you can view the completion percentage of the job, and a list of pending, currently running, and completed job steps. You can also view errors, warnings, and logs. The tail of the log for the currently running job step is displayed.

Administering Pluggable Databases Using Enterprise Manager

This section provides information about performing PDB administration tasks using Cloud Control. It contains the following:

- [Switching Between Pluggable Databases Using Enterprise Manager](#)
- [Altering Pluggable Database State Using Enterprise Manager](#)

Switching Between Pluggable Databases Using Enterprise Manager

If you are performing a task such as granting user privileges or performance reporting, and you need to perform the same task on another PDB, then you can switch to another PDB using Cloud Control. To switch between PDBs while staying on the same feature page, follow these steps:

1. From the current PDB, select any PDB scope page (such as, Manage Advanced Queues).

In the upper-left corner of the window, the name of the PDB will update to display a context switcher as a drop-down menu.
2. Click the context switcher to display the drop-down menu. This menu shows the PDBs most recently used.

Select a PDB from this list.
3. The page will update to show the System Queues.
4. Click the context switcher to display the drop-down menu. If the menu does not show the PDBs you want, then select **All Containers**.
5. A Switch Container window will pop up to display all available PDBs for the monitored target.

Select a PDB from this list and click **OK**.
6. The page will update to show data for the selected PDB.

Altering Pluggable Database State Using Enterprise Manager

To change the state of a single-instance PDB to Open or Close using Cloud Control, follow these steps:

1. From the **Oracle Database** menu, select **Control**, then select **Open/Close Pluggable Database**.
2. From the Open/Close PDB page, select a PDB from the list.
3. Click the **Action** drop-down menu and select the appropriate actions. Your choices are **Open**, **Open Read Only**, and **Close**.
4. In the Confirmation dialog window, click **Yes** to complete the change. A Processing dialog window appears to show you the progress of your choice.
5. Once state change completes, the Open/Close PDB page will update to show the new state of the PDB.

To change the state of a PDB in a Cluster/RAC to Open or Close, follow these steps:

1. From the **Oracle Database** menu, select **Control**, then **Open/Close Pluggable Database**.
2. From the Open/Close PDB page, select a PDB from the list. The RAC instances are shown along with the PDB's current state on those instances.
3. Once you select a PDB, a panel appears below the list to show the state of the PDBs on the different RAC instances. The open and close options apply to the PDBs on the RAC instance's panel. You can open or close a PDB on any number of available RAC instances.
4. In the Confirmation dialog window, click **Yes** to complete the change. A Processing dialog window appears to show you the progress of your choice.
5. Once state change completes, the Open/Close PDB page will update to show the new state of the PDB.

Cloning Oracle Databases and Pluggable Databases

Enterprise Manager Cloud Control enables you to clone databases using the Full Clone method, or by using the classic cloning wizard which enables you clone databases using RMAN backup, staging areas, or an existing backup.

This chapter outlines the following procedures which you can use to create a database clone. For other advanced cloning methods, see *Enterprise Manager Cloud Administration Guide*.

- [Creating a Full Clone Database](#)
- [Creating a Full Clone Pluggable Database](#)
- [Cloning Databases Using the Classic Cloning Wizard](#)

Creating a Full Clone Database

To create a Full Clone database, you can use either of the following solutions:

- [Creating a Full Clone Database Using the Clone Wizard](#)
- [Creating a Full Clone Database Using EM CLI](#)

Creating a Full Clone Database Using the Clone Wizard

To create a full clone database, follow these steps:

1. On the Databases page, you can access the Full Clone database wizard by following any one method:
 - Select the database that you want to clone from the list of the databases displayed. On the Database home page, click the **Database** menu, select **Cloning**, and then select **Create Full Clone**.
 - Right click on the database target name, select **Database**, select **Cloning**, and then select **Create Full Clone**.
 - Right click on the database target name, select **Database**, select **Cloning**, and then select **Clone Management**. On the Clone Management page, in the Full Clone Databases box, click **Create**.
2. On the Create Full Clone Database: Source and Destination page, do the following:
 - In the Source section, launch the credentials selector by selecting the search icons for SYSDBA Database and Database Host credentials. Click **OK**.
 - In the Data Time Series section, select **Now** or **Prior Point in Time**.

If you selected Now, specify or search and select the SYSASM ASM Credentials. Now refers to Live Clone.

 Note:

- Live clone (RMAN duplicate) is supported only if the source is in ARCHIVELOG mode.
- For a Container database only **Now** option (live clone) in data-time series is supported.

If you selected Prior Point in Time, a carousel of RMAN Backup images appear. Select the appropriate RMAN backup by clicking **Select** on the image.

You can create full clones by selecting a backup and optionally modify the time and SCN to do a point in time restore. The **Select Time** option has the minimum limit set to the current backups time and maximum time limit set to the next backup time. You can modify this in case you have to create a new clone between these two time periods. Similarly, you can do the same for SCN by selecting the **Select SCN** option.

3. In the Destination Database Definition section, do the following:

- Specify a display name.
- Specify a global database name and SID.

A database is uniquely identified by a Global Database Name. The typical form of a name is name.domain. A database is referenced by at least one Oracle instance which is uniquely identified by SID.

- Select one of the following types of databases:

- **Single Instance Database.**

In the Hosts section, specify the Oracle Home location. The host gets specified by default. Next, select the Database Home credentials and the SYSASM ASM credentials.

- **RAC Database**

In the Hosts section, specify or select the cluster target. The Oracle Home location gets specified by default. Next, specify the Database Host credentials, and the SYSASM ASM credentials.

In the Nodes section, select the cluster and Oracle Home to display one or more hosts on which the administrator managed Oracle RAC database will be created.

- **RAC One Node Database**

In the Hosts section, specify or select the cluster target. The Oracle Home location gets specified by default. Next, specify the Database Host credentials, and the SYSASM ASM credentials.

In the Nodes section, select the cluster and Oracle Home to display one or more hosts on which the administrator managed Oracle RAC database will be created.

 Note:

Oracle supports inline patching as part of clones. When the destination home selected has patches applied such as the latest CPU or PSU, then the cloned database is automatically brought up with that level.

Click **Next**.

4. On the Create Full Clone Database: Configuration page, do the following:

- In the Database Files Location, specify the location where you want the data files, temp files, redo log files, and control files to be created. You can select **File System** or **Automatic Storage Management (ASM)**, and then specify the common location for the database files.

The **Use Oracle Optimal Flexible Architecture-compliant directory structure (OFA)** option enables you to configure different locations for:

- Archive and Redo logs
- Data files
- Control file
- Temporary file
- In the Recovery Files location, specify the location where you want the recovery files, such as archived redo logs, RMAN backups, and other related files to be created. You can choose to use the fast area recovery by selecting **Use Fast Recovery Area**. If you do, specify the fast recovery area size. The fast recovery area size is defaulted to that of source.
- In the Listener Configuration section, select the listener targets running under the new Oracle Home, to register the clone database.
- In the Database Credentials section, specify passwords for the SYS, SYSTEM, and DBSNMP administrative users in the clone database. You can choose to have the same password for all the three users or a different password for each.
- In the Parallelism section, enter the number of parallel channels used by Recovery Manager (RMAN) to copy the database files in the **Degree of Parallelism** text box.

Degree of parallelism is the number of parallel channels used by Recovery Manager (RMAN) to copy the database files. It is calculated as the CPU count of the host. Increased parallelism may speed the process if sufficient system resources are available.

Click **Next**.

5. On the Create Full Clone Database: Initialization Parameters page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click **Edit** to modify the value of the parameter.

Click **Next**.

6. On the Create Full Clone Database: Post Processing page, specify the following:

- **Masking Definition:** Data masking is the process of masking sensitive data in test or non-production databases. The masking definition defines the columns to be masked in the format of the masked data. Specify the masking definition to be applied after the database is cloned.

 Note:

The masking definition can be used only when you have a Subset-Masking license pack.

- **Custom Scripts:** Specify the custom scripts that need to be executed before and after the database is created.

For more information on how to store and use custom scripts in the Software Library, refer to [Using Custom Scripts Stored in the Software Library](#). To know the formats of custom scripts, you can view the sample scripts in the Database as a Service/Samples/Pre or Post Database Script folder in the Software Library. To navigate to the Software Library, from **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library**.

- **Create Data Profile:** This option enables you to automatically take a backup of the new cloned instance once it is created. When the clone or the Test Master is refreshed, this section displays the existing profiles created for the database. You can select the profile that has to be refreshed along with the database.
- **Create as Test Master:** This option if you want to create the cloned database as a Test Master database.

Click **Next**.

7. On the Create Full Clone Database: Schedule page, specify a unique deployment procedure instance name. You can choose to start the deployment procedure immediately or at a later time,

In the Notification Details section, you can choose to set the following notifications:

- Scheduled
- Running
- Action Required
- Suspended
- Succeeded
- Problems

Click **Next**.

8. On the Create Full Clone Database: Review page, verify the details of the source database, the data source of the clone, and the destination database.

Click **Submit**.

Creating a Full Clone Database Using EM CLI

To create a full clone of a database execute the verb `emcli db_clone_management -createFullClone -input_file="location of file containing properties for creating the database clone"`.

Sample properties file:

```
CLONE_TYPE=DUPLICATE
SRC_DB_TARGET_NAME=xyy.us.example.com
SRC_DB_TARGET_TYPE=oracle_database
SRC_DB_CRED=NC_DB_CRED:SYSCO
SRC_HOST_NORMAL_NAMED_CRED=NC_HOST_CRED:SYSCO
COMMON_GLOBAL_DB_NAME=clonedb.example.com
COMMON_DB_SID=clonedb
DB_TARGET_NAME=clonedb.xyy.example.com
DATABASE_TYPE=dbTypeSI
TARGET_HOST_LIST=desthost.example.com
ORACLE_HOME_NAME=Oradb12Home2_29
ORACLE_HOME_LOC=/scratch/app/product/11.2.0./dbhome_1
ORACLE_BASE_LOC=/scratch/app/base
HOST_NORMAL_NAMED_CRED=NC_HOST_CRED2:SYSCO
DB_STORAGE_TYPE=FS
DB_FILE_LOC=/scratch/app/oradata
FRA_STORAGE_TYPE=FS
FLASH_REC_AREA=/scratch/user/app/fra
FRA_SIZE=4395
ARCHIVE_LOG_MODE=NO
DEST_LISTENER_SELECTION=DEST_DB_HOME
LISTENER_PORT=1526
ENABLE_LIVE_CLONE=true
DB_ADMIN_PASSWORD_SAME=true
DATABASE_PASSWORDS=right1
DB_TEMPLATE_STAGE=/tmp
```

Note:

The sample file shown above may not be valid for different types of full clone. As in, you may need to have different parameters for a live clone, or a point-in-time clone. Similarly, for a Single Instance, or a RAC clone. To know the necessary parameters to be added in the input file for a particular type of clone, enter the command `emcli db_clone_management -print_properties`.

To verify the status of the database clone creation, execute the verb `emcli get_instance_status -instance={instance GUID}`.

Refresh a Cloned Database

Once the Enterprise Manager Plugins are deployed, the 'Refresh' option is enabled. To deploy the plugins perform the following steps:

1. Go to **Setup** and select **Extensibility**, and then click **Plugins**.

2. From **Servers, Storage and Network** menu, deploy the **Oracle Virtualization** plugin on the Management Server and Management Agents.
3. From the **Cloud** menu, deploy the **Oracle Cloud Application** plugin on the Management Server.

When these plugins are successfully deployed you can refresh a cloned database by performing the following steps:

1. Navigate to Oracle Database and click **Cloning**.
2. Click **Cloning Management**, and then select a cloned database from the list.
3. Click **Refresh**.

Creating a Full Clone Pluggable Database

To create a full clone PDB, you can use either of the following solutions:

- [Creating a Full Clone Pluggable Database Using the Clone Wizard](#)
- [Creating a Full Clone Pluggable Database Using EM CLI](#)

Creating a Full Clone Pluggable Database Using the Clone Wizard

You can create a full clone of a PDB using the new Clone PDB Wizard. To create a full clone PDB, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source CDB (the CDB that the source PDB is a part of) in the list, then click the name of the PDB that you want to clone.
4. From the **Oracle Database** menu, select **Cloning**, then select **Create Full Clone**.

Alternatively, in Step 3, you can right click the name of the PDB that you want to clone, select **Oracle Database**, select **Cloning**, then select **Create Full Clone**.

5. On the Source and Destination: Create Full Clone Pluggable Database page, do the following:
 - Specify the SYSDBA credentials for the source CDB. You can choose to use the preferred credentials, use a saved set of named credentials, or specify a new set of credentials.
 - In the Pluggable Database Destination section, specify a name, and a display name for the PDB clone. Enterprise Manager uses the display name to identify the PDB clone target.
 - Select the Refreshable clone check box to create a manually Refreshable PDB.

 Note:

- Refreshable clone feature is only applicable to Oracle Database version 12.2 and above.
- To create a Refreshable clone you have to select a different Container Database (CDB). You will not be able to create one in the same CDB.
- The Refreshable PDB clone is created in read-only mode and can only be either in read-only or mounted state. The clone cannot be opened in read-write mode. Therefore, data masking and post SQL scripts cannot be run on the clone.

- In the PDB Administrator Credentials section, specify the credentials of the admin user account that you want to use to administer the PDB clone.
- To clone the PDB to a CDB different from the source CDB, select **Clone the Pluggable Database into a different Container Database**, then specify the destination CDB.
- In the Credentials section, specify the destination CDB host credentials. If you chose to clone the PDB to a CDB different from the source CDB, specify the SYSDBA credentials for the destination CDB. Also, if the destination CDB is using Automatic Storage Management (ASM) to manage disk storage, you must specify the ASM credentials.

 Note:

For Database versions 12.2 and above, if both the source and destination Container Databases are in Archive Log mode and have Local Undo configured, then a Hot clone is performed. Hot clone is the ability to clone a Pluggable Database while the source is open in read-write mode.

6. You can specify other configuration details, mask data, as well as schedule the cloning process, by selecting the **Advanced** option. Follow the rest of the steps, if you have selected the Advanced option. If you do not need to specify anymore details, click **Clone**. This submits the deployment procedure to clone a PDB to a CDB that is deployed in a public cloud setup.
7. On the Create Full Clone Pluggable Database: Source and Destination page, verify the details specified, and then click **Next**.
8. On the Create Full Clone Pluggable Database: Configuration page, do the following:
 - In the Database Files Location section, specify the storage location where the datafiles of the PDB clone must be stored. If the destination CDB is using ASM to manage disk storage, specify the disk group where the datafiles of the PDB clone must be stored.
 - To ensure that only the source PDB data model definition is cloned (and the source PDB data is not cloned), select **Exclude User Data**.

- In the Advanced Configuration section, specify the storage limits for the maximum size of the PDB clone, and the maximum size of a shared tablespace within the PDB clone. By default, no limits are placed on the values for these attributes.
- In the Miscellaneous section, select the logging option that you want to use for the tablespaces created within the PDB clone.

Click **Next**.

9. On the Create Full Clone Pluggable Database: Post Processing page, do the following:

- In the Data Masking section, specify the data masking definition that you want to apply after cloning the PDB. Data masking masks sensitive data in a database.

For information on how to create a data masking definition, see *Oracle Data Masking and Subsetting Guide*. Note that you can apply a data masking definition only if you have the Subset-Masking license pack.

- In the Custom Scripts section, for **Pre Script** and **Post Script**, specify the Oracle Software Library components that contain the scripts that you want to run before cloning, and after cloning the PDB respectively. Also, for **SQL Script**, specify the SQL scripts that you want to run after cloning the PDB. For **Run As User**, select the user account that you want to use to run the SQL scripts.

Click **Next**.

10. On the Create Full Clone Pluggable Database: Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive e-mail notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

11. On the Create Full Clone Pluggable Database: Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Clone** to submit the deployment procedure to create a full clone of the source PDB.

Creating a Full Clone Pluggable Database Using EM CLI

To create a full clone of a pluggable database, execute the verb `emcli pdb_clone_management -input_file=data:/xyz/sdf/pdb_clone.props`, where `pdb_clone.props` is the properties file.

Sample properties file (pdb_clone.props):

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLCO_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
```

```
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
```

 Note:

- If you are creating a full clone PDB on Oracle cloud, wherein both the destination PDB and the source PDB are in different CDBs, while both the CDBs are on Oracle Cloud, then ensure that the source PDB is in read-write mode. This is necessary since a database link is created in the destination CDB for cloning the PDB, and a temporary user is created in the source PDB for using the database link. If there is an existing database link in the destination CDB that connects to the source PDB, then use the parameter `EXISTING_DB_LINK_NAME` to provide the database link name in the properties file.
- If the PDB full clone (on-premise) needs to be performed using a backup, then the following parameters should be added in the properties file.

```
BACKUP_TYPE=RMAN
DEST_STAGE_DIR=/tmp/bkp_common
```

If the temporary backup location is shared across source and destination, then add the `SKIP_DATA_TRANSFER` parameter and set it to 'Y'. For example, `SKIP_DATA_TRANSFER=Y`. This indicates that the data transfer step should be skipped in the clone procedure.

Cloning Databases Using the Classic Cloning Wizard

You can clone databases using the older cloning wizard. This section consists of the following:

- [Overview of Classic Cloning Methods](#)
- [Cloning an Oracle Database Using Recovery Manager \(RMAN\) Backup](#)
- [Cloning an Oracle Database Using Staging Areas](#)
- [Cloning an Oracle Database Using an Existing Backup](#)

Overview of Classic Cloning Methods

You can use the Enterprise Manager Clone Database wizard to clone an Oracle database instance to an existing Oracle home. After you have an Oracle database instance in a known state (for example, you've configured it, tuned it, and tested it), you may want to clone that database to another existing Oracle home.

The following table lists the cloning methods and their cloning process:

Table 15-1 Oracle Database Cloning Methods

Cloning Method	Cloning Process
Cloning an Oracle Database Using Recovery Manager (RMAN) Backup	• Connects source and destination Oracle instances • Copies database files using RMAN duplicate feature • Recovers and opens the cloned database
Cloning an Oracle Database Using Staging Areas	• Backs up each database file and stores it in a staging area • Transfers each backup file from source to destination • Restores each backup file to the specified locations • Recovers and opens the cloned database
Cloning an Oracle Database Using an Existing Backup	• Creates cloned database as of specified point-in-time or SCN • Validates backups prior to the clone operation • Transfers required archived redo log files to destination host • Recovers and opens the cloned database

Cloning an Oracle Database Using Recovery Manager (RMAN) Backup

To clone an Oracle database using RMAN backup, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, select a database that you want to clone.
3. On the Database target page, from the **Oracle Database** menu, select **Provisioning**, and then click **Clone and Refresh Database**.
4. On the Clone and Refresh page, click the **Switch to Classic Clone** link.
5. On the Database Login page, enter your credentials. Click **Login**.

The Clone Database wizard opens.

6. On the Clone Database page: Source Type page, select **Online Backup** and **Use Recovery Manager (RMAN) to copy database files**.

Click **Continue**.

Note:

When you use RMAN backup to clone a database, the source database will be duplicated directly to the specified destination Oracle Home. No staging areas are required.

7. On the Clone Database: Source Options page, in the Degree of Parallels box, enter the number of parallel channels to be used by RMAN to copy the database files. The default number is 2.



Note:

Degree of parallelism is the number of parallel channels used by Recovery Manager (RMAN) to copy the database files. It is calculated as the CPU count of the host. Increased parallelism may speed the process if sufficient system resources are available.

8. In the Source Host Credentials section, enter the credentials of the user who owns the source database Oracle server installation. You can either select **Named** credential or **New** credential.

If you select **New** credential, enter the Username and Password. You can select the **Set as Preferred Credentials** checkbox, if you want to use these set of credentials again. Click **Test** to check if your credentials are valid.

Click **Next**.

9. On the Clone Database: Select Destinations page, in the Destination Oracle Home section, click **Search**.



Note:

The Oracle Home should exist on the specified host and should match the version of the source database.

On the Destination Oracle Home page that appears, search and select the destination Oracle Home. Click **Next**.

10. In the Destination Host Credentials section, enter the credentials of the user who owns the Oracle Home specified in the Destination Oracle Home section.
11. In the Destination Database section, do the following: specify the global database name, the instance name, and for select **file system** as the database storage. Click **Next**.

- Specify the global database name.
For example: clone1.example.com
- Specify a unique instance name.
For example: clone1
- Select **File System** as the database storage.
- Click **Next**.

12. On the Clone Database: Destination Options page, select **Use Database Area and Fast Recovery Area**.

Click **Next**.

13. On the Clone Database: Database Configuration page, in the Listener Configuration section, specify the name and port of the listener that will be used for the cloned database. If a new name and port are specified that are not in use by an exiting listener, a new listener using the specified port will be created.

In the Database Registration section, select Register the cloned database as an Enterprise Manager target monitored by using DBSNMP. Enter the target database name.

Click **Next**.

14. On the Clone Database: Schedule page, specify a name description for the clone job. You can choose to run the clone job immediately or you can specify a later time and date for the job to run.

Click **Next**.

15. On the Clone Database: Review page, review the details and configuration of the source database, the destination database, and the database storage. You can view the database storage files by clicking on **View Source Database Files**.

Click **Submit Job**. T

he clone database job is now submitted. When the job completes, a Clone Database: Confirmation page appears. To view the status of the job, click **View Status**. To exit the page click **OK**.

Cloning an Oracle Database Using Staging Areas

To clone an Oracle database by copying database files via staging areas, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, select a database that you want to clone.
3. On the Database target page, from the **Oracle Database** menu, select **Provisioning**, and then click **Clone Database**.
4. On the Clone and Refresh page, click the **Switch to Classic Clone** link.
5. On the Database Login page, enter your credentials. Click **Login**.

The Clone Database wizard opens.

6. On the Clone Database page: Source Type page, select **Online Backup** and **Copy database files via staging areas**.

Click **Continue**.

Note:

This method requires staging areas on both the source and the destination hosts.

7. On the Clone Database: Source Options page, in the Staging Area section, enter the Staging Area Location.



Note:

A staging area on the source host is required in order to clone a running database. A backup is performed on the database and the backup files are stored in the staging area.

8. Select if you want to delete or retain the staging area after the cloning operation.
By retaining the staging area after a cloning operation, you avoid doing another backup later. However, this option requires a minimum disk space of 2230 MB.
9. In the Source Host Credentials section, enter the credentials of the user who owns the source database Oracle server installation. You can either select **Named** credential or **New** credential.
If you select **New** credential, enter the Username and Password. You can select the **Set as Preferred Credentials** checkbox, if you want to use these set of credentials again. Click **Test** to check if your credentials are valid.
Click **Next**.
10. On the Clone Database: Select Destinations page, in the Destination Oracle Home section, click the Search icon.



Note:

The Oracle Home should exist on the specified host and should match the version of the source database.

On the Destination Oracle Home page that appears, search and select the destination Oracle Home. Click **Next**.

11. In the Destination Host Credentials section, enter the credentials of the user who owns the Oracle Home specified in the Destination Oracle Home section.
12. In the Destination Database section, do the following: specify the global database name, the instance name, and for select **file system** as the database storage. Click **Next**.
 - Specify the global database name.
For example: clone1.example.com
 - Specify a unique instance name.
For example: clone1
 - Select **File System** as the database storage.
 - Click **Next**.
13. On the Clone Database: Destination Options page, select **Use Database Area and Fast Recovery Area**.
Click **Next**.
14. On the Clone Database: Database Configuration page, in the Listener Configuration section, specify the name and port of the listener that will be used for the

cloned database. If a new name and port are specified that are not in use by an existing listener, a new listener using the specified port will be created.

In the Database Registration section, select Register the cloned database as an Enterprise Manager target monitored by using DBSNMP. Enter the target database name.

Click **Next**.

15. On the Clone Database: Schedule page, specify a name description for the clone job. You can choose to run the clone job immediately or you can specify a later time and date for the job to run.

Click **Next**.

16. On the Clone Database: Review page, review the details and configuration of the source database, the destination database, and the database storage. You can view the database storage files by clicking on **View Source Database Files**.

Click **Submit Job**. The

clone database job is now submitted. When the job completes, a Clone Database: Confirmation page appears. To view the status of the job, click **View Status**. To exit the page click **OK**.

Cloning an Oracle Database Using an Existing Backup

To clone an Oracle database using an existing backup, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. On the Databases page, select a database that you want to clone.
3. On the Database target page, from the **Oracle Database** menu, select **Provisioning**, and then click **Clone Database**.
4. On the Clone and Refresh page, click the **Switch to Classic Clone** link.
5. On the Database Login page, enter your credentials. Click **Login**.

The Clone Database wizard opens.

6. On the Clone Database page: Source Type page, select **Existing BackUp**.
Click **Continue**.
7. On the Clone Database: Source Host Credentials page, select the backup that you want to use.
8. In the Source Host Credentials section, enter the credentials of the user who owns the source database Oracle server installation. You can either select **Preferred**, **Named** or **New** credential.

If you select **New** credential, enter the Username and Password. You can select the **Set as Preferred Credentials** checkbox, if you want to use these set of credentials again. Click **Test** to check if your credentials are valid.

Click **Next**.

9. On the Clone Database: Backup Details page, in the Point In Time section, specify a time or System Change Number (SCN). This will help identify backups necessary to create the clone database.

**Note:**

If the existing backup does not have all necessary archive logs, Enterprise Manager will transfer them from the source host to the destination host as part of the clone operation.

10. Oracle database backups can be encrypted using a database wallet, password, or both. If the backups are encrypted, specify the encryption mode and password as needed, in the Encryption section. By default, the encryption mode is set as **None**.

Click **Next**.

11. In the Destination Host Credentials section, enter the credentials of the user who owns the Oracle Home specified in the Destination Oracle Home section.
12. In the Destination Database section, do the following: specify the global database name, the instance name, and for select **file system** as the database storage. Click **Next**.
 - Specify the global database name.
For example: clone1.example.com
 - Specify a unique instance name.
For example: clone1
 - Select **File System** as the database storage.
13. In the Parallelism section, in the Degree of Parallels box, enter the number of parallel channels to be used by RMAN to copy the database files. The default number is 2.

**Note:**

Degree of parallelism is the number of parallel channels used by Recovery Manager (RMAN) to copy the database files. It is calculated as the CPU count of the host. Increased parallelism may speed the process if sufficient system resources are available.

Click **Next**.

14. On the Clone Database: Destination Database Settings page, in the Memory Parameters section, select **Configure Memory Management** and then from the drop-down list select **Automatic Shared Memory Management**.

The database automatically sets the optimal distribution of memory across the System Global Area (SGA) components. The distribution of memory will change from time to time to accommodate changes in the workload. Also, specify the aggregate Program Global Area (PGA) size.

15. In the Listener Configuration section, specify the name and port of the listener to be configured for the database. If the listener specified does not exist at the destination Oracle Home, it will be created.

 Note:

If you are going to convert the cloned database RAC at a later point, it is recommended that you specify storage location shared across all hosts in the cluster.

16. In the Recovery Files section, specify the location where recovery-related files such as, archived redo log files, RMAN backups, and the like are to be created.

Click **Next**.

17. On the Clone Database: Storage Locations page, in Database Files Location section, specify the location where datafiles, tempfiles, redo log files, and control files are to be created.

In the Database Registration section, select Register the cloned database as an Enterprise Manager target monitored by using DBSNMP. Enter the target database name.

Click **Next**.

18. On the Clone Database: Schedule page, specify a name description for the clone job. You can choose to run the clone job immediately or you can specify a later time and date for the job to run.

Click **Next**.

19. On the Clone Database: Review page, review the details and configuration of the source database, the destination database, and the database storage. You can view the database storage files by clicking on **View Source Database Files**.

Click **Submit Job**. T

The clone database job is now submitted. When the job completes, a Clone Database: Confirmation page appears. To view the status of the job, click **View Status**. To exit the page click **OK**.

Database High Availability REST API

This chapter provides details of the REST APIs for Database High Availability supported in Enterprise Manager.

- [Data Guard](#)
- [Backup and Recovery](#)

Data Guard

Data Guard Switchover REST API

Table 16-1 Data Guard Switchover

Feature	Description
URL	<code>https://<OMS_HOST>:<OMS_CONSOLE_PORT>/em/websvcs/restful/db/dataguard/switchover</code>
Request header	Authorization: basic Content-Type: application/json
Body	<pre>{ "primaryTargetName" : "<primary target name>", "primaryTargetType" : "oracle_database rac_database", "standbyTargetName" : "<standby target name>", "standbyTargetType" : "oracle_database rac_database", ["primaryDbCredsName" : "<primary database credential name>"], ["primaryHostCredsName" : "<primary database host cre- dential name>"], ["standbyDbCredsName" : "<standby database credential name>"], ["standbyHostCredsName" : "<standby database host cre- dential name>"], ["swapJobs" : "false true"], ["swapThresholds" : "false true"], ["tdewalletCredsName" : "<transparent data encryption wallet credentials of the primary database>"] ["tdewalletPassword" : "<transparent data encryption wallet password of the primary database>"] }</pre>
Request method	POST

Table 16-1 (Cont.) Data Guard Switchover

Feature	Description
Sample Body	<pre>{ "primaryTargetName" : "dbname", "primaryTargetType" : "oracle_database", "standbyTargetName" : "sdbname", "standbyTargetType" : "oracle_database", "primaryDbCredsName" : "SYS", "primaryHostCredsName" : "AIMEDB_N", "standbyDbCredsName" : "SYS", "standbyHostCredsName" : "AIMEDB_N", "swapJobs" : "true", "swapThresholds" : "true", "tdewalletCredsName" : "WC1" }</pre>
Sample Response	<pre>{ "status": "Succeeded", "submissionName": "dbname_DataguardSwitch- over_091418163542", "executionGUID": "75D3A124B36D1A4DE0536A8A7B0A80E0", "databaseName": "dbname", "message": "Deployment procedure for switchover to dbname:oracle_database successfully submitted." }</pre>
Supported Since Release	EM DB 13.3.2.0.0

**Note:**

Parameters listed within a square bracket are optional.

Table 16-2 Data Guard Switchover Parameters

Parameter	Description
<i>primaryTargetName</i>	Primary database Enterprise Manager target name. Can be either a single-instance database or a cluster database.
<i>primaryTargetType</i>	Primary database target type. Allowed values are either single-instance database (oracle_database) or cluster database (rac_database).
<i>standbyTargetName</i>	Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database.
<i>standbyTargetType</i>	Standby database target type. Allowed values are either single-instance database (oracle_database) or cluster database (rac_database).

Table 16-2 (Cont.) Data Guard Switchover Parameters

Parameter	Description
<i>primaryDbCredsName</i>	Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used.
<i>primaryHostCredsName</i>	Primary database host named credential for an operating system user who can access the primary database Oracle home. If this argument is not specified, preferred credential will be used.
<i>standbyDbCredsName</i>	Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used.
<i>standbyHostCredsName</i>	Standby database host named credential for an operating system user who can access the standby database Oracle home.
<i>swapJobs</i>	Specifies that scheduled jobs on either the primary or standby database that are of a transferable type (Backup, RMAN Script, and SQL Script) will be moved to the other database after the switchover is complete. Only jobs that use global named or preferred credentials can be transferred. Default Value: Not enabled.
<i>swapThresholds</i>	Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled
<i>tdeWalletCredsName</i>	Transparent Data Encryption wallet credentials for the primary database. Use create_named_credential verb to create these credentials as shown below: emcli create_named_credential -cred_name=WC1 -cred_type=GenericPassword -auth_target_type='<system>' -attributes="GENERIC_PASSWORD:<Primary Database TDE Wallet Password>"
<i>tdeWalletPassword</i>	Transparent Data Encryption wallet password for the primary database.

Convert Standby Database REST API**Table 16-3 Convert Standby Database**

Feature	Description
URL	https://<OMS_HOST>:<OMS_CONSOLE_PORT>/em/websvcs/restful/db/dataguard/convertStandby
Request header	Authorization: basic Content-Type: application/json

Table 16-3 (Cont.) Convert Standby Database

Feature	Description
Body	<pre>{ "dbTargetName" : "<standby target name>", "dbTargetType" : "oracle_database rac_database", ["dbHostCredName" : "<standby database host credential name>",] ["dbCredName" : "<standby database credential name>",] ["convertToDbRole" : "PHYSICAL STANDBY SNAPSHOT STAND- BY"] }</pre>
Sample Body	<pre>{ "dbTargetName": "s11gfss", "dbTargetType": "oracle_database", "dbHostCredName": "GIMAN_N", "dbCredName": "SYSDBA_CREDS", "convertToDbRole": "snapshot" }</pre>
Sample Response	<pre>{ "status": "Succeeded", "submissionName": "dbname_DGConvertStandby_092318020955", "executionGUID": "76815E6ABC9465F5E053C365F20A61AF", "databaseName": "dbname", "message": "Deployment procedure to convert the standby database successfully submitted." }</pre>
Request method	POST
Supported Since Re-lease	EM DB 13.3.2.0.0

Table 16-4 Convert Standby Database Parameters

Parameter	Description
<i>dbTargetName</i>	Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database.
<i>dbTargetType</i>	Standby database target type. Allowed values are either single-instance database (oracle_database) or cluster database (rac_database).
<i>dbHostCredName</i>	Standby database host named credential for an operating system user who can access the standby database Oracle home. If this argument is not specified, preferred credential will be used.
<i>dbCredName</i>	Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used.

Table 16-4 (Cont.) Convert Standby Database Parameters

Parameter	Description
<i>convertToDbRole</i>	Standby database role to be converted to given role. Possible values are Physical Standby and Snapshot Standby.

Backup and Recovery

Configure Backup REST API

Table 16-5 Configure Database to Send Backup

Feature	Description
URL	<code>https://<OMS_HOST>:<OMS_CONSOLE_PORT>/em/websvcs/restful/db/backup/configure</code>
Request header	Authorization: basic Content-Type: application/json
Body1	Configure database to be protected by a Recovery Appliance, with the ability to send backups to the Recovery Appliance.

```
{
  "dbTargetName":"<database target name>",
  "dbTargetType":"oracle_database|rac_database",
  "configureRABackup":"true"
  "raTargetName":"<Recovery Appliance target name>"
  "raUser":"<Recovery Appliance database user name>"
  ["dbCredName":"<database named credential>"],
  ["dbHostCredName":"<database host named credential>"],
  ["enableRedoShip":"true"]
  ["forceBackupModuleInstall":"true"]
}
```

Sample Body

```
{
  "configureRABackup":"true",
  "dbTargetType":"oracle_database",
  "dbTargetName":"zd2",
  "raTargetName":"ra",
  "raUser":"orsadmin",
  "dbCredName": "RMAN",
  "dbHostCredName":"AIMEDB_N"
}
```


Table 16-5 (Cont.) Configure Database to Send Backup

Feature	Description
Sample Response	<pre>{ "status": "Succeeded", "submissionName": "BACKUP_CONFIG_011419155235", "executionGUID": "7F69505339B4413FE053EB987B0A226E", "databaseName": "zd2", "message": "Deployment procedure to configure high availability settings to database successfully submitted." }</pre>
Body2	Configure database to send backups to Oracle Database Backup Cloud Service. <pre>{ "dbTargetName": "<database target name>", "dbTargetType": "oracle_database rac_database", "configureCloudBackup": "true", ["dbCredName": "<database named credential>"], ["dbHostCredName": "<database host named credential>"], ["forceBackupModuleInstall": "true"], ["cloudAccount": "<Hybrid Cloud Console cloud service account name>"] }</pre>
Sample Body	<pre>{ "configureCloudBackup": "true", "dbTargetType": "oracle_database", "dbTargetName": "dgprim", "dbHostCredName": "GIMAN_N", "dbCredName": "SYS", "cloudAccount": "endpoint:https://storage.example.com/v1/Storage-emdev2;identityDomain:emdev2;username:xxxx;password:xxxx;container:db19_cntr" }</pre>
Sample Response	<pre>{ "status": "Succeeded", "submissionName": "OPC_BACKUP_CONFIG_010919120757", "executionGUID": "7F000E71E0DE624BE053EB987B0ADF02", "databaseName": "dgprim", "message": "Deployment procedure to configure high availability settings to database successfully submitted." }</pre>

Table 16-5 (Cont.) Configure Database to Send Backup

Feature	Description
Body3	Install an RMAN Backup Module in the Oracle Home of the specified database or install the Recovery Appliance file system backup agent on the host of the database. <pre>{ "dbTargetName": "<database target name>", "dbTargetType": "oracle_database rac_database", "installSoftware": "true", "moduleType": "ra cloud", "installBackupModule": "true", ["dbHostCredName": "<database host named credential>"], ["forceBackupModuleInstall": "true"] }</pre>
Sample Body	<pre>{ "installSoftware": "true", "dbTargetType": "oracle_database", "dbTargetName": "dgprim", "moduleType": "cloud", "installBackupModule": "true", "dbHostCredName": "GIMAN_N" }</pre>
Sample Response	<pre>{ "status": "Succeeded", "submissionName": "INSTALL_BKP_MODULE_011019111214", "executionGUID": "7F1502D6C5CB5E53E053EB987B0AA38E", "databaseName": "dgprim", "message": "Deployment procedure to configure high availability settings to database successfully submitted." }</pre>
Request method	POST
Supported Since Release	EM DB 13.3.2.0.0

Table 16-6 Configure Backup Parameters

Parameter	Description
Body1	
<i>dbTargetName</i>	A single-instance or cluster database target name for the database that will be configured to send backups to the Recovery Appliance.
<i>dbTargetType</i>	The type of the target specified by "dbTargetName", either a single-instance database (oracle_database) or a cluster database (rac_database).

Table 16-6 (Cont.) Configure Backup Parameters

Parameter	Description
Body1	
<i>configureRABackup</i>	Configure database to be protected by a Recovery Appliance, with the ability to send backups to the Recovery Appliance. Any one option <i>configureRABackup</i> , <i>configureCloudBackup</i> or <i>installSoftware</i> must be provided. The possible value is "true".
<i>dbCredName</i>	The name of an existing Enterprise Manager database named credential that can be used to connect to the specified target database. If this argument is not specified, preferred credential will be used.
<i>dbHostCredName</i>	The name of an existing Enterprise Manager database host named credential that can be used to run operating system commands on the specified target database hosts. The credential should be for a user that has write permission for Oracle Home. If this argument is not specified, preferred credential will be used.
<i>raTargetName</i>	The target name of the Recovery Appliance that the specified database will be configured to send backups to.
<i>raUser</i>	The name of the Recovery Appliance database user that will be used by the specified database to send backups to and ship redo to the Recovery Appliance. This must be a virtual private catalog user, not the Recovery Appliance administrator user.
<i>enableRedoShip</i>	Enable realtime redo transport to the Recovery Appliance from the specified target database. Default value: True
<i>forceBackupModuleInstall</i>	Force installation of the version of the Recovery Appliance Backup Module stored in the Enterprise Manager software library into the Oracle Home of the specified target database, regardless of whether there is an existing Backup Module installed in the Oracle Home. This will overwrite any existing Backup Module, so this option should be selected only if it is known that the Backup Module version in the software library is at the same or later version than the version installed on the specified target database. If this option is not specified, the Backup Module will be installed only if there is no existing Backup Module installed in the Oracle Home. Default value: True
Body2	
<i>dbTargetName</i>	A single-instance or cluster database target name for the database that will be configured to send backups
<i>dbTargetType</i>	The type of the target specified by "dbTargetName", either a single-instance database (oracle_database) or a cluster database (rac_database).
<i>configureCloudBackup</i>	Configure database to send backups to Oracle Database Backup Cloud Service. Any one option <i>configureRABackup</i> , <i>configureCloudBackup</i> or <i>installSoftware</i> must be provided. The possible value is "true".
<i>dbCredName</i>	The name of an existing Enterprise Manager database named credential that can be used to connect to the specified target database. If this argument is not specified, preferred credential will be used.

Table 16-6 (Cont.) Configure Backup Parameters

Parameter	Description
Body1	
<i>dbHostCredName</i>	The name of an existing Enterprise Manager database host named credential that can be used to run operating system commands on the specified target database hosts. The credential should be for a user that has write permission for Oracle Home. If this argument is not specified, preferred credential will be used.
<i>forceBackupModuleInstall</i>	Force installation of the version of the Oracle Database Cloud Backup Module stored in the Enterprise Manager software library into the Oracle Home of the specified target database, regardless of whether there is an existing Backup Module installed in the Oracle Home. This will overwrite any existing Backup Module, so this option should be selected only if it is known that the Backup Module version in the software library is at the same or later version than the version installed on the specified target database. If this option is not specified, the Backup Module will be installed only if there is no existing Backup Module installed in the Oracle Home.
<i>cloudAccount</i>	The name of a cloud account configured under the Enterprise Manager Hybrid Cloud Setup console. All information needed to connect to the Backup Service will be obtained from the account and its associated Backup Service settings, including the service name, identity domain, username, password, and container (optional). If this argument is not specified, the global cloud account settings configured in the Database Backup Cloud Service Settings page will be used. (This argument is applicable only if the Hybrid Cloud Setup console has been used to configure accounts and Backup Service settings.). This is field should be provided with "configureCloudBackup" option. Example- "cloudAccount"."endpoint:https://storage.oracle.com/ Storageemdev2;identityDomain:xxx;username:xxxx;password:xxxx;container:xxxx"
Body3	
<i>dbTargetName</i>	Same as -configureRABackup usage above.
<i>dbTargetType</i>	Same as -configureRABackup usage above.
<i>installSoftware</i>	Install an RMAN Backup Module in the Oracle Homes of one or more databases or install the Recovery Appliance file system backup agent on the host of one database. Any one option configureRABackup, configureCloudBackup or installSoftware have be provided. The possible value is "true".
<i>dbHostCredName</i>	Same as -configureRABackup usage above.
<i>installBackupModule</i>	Install the version of the Backup Module (of the type specified by the moduleType option) stored in the Enterprise Manager software library in the Oracle Home of the specified target database. By default, the module will be installed only if there is not an existing module in the database Oracle Home, unless the force_backup_module_install flag is specified.
<i>moduleType</i>	The type of RMAN backup module to be installed. Possible values: • ra: Recovery Appliance Backup Module • cloud: Oracle Database Cloud Backup Module

Table 16-6 (Cont.) Configure Backup Parameters

Parameter	Description
Body1	
<i>forceBackupModuleInstall</i>	Used in combination with installBackupModule option. Same as -configureRABackup usage above.

Backup REST API

This section provides the details of the Backup REST API.

Table 16-7 Backup

Feature	Description
URL	https://<OMS_HOST>:<OMS_CONSOLE_PORT>/em/websvcs/restful/db/backup
Request header	Authorization: basic Content-Type: application/json
Body	Configure database to be protected by a Recovery Appliance, with the ability to send backups to the Recovery Appliance. <pre>{ "dbTargetName": "<database target name>", "dbTargetType": "oracle_database rac_database oracle_cloud_dbcs", ["dbCredName": "<database named credential>"], ["dbHostCredName": "<database host named credential>"], ["scope": "wholeDB tablespace datafile archivedLog recoveryFilesOnDisk cdbRoot pluggableDbs"], ["scopeValue": "<comma-separated list of values>"], ["backupType": "full incr0 incr1" ["incrType": "diff cumulative"]], ["destinationMedia": "disk tape ra cloud"], ["skipBackupArchivedLogs": "true false"], ["deleteArchivedLogsAfterBackup": "true false"], ["deleteObsoleteBackups": "true false"], ["maxFilesPerSet": "<numeric value>"], ["maxCorruptions": "<numeric value>"], ["sectionSize": "<numeric value>"], ["backupEncryptionMode": "wallet password both"], ["rmanEncryptionCredName": "<RMAN encryption named credential>"], ["encryptionAlgorithm": "AES128 AES192 AES256"], ["cloudAccount": "<cloud storage container url with endpoint, identityDomain, username, password and container name>"] }</pre>
Request method	POST
Supported Since Release	EM DB 13.3.2.0.0

Table 16-7 (Cont.) Backup

Feature	Description
Sample Body	<pre> { "dbTargetType":"oracle_database", "dbTargetName":"dgprim", "dbHostCredName":"GIMAN_N", "dbCredName": "SYS", "scope":"wholeDB", "backupType":"full", "backupEncryptionMode":"password", "encryptionAlgorithm":"AES128", "rmanEncryptionCredName":"WC1", "cloudAccount":"endpoint:https://storage.example.com/v1/ Storage-emdev2;identityDomain:emdev2;username:xxxx;pass- word:xxx;container:db19_cntr", "destinationMedia":"cloud", "maxFilesPerSet":2 } </pre>
Sample Response	<pre> { "status":"Succeeded", "submissionName":"Backup_DB_from_OMC_010919152109", "executionGUID":"7F044AF2D1F60905E053EB987B0A569B", "databaseName":"dgprim", "message":"Backup to Oracle Cloud deployment procedure successfully submitted." } </pre>

Table 16-8 Backup Parameters

Parameter	Description
Body1	
<i>dbTargetName</i>	A single-instance, cluster database, or database cloud service target name for the database that will be backed up.
<i>dbTargetType</i>	The type of the target specified by "dbTargetName", either a single-instance database (oracle_database), a cluster database (rac_database), or a database cloud service (oracle_cloud_dbcs).
<i>dbCredName</i>	The name of an existing Enterprise Manager database named credential that can be used to connect the specified target database. If the operating system user specified by the "dbHostCredName" argument is a member of the operating system DBA group, the database credential can specify a Normal role user. Otherwise, the database credential must specify a SYSDBA role user. If this argument is not specified, then preferred credential will be used.
<i>dbHostCredName</i>	The name of an existing Enterprise Manager database host named credential that can be used to run the RMAN command on the specified target database host. If this argument is not specified, then preferred credential will be used.

Table 16-8 (Cont.) Backup Parameters

Parameter	Description
<i>scope</i>	The scope of the backup. If this argument is not specified, the default value is wholeDB. Possible values are: - wholeDB: All data files in the database. - tablespace: One or more data files associated with the tablespace name. - datafile: One or more data files. - archivedLog: Exactly one copy of each distinct log sequence number. - recoveryFilesOnDisk: All recovery files on disk, whether they stored in the fast recovery area or other locations on disk. - cdbRoot: The data files in the whole container root. - pluggableDbs: The data files in one or more pluggable databases.
<i>scopeValue</i>	A comma-separated list of values to back up. It is required when the -scope value is tablespace datafile pluggableDbs.
<i>backupType</i>	The type of backup. If this argument is not specified, the default value is full. Possible values: - full: Back up all data blocks in data files being backed up. - incr0: Incremental level 0. Similar to a full backup, but can be used as the base of an incremental backup strategy. - incr1: Incremental level 1. Back up only changed blocks.
<i>incrType</i>	The type of incremental backup. Used in combination with -"backupType": "incr". If this argument is not specified, the default value is diff. Possible values: - diff: Back up blocks updated since the most recent level 0 or level 1 incremental backup. - cumulative: Back up all blocks changed since the most recent level 0 backup.
<i>destinationMedia</i>	The destination media for this backup. If this argument is not specified, the default value is disk. Possible values: - disk: Backup to disk (not allowed when -scope="recoveryFilesOnDisk"). - tape: Backup to a SBT (system backup to tape) device. - ra: Backup to a Recovery Appliance. - cloud: Backup to Database Backup Cloud Service.
<i>backupEncryptionMode</i>	Specify to encrypt the backup using the Oracle Encryption Wallet, a user-supplied password, or both. Possible values: - wallet: Oracle Encryption Wallet - password: user-supplied password - both: Oracle Encryption Wallet and user-supplied password This option can be specified when -scope="wholeDB tablespace datafile archivedLog recoveryFilesOnDisk cdbRoot pluggableDbs".
<i>rmanEncryptionCredentialName</i>	The name of an existing Enterprise Manager generic named credential that contains the encryption password. Required when "backupEncryptionMode": "password" is specified.
<i>encryptionAlgorithm</i>	The name of the encryption algorithm to use when encrypting backups. Used in combination with "backupEncryptionMode". Possible values are AES128, AES192, and AES256. If this argument is not specified, the default value is AES256.

Table 16-8 (Cont.) Backup Parameters

Parameter	Description
<i>skipBackupArchivedLogs</i>	Do not back up all archived logs on disk that have not been backed up. The default behavior if this argument is not specified is to back up all archived logs that have not been backed up. Possible values are true false. This option can be specified when <code>-scope="wholeDb tablespace datafile cdbRoot pluggableDbs"</code> .
<i>delete_Archived_Logs_After_Backup</i>	Delete all archived logs from disk after they are successfully backed up. Ignored if <code>"-skip_Backup_Archived_Logs"</code> is specified. Possible values are true false. This option can be specified when <code>-scope="wholeDB tablespace datafile cdbRoot pluggableDbs"</code> .
<i>deleteObsolateBackups</i>	Delete backups that are no longer required to satisfy the retention policy. Possible values are true false. This option can be specified when <code>-scope="wholeDB tablespace datafile cdbRoot pluggableDbs"</code> .
<i>maxFilesPerSet</i>	The maximum number of files to include in each backup set. This option can be specified when <code>-scope="wholeDB tablespace datafile archivedLog cdbRoot pluggableDbs"</code> .
<i>maxCorruptions</i>	The maximum number of physical corruptions allowed in data files. This option can be specified when <code>-scope="datafile"</code> .
<i>sectionSize</i>	The size in MB of each backup section produced during a data file backup, resulting in a multi-section backup where each backup piece contains one file section. This option cannot be used if the Maximum Backup Piece Size is set in the RMAN configuration. This option can be specified when <code>- scope="wholeDB tablespace datafile archivedLog cdbRoot pluggableDbs"</code> .
<i>cloudAccount</i>	The name of a cloud account configured under the Enterprise Manager Hybrid Cloud Setup console. All information needed to connect to the Backup Service will be obtained from the account and its associated Backup Service settings, including the service name, identity domain, username, password, and container (optional). If this argument is not specified, the global cloud account settings configured in the Database Backup Cloud Service Settings page will be used. (This argument is applicable only if the Hybrid Cloud Setup console has been used to configure accounts and Backup Service settings.). Example - <code>"cloudAccount": "endpoint:https://storage.oracle.com/Storageemdev2;identityDomain:xxx;username:xxxx;password:xxxx;container:xxxx"</code>

Part IV

Hybrid Cloud Management

This section covers the following:

- [Hybrid Solution: Oracle Cloud Infrastructure Management](#)
- [Hybrid Solution: Oracle Cloud Platform Management](#)

Hybrid Solution: Oracle Cloud Infrastructure Management

This chapter contains the following sections:

- [Hybrid Cloud Cloning Overview](#)
- [Cloning in Hybrid Cloud Use Cases](#)
- [Prerequisites for Cloning in Oracle Cloud](#)
- [Cloning to Oracle Cloud](#)
- [Cloning to Oracle Cloud Using Oracle Storage Service](#)
- [Cloning from Oracle Cloud](#)
- [Cloning Within Oracle Cloud](#)

Hybrid Cloud Cloning Overview

While managing the IT infrastructure of your enterprise, your Oracle Cloud applications may need to utilize and analyze data stored on-premise. In such a situation, you may need to migrate the data stored on-premise to Oracle Cloud.

Hybrid Cloud management enables you to monitor certain Oracle Cloud targets using an on-premise Cloud Control monitored by Enterprise Manager. The new cloning wizards leverage Hybrid Cloud management and enable you to clone an on-premise database, schema, or a PDB to a database or a PDB that is deployed in Oracle Cloud. Effectively, this enables you to copy or migrate your on-premise data to Oracle Public Cloud.

Additionally, you can also clone databases, schemas, or PDBs that are deployed in Oracle Cloud to a database or PDB that is deployed on-premise, as well as clone databases and PDBs within Oracle Cloud.

Cloning in Hybrid Cloud Use Cases

The following table lists the use cases covered when cloning in Hybrid Cloud:

Table 17-1 Cloning in Hybrid Cloud Use Cases

Cloning Use case	Cloning Solutions
Clone to Oracle Cloud	• Cloning a PDB to Oracle Cloud • Cloning Schema(s) to a DB or PDB on Oracle Cloud • Cloning a DB to a PDB on Oracle Cloud • Cloning a DB to Oracle Cloud (Oracle Compute Service) (This option is available only on the Oracle Compute Service.)
Clone to Oracle Cloud Using OSS	Cloning to Oracle Cloud Using Oracle Storage Service

Table 17-1 (Cont.) Cloning in Hybrid Cloud Use Cases

Cloning Use case	Cloning Solutions
Clone from Oracle Cloud	• Cloning a PDB from Oracle Cloud • Cloning Schema(s) from Oracle Cloud to a DB or PDB • Cloning a DB from Oracle Cloud to an On-Premise PDB • Cloning a DB from Oracle Cloud (Oracle Compute Service) (This option is available only on the Oracle Compute Service.)
Clone Within Oracle Cloud	• Cloning a PDB Within Oracle Cloud • Cloning a DB Within Oracle Cloud (Oracle Compute Service) (This option is available only on the Oracle Compute Service.)

Prerequisites for Cloning in Oracle Cloud

The following are the prerequisites for cloning in Oracle Cloud:

- The on-premise Cloud Control instance must have the latest patches applied from MOS (Doc ID 1549855.1)
- If you are cloning a database to Oracle Cloud, you must ensure that a Management Agent has been deployed on the destination host. Also, the destination target must be discovered.

Note:

If you do not have Internet access, the Management Agent must be configured in Hybrid Mode. For details on configuring the Management Agent in hybrid mode and deploying it on an Oracle Cloud target, see the Enabling Hybrid Cloud Management in the *Oracle Enterprise Manager Cloud Administrator's Guide*.

- Cloning is supported only if Oracle Software Library is **not** configured with an up-load location of the OMS Agent storage type.
- It is recommended that you use a Test Master database or a Test Master pluggable database for cloning to Oracle Cloud.

To create a Test Master database, refer to the *Oracle Enterprise Manager Cloud Administration Guide*.

To create a Test Master pluggable database, refer to the *Oracle Enterprise Manager Cloud Administration Guide*.

- The on-premise database and the database on Oracle Cloud should not be encrypted, should possess the same character set, and should have the same patch set level.
- When cloning from on-premise to Oracle Cloud, the SELINUX security policy must set to **permissive**. If the security policy is enforced or set to **enforcing**, the cloning procedure may fail in the SecureCopyFiles step in the non-advanced wizard mode. You will need to configure the SELinux to allow RSYNC from the Agent (script).

Cloning to Oracle Cloud

To clone a database, schema(s), or a pluggable database from on-premise to Oracle Cloud, refer to the following use cases:

- [Cloning a DB to Oracle Cloud \(Oracle Compute Service\)](#)
- [Cloning a PDB to Oracle Cloud](#)
- [Cloning Schema\(s\) to a DB or PDB on Oracle Cloud](#)
- [Cloning a DB to a PDB on Oracle Cloud](#)


Cloning a DB to Oracle Cloud (Oracle Compute Service)

You can clone an on-premise database using one of the following options:

- [Cloning a DB to Oracle Cloud Using the Clone Wizard](#)
- [Cloning a DB to Oracle Cloud Using EMCLI](#)

This figure shows how data is transferred between the on-premise host and Oracle Cloud.

Figure 17-1 Cloning a Database to Oracle Cloud



Cloning a DB to Oracle Cloud Using the Clone Wizard

You can create a full clone of an on-premise database from Enterprise Manager to Oracle Public Cloud. While cloning a database instance, you can optionally choose to clone the Oracle Home. This allows you to move your on-premise database to Oracle Public Cloud. To clone an on-premise database to cloud, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Select the database to be cloned and right click the mouse button.

4. From the **Oracle Database** menu, select **Cloning**, then select **Clone to Oracle Cloud**.
5. On the Source and Destination: Clone to Oracle Cloud page, do the following:
 - In the Source Credentials section, specify or search and select the database and Oracle Home credentials. If the source database is encrypted, specify the Wallet Password. Click OK.

 Note:

If the source database is present on a hybrid gateway, the backups and oracle home binaries can be directly be transferred to the Oracle Public Cloud machine. If the source database is present on any other host machine, the Hybrid Gateway details must be specified in this page.

- In the Data Time Series section, the Prior Point in Time option is selected by default and a timeline of RMAN Backup images appear. Select the appropriate RMAN backup by clicking Select on the image.
- In the Source Backup Location field, the location of the source backup you have selected appears.
- In the Destination section, in the Clone to field, the Select Compute Cloud Service is selected by default. This indicates that you are requesting resources on Oracle Public Cloud to clone your on-premise database.
- In the Database Definition section, specify the Display Name for the database, the Global Database Name, SID.
- In the Hosts section, you can choose:
 - Clone Oracle Home: If the Oracle Home is not pre-installed on the Oracle Public Cloud machine on which the database is being cloned, select the Clone Oracle Home option.
 - Use an Existing Home: If the Oracle Home is already cloned, select the Use an Existing Home option
- If you select the Clone Oracle Home option, select the Oracle Public Cloud host machine on which the Oracle Home is to be provisioned. The Shape (size) of the Oracle Public Cloud machine and the Site (Oracle Cloud) is displayed. Specify the following:
 - Database Host Credentials: Specify the database host credentials.

 Note:

The credentials for accessing the Oracle Public Cloud machine must be SSH key based credentials.

- Run root scripts that require root privileges: You can either specify the root credentials or choose to enter them later to run the root scripts.

In the Stage Location field, the location on the Oracle Public Cloud machine where the backups will be staged before cloning is displayed.

- If you select the Use an Existing Home option, specify the following:
 - Oracle Home Location: Select the location of the Oracle Home on an Oracle Public Cloud machine already managed in Enterprise Manager. The host name of the machine in which the Oracle Home is present is displayed.
 - Database Host Credentials: Specify the database host credentials.

In the Stage Location field, the location on the Oracle Public Cloud machine where the backups will be staged before cloning is displayed.

 Note:

All credentials specified for the Oracle Public Cloud machine must be SSH credentials.

- Hybrid Gateway Server: Specify the host name and credentials for the Hybrid Gateway Server (configured Enterprise Manager Management Agent) which transfers data from Enterprise Manager to Oracle Public Cloud.

Data is first transferred to the stage location on the Hybrid Gateway Server and then copied to the Oracle Public Cloud host. This stage location (/tmp by default) must contain enough space to store the backups and binaries.

6. Click **Next**. The Configuration page appears. Specify the following details:

- **Software Location:** Specify the Oracle Base and Oracle Home location where the database is to be cloned.
- **Oracle Database User Groups:** Specify the name of the Oracle Database User Group performing this cloning operation.
- **Database Files Location:** Specify the location where you want the data files, temp files, redo log files, and control files to be created.
- **Recovery Files Location:** Specify the location where you want the recovery files, such as archived redo logs, RMAN backups, and other related files to be created. You can choose to use the fast recovery area. If you do, specify the fast recovery area size. In the Database Mode section, you can enable Read Only.
- Listener Configuration: In the Hosts section on the previous page, if you have selected:
 - **Clone Oracle Home:** Specify the Listener Name and Port to create a new listener with which the database will be registered.
 - **Use an Existing Home:** Select the listener targets running under the new Oracle Home, to register the clone database. You can add new listeners by clicking Add, and specifying the listener name and port.
- **Database Credentials:** Specify passwords for the SYS, SYSTEM, and DBSNMP administrative users in the clone database. You can choose to have the same password for all the three users or a different password for each.
- **Transparent Data Encryption:** If the source database is encrypted, specify the location in which the Wallet is to be placed on the Oracle Public Host machine. Select the Use same password as source checkbox to use the same password as the host database. Select the Reset/Regenerate the Master Encryption Key to reset the encryption key.

- **Parallelism:** Setting a higher Degree of Parallelism will ensure faster recompilation of objects. The default setting is the number of CPU count of the host. Click Next.
7. On the Initialization Parameters page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click Edit to modify the value of the parameter. Click Next.
 8. On the Post Processing page, you can select masking rules, subsetting rules, or custom scripts such as pre script, post script, and post SQL script. Click Next.
 9. On the Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning deployment procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive e-mail notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

10. On the Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Submit** to launch the procedure to clone the database, transfer the selected backups to the destination host via the hybrid gateway. Once cloned, this database can be managed as an Enterprise Manager target.

To refresh the cloned database, from the **Oracle Database** menu, select **Cloning**, then select **Clone Management**. On the Clone Management Dashboard, select the database and click **Refresh**. Accept the default values in the wizard and optionally change the Initialization Parameters and Configuration details and click **Refresh** to refresh the database.

Cloning a DB to Oracle Cloud Using EMCLI

To clone a database to Oracle Cloud using EMCLI, enter the following command:

```
emcli db_clone_management -cloneToCloud -input_file=data:/xyz/sdf/
clone2cloud.props
where clone2cloud.props is the properties file which provides the cloning parameters
and their values.
```

Sample Properties File

```
#-----#
#                               Source Details                               #
#-----#
SRC_DB_TARGET_NAME=dbprod.example.com
SRC_DB_TARGET_TYPE=database1
SRC_HOST_NORMAL_NAMED_CRED=SRCHOSTCRED:SYSMAN
SRC_DB_CRED=DBPROD_CRED:SYSMAN
SRC_WORK_DIR=/tmp
PROFILE_INTERNAL_PATH=Database Provisioning Profiles/11.2.0.4.0/linux_x64/Database
db11.us.xyz.com Profile 22-02-2016 05:13 PM
PROFILE_VERSION_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_DatabaseConfigProfile:2C5B357D5FF67E2AE053F17CB10A53F1:0.1
SOURCE_BACKUP_LOCATION=/scratch/dbprod/app/dbprod/fast_recovery_area/back-
up1456141446021
```

```
#-----#
#           Hybrid Gateway Details
#-----#
FORWARDER_HOST=slc04wim.example.com
FORWARDER_HOST_CREDS=HOST:SYSMAN
FORWARDER_STAGE_LOCATION=/scratch/fwd

#-----#
# Destination Details
#-----#
TARGET_HOST_LIST=129.124.22.34
HOST_NORMAL_NAMED_CRED=OPC_CRED:SYSMAN
HOST_PRIV_NAMED_CRED=HOST_CRED_ROOT:SYSMAN
DEST_LOCATION=/u03/backup
DEST_WORK_DIR=/tmp
#-----#
# Database Definition
#-----#
COMMON_DB_SID=Dcln1
COMMON_DOMAIN_NAME=us.xyz.com
COMMON_GLOBAL_DB_NAME=Dcln1.xyz.com
DATABASE_TYPE=dbTypeSI
DB_ADMIN_PASSWORD_SAME=true
DATABASE_PASSWORDS=welcome1

#-----#
# Software Configuration
#-----#
CLONE_HOME=Y # mention it as N if only database cloning is required
#All these below parameters are only required if a new oracle home provisioning
needs to be done.
OINSTALL_GROUP=dba
OSBACKUPDBA_GROUP=dba
OSDBA_GROUP=dba
OSDGDBA_GROUP=opc
OSKMDBA_GROUP=opc
OSOPER_GROUP=opc
ORACLE_BASE_LOC=/u03/home/app/base
ORACLE_HOME_LOC=/u03/home/app/base/product/11.2.0/dbhome_1

#-----#
# Database Configuration
#-----#
ARCHIVE_LOG_MODE=YES
DB_FILE_LOC=/u03/home/oradata
DB_TARGET_DISPLAY_NAME=Dcln1
DB_TARGET_NAME=Dcln1.xyz.com
FLASH_REC_AREA=/u03/home/app/home/fast_recovery_area
FLASH_REC_AREA_SIZE=4182
LISTENER_PORT=2887
#Only needed if we are provisioning a new home
LISTENER_NAME=LIST_2887

#-----#
# Data file transfer configuration parameters
#-----#
# This will be moved to emctl metadata. for now its a variable in DP
NUMBER_OF_THREADS=10
# This will be moved to emctl metadata. for now its a variable in DP
# The chunk size in bytes. For now default is 250MB
```



```
MIN_CHUNK_SIZE=262144000
#Transfer Unit can be FILE or CHUNK
TRANSFER_UNIT=FILE
#Maximum speed in Kbps
BANDWIDTH_LIMIT=2048
#Time out in seconds for data transfer
TIME_OUT=1200

#-----#
#   Wallet Configuration
#-----#
DEST_PATH_TO_WALLET=/u03/wallet
WALLET_PASSWORD=welcome123
NEW_WALLET_PASSWORD=welcome1
REKEY_REQUIRED=true
```

Cloning a PDB to Oracle Cloud

To clone an on-premise PDB to a PDB on Oracle Cloud, you can use either of the following solutions:

- [Cloning a PDB to Oracle Cloud Using the Clone Wizard](#)
- [Cloning a PDB to Oracle Cloud Using EM CLI](#)

Cloning a PDB to Oracle Cloud Using the Clone Wizard

To clone a PDB to a CDB deployed in Oracle Cloud, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source CDB (the CDB that the source PDB is a part of) in the list, then click the name of the PDB that you want to clone.
4. From the **Oracle Database** menu, select **Cloning**, then select **Clone to Oracle Cloud**.

Alternatively, in Step 3, you can right click the name of the PDB that you want to clone, select **Oracle Database**, select **Cloning**, then select **Clone to Oracle Cloud**.

5. On the Source and Destination: Clone to Oracle Cloud page, do the following:
 - In the Credentials section, specify the SYSDBA credentials for the source CDB, and the host credentials for the source CDB. You can choose to use the preferred credentials, use a saved set of named credentials, or specify a new set of credentials.
 - In the Pluggable Database Definition section, specify a name, and a display name for the PDB clone. Enterprise Manager uses the display name to identify the PDB clone target.
 - In the PDB Administrator Credentials section, specify the credentials of the Admin user account that you want to use to administer the PDB clone.
 - In the Container Database section, specify the destination CDB that is deployed in Oracle Cloud (the CDB that the PDB clone must be a part of).
 - In the Credentials section, specify the SYSDBA credentials for the destination CDB, and the host credentials for the destination CDB.

6. If you do not need to specify anymore details, click **Clone**. This submits the deployment procedure to clone a PDB to a CDB that is deployed in Oracle Cloud.
To specify other configuration details, mask data, as well as schedule the cloning process, click **Advanced**.
Follow the rest of the steps, if you have selected the Advanced option.
7. On the Clone to Oracle Cloud: Source and Destination page, verify the details, and then click **Next**.
8. On the Clone to Oracle Cloud: Configuration page, in the Database Files Location section, specify the storage location where the datafiles of the PDB clone must be stored.

In the Advanced Configuration section, specify the storage limits for the maximum size of the PDB clone, and the maximum size of a shared table space within the PDB clone. By default, no limits are placed on the values for these attributes.

In the Miscellaneous section, select the logging option that you want to use for the table spaces created within the PDB clone.

Click **Next**.
9. On the Clone to Oracle Cloud: Post Processing page, in the Data Masking section, specify the data masking definition that you want to apply after cloning the PDB. Data masking masks sensitive data in a database.

For information on how to create a data masking definition, see *Oracle Data Masking and Subsetting Guide*. Note that you can apply a data masking definition only if you have the Subset-Masking license pack.

In the Custom Scripts section, for **Pre Script** and **Post Script**, specify the Oracle Software Library components that contain the scripts that you want to run before cloning, and after cloning the PDB respectively. Also, for **SQL Script**, specify the SQL scripts that you want to run after cloning the PDB. For **Run As User**, select the user account that you want to use to run the SQL scripts.

Click **Next**.
10. On the Clone to Oracle Cloud: Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning deployment procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive email notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.
11. On the Clone to Oracle Cloud: Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Clone** to submit the deployment procedure to clone a PDB to a CDB that is deployed in Oracle Cloud.

Cloning a PDB to Oracle Cloud Using EM CLI

You can clone an on-premise pluggable database to Oracle Cloud. Before you proceed with the EM CLI command, it is recommended that you create a Test Master of the on-premise PDB and use the Test Master to create a clone. This is recommended so as to mask the data before it can be transferred over the internet.

To clone an on-premise pluggable database to Oracle Cloud, enter the EM CLI verb `emcli pdb_clone_management -input_file=data:/xyz/sdf/pdb_clone.props -cloneToOracleCloud`, where `pdb_clone.props` is the properties file which provides the cloning parameters and their values.

There are three methods in which you can clone a pluggable database. The difference between each of these methods is in the configuration of certain parameters in the properties file. The 3 methods and the details of the configuration parameters are explained below:

- **Operating System (OS) Image backup**

Takes a backup of the source PDB and creates a new PDB. The `BACKUP_TYPE` parameter should specify the type of backup. The allowed values for `BACKUP_TYPE` are `OSIMAGE`, `RMAN` and `TAR`. The `EXISTING_BACKUP` and `EXISTING_BACKUP_METADATA` parameters should not be provided.

Sample properties file:

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
BACKUP_TYPE=OSIMAGE
```

- **Existing backup**

Uses an existing backup of the source PDB and creates a new PDB. The `BACKUP_TYPE` parameter should specify the type of backup. The allowed values for `BACKUP_TYPE` are `OSIMAGE`, `RMAN` and `TAR`. The `EXISTING_BACKUP` parameter should specify the location with the backup name and `EXISTING_BACKUP_METADATA` should specify the location and the metadata file name for the backup.

Sample properties file:

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
EXISTING_BACKUP=/user1/pdbbackup/PDB1_Backup_14297779
EXISTING_BACKUP_METADATA=/user1/pdbbackup/PDB1_Backup_14297779/PDB1.xml
BACKUP_TYPE=RMAN
```

 Note:

To create a PDB backup, enter the verb `emcli pdb_backup -input-File="loaction of file containing properties required for taking backup of PDB"`, where the sample contents of the properties file is as follows:

```
TARGET_HOST_LIST=xyz.abccorp.com
HOST_NORMAL_NAMED_CRED=XYZ_CRED:CRED_OWNER
SRC_CDB_NAMED_CRED=CDB1_CRED:CRED_OWNER
SRC_CDB_TARGET_NAME=CDB1
SRC_CDB_TARGET_TYPE=oracle_database
SRC_PDB_TARGET_NAME=CDB1_PDB1
BACKUP_LOCATION=/user1/pdbbackup
WORK_DIR_LOCATION=/tmp
ORACLE_HOME_LOC=/scratch/d121hmcasm/product/12.1.0/dbhome_1
```

- **Unplug/plug**

Unplugs the source PDB and creates a new PDB at the destination using the unplugged source, and then plugs the source back. *EXISTING_BACKUP*, *EXISTING_BACKUP_METADATA* and *BACKUP_TYPE* parameters should not be provided.

Sample properties file:

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLCO_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
```

 Note:

For all the three methods stated above, in case the destination PDB data files location is ASM then add the parameter *DEST_STAGE_DIR* who's value will be used as the destination while transferring the source PDB data files. This parameter is optional, if it is not provided a temporary directory will be used. For Linux systems the temporary directory is `/tmp`.

Cloning Schema(s) to a DB or PDB on Oracle Cloud

You can clone a schema that is on-premise to Oracle Cloud either as a database or a pluggable database using EM CLI verbs.

 Note:

As a prerequisite, it is recommended that you create a Test Master of the schema database and to use the schema of the Test Master to create a clone.

To clone a schema that is on-premise to Oracle Cloud either as a database or a pluggable database, follow the steps below:

1. Enter the EM CLI verb `emcli describe_dbprofile_input -data_mode=EXPORT`. The output provides all profile creation input variables.

 Note:

Export is supported only for database and schema whereas import is supported for both database and PDBs. This indicates that the source should always be a database or schema and the destination can either be a database or a PDB.

2. Use the input variables to create a properties file with values for all the variables.
3. Create the directory object `SCHEMAS_DUMP_DIR`.
4. Export data from the source database by creating a database profile. To do so, enter the verb `emcli create_dbprofile - input_file=data:<properties file name along with path>`.

 Note:

Use the properties file created in the previous step for this verb.

Sample properties file:

```
#-----#
# SOURCE                                     #
#-----#
REFERENCE_DATABASE=SS_REF_TD_DB
REFERENCE_DATABASE_TYPE=oracle_database
REF_DB_CREDENTIALS=CRED_DB:sysman
REF_HOST_CREDENTIALS=CRED_HOST:sysman

#-----#
# DATA CONTENT DETAILS                     #
#-----#
DATA_CONTENT_MODE=EXPORT
DATA_CONTENT=METADATA_AND_DATA

#-----#
# EXPORT DETAILS                           #
#-----#
EXPORT.EXPORT_TYPE=SELECTED_SCHEMAS
EXPORT.SCHEMA_INCLUDE_LIST.0=HR
```

```

EXPORT.SCHEMA_INCLUDE_LIST.1=PM
EXPORT.SCHEMA_INCLUDE_LIST.2=OE
EXPORT.SCHEMA_INCLUDE_LIST.3=IX
EXPORT.SCHEMA_INCLUDE_LIST.4=SH
EXPORT.SCHEMA_INCLUDE_LIST.5=BI
EXPORT.DEGREE_OF_PARALLELISM=1
EXPORT.DUMP_DIRECTORY_LIST.0=directory=SCHEMAS_DUMP_DIR,file_name=samplsche-
mas.dmp,max_size=100
EXPORT.LOG_FILE_DIRECTORY=directory=SCHEMAS_DUMP_DIR,file_name=samplschemas.log

```

```

#-----#
#          PROFILE DETAILS          #
#-----#
PROFILE_NAME=Export Dump of Sample schemas10
PROFILE_VERSION=11.2.0.4.0
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/
WORKING_DIRECTORY=/tmp

```

5. Enter the verb to transfer data: `emcli data_transfer -input_file=data:/u01/files/data_trans.props.`

Sample properties file:

```

#-----#
# SOURCE                                #
#-----#
SRC_HOST_CREDS=NC_HOST_SRAY
SOURCE_LOCATION=/tmp/newp/PDB_Backup_142838
SRC_HOST=b12.idc.example.com

#-----#
# DESTINATION                          #
#-----#
DEST_HOST_CREDS=NC_HOST_SRAY
DEST_LOCATION=/scratch/sray/app3/sray/oradata/migda
DEST_HOST=slo.us.example.com

#-----#
# HYBRID GATEWAY / FORWARDER          #
#-----#
FORWARDER_HOST=slo.us.example.com
FORWARDER_CRED=ACD_NY:SYSCO
WORKING_DIRECTORY=/tmp
#-----#
# Data file transfer configuration parameters
#-----#
NUMBER_OF_THREADS=10
#The chunk size in bytes. For now default is 250MB
MIN_CHUNK_SIZE=262144000
#Transfer Unit can be FILE or CHUNK
TRANSFER_UNIT=FILE
#Maximum speed in Kbps
BANDWIDTH_LIMIT=2048
#Time out in seconds for data transfer
TIME_OUT=1200

```

 Note:

Remove the Hybrid Gateway parameters if the SSH connection exists between the source and the destination hosts.

6. Enter the verb to import data in to the destination database: `emcli dbimport -input_file=data:/u01/files/dbimport.props.`

 Note:

To clone the destination to database or pluggable database, ensure you provide the required value in the `DESTINATION_TARGET_TYPE` option in the properties file. For database, enter `oracle_database`, and for PDB enter `oracle_pdb`.

Sample properties file:

```
#-----#
#          DESTINATION          #
#-----#
DESTINATION_TARGET=SS_OPC_DB
DESTINATION_TARGET_TYPE=oracle_database
DATABASE_CREDENTIAL=CRED_DB:sysman
HOST_NAMED_CREDENTIAL=CRED_HOST:sysman

#-----#
#          PROFILE              #
#-----#
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/Export Dump
of Sample schemas10

#-----#
#          SCHEMA DETAILS      #
#-----#
REMAP_SCHEMA_LIST.0=HR:HR
REMAP_SCHEMA_LIST.1=OE:OE
REMAP_SCHEMA_LIST.2=PM:PM
REMAP_SCHEMA_LIST.3=IX:IX
REMAP_SCHEMA_LIST.4=SH:SH
REMAP_SCHEMA_LIST.5=BI:BI
REMAP_TABLESPACE_LIST.0=EXAMPLE:MYTBSP1
REMAP_TABLESPACE_LIST.1=USERS:MYTBSP1
REMAP_TABLESPACE_LIST.2=SYSTEM:MYTBSP1
DEGREE_OF_PARALLELISM=1
DUMP_FILE_LIST.0=/scratch/ae/dumpdir/sampschemas.dmp
IMPORT_LOG_FILE_DIRECTORY=DATA_PUMP_DIR
```

Cloning a DB to a PDB on Oracle Cloud

You can clone a database that is on-premise to a pluggable database on Oracle Cloud using EM CLI verbs.

 Note:

As a prerequisite it is recommended that you create a Test Master of the database and use the Test Master to create a clone.

To clone a database that is on-premise to Oracle cloud as a pluggable database, follow the steps below:

1. Enter the EM CLI verb `emcli describe_dbprofile_input -data_mode=EXPORT`. The output provides all profile creation input variables.

 Note:

Export is supported only for database and schema whereas import is supported for both database and PDBs. This indicates that the source should always be a database or schema and the destination can either be a database or a PDB.

2. Use the input variables to create a properties file with values for all the variables.
3. Create the directory object `SCHEMAS_DUMP_DIR`.
4. Export data from the source database by creating a database profile. To do so, enter the verb `emcli create_dbprofile - input_file=data:<properties file name along with path>`.

 Note:

Use the properties file created in the previous step for this verb.

Sample properties file:

```
#-----#
# SOURCE                                     #
#-----#
REFERENCE_DATABASE=SS_TM_DB
REFERENCE_DATABASE_TYPE=oracle_database
REF_DB_CREDENTIALS=CRED_DB:sysman
REF_HOST_CREDENTIALS=CRED_HOST:sysman

#-----#
# DATA CONTENT DETAILS                       #
#-----#
DATA_CONTENT_MODE=EXPORT
DATA_CONTENT=METADATA_AND_DATA

#-----#
# EXPORT DETAILS                             #
#-----#
EXPORT.EXPORT_TYPE=FULL_DATABASE
EXPORT.DEGREE_OF_PARALLELISM=1
EXPORT.DUMP_DIRECTORY_LIST.0=directory=SCHEMAS_DUMP_DIR,file_name=samplsche-
```



```
mas.dmp,max_size=100
EXPORT.LOG_FILE_DIRECTORY=directory=SCHEMAS_DUMP_DIR,file_name=samplschemas.log
```

```
#-----#
#          PROFILE DETAILS          #
#-----#
PROFILE_NAME=Export Dump of Sample schemas10
PROFILE_VERSION=11.2.0.4.0
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/
WORKING_DIRECTORY=/tmp
```

5. Enter the verb to transfer data: `emcli data_transfer -input_file=data:/u01/files/data_trans.props.`

Sample properties file:

```
#-----#
# SOURCE                               #
#-----#
SRC_HOST_CREDS=NC_HOST_SRAY
SOURCE_LOCATION=/tmp/newp/PDB_Backup_1428003803938
SRC_HOST=b12.idc.example.com

#-----#
# DESTINATION                          #
#-----#
DEST_HOST_CREDS=NC_HOST_SRAY
DEST_LOCATION=/scratch/sray/app3/sray/oradata/migda
DEST_HOST=slo.us.example.com

#-----#
# HYBRID GATEWAY / FORWARDER          #
#-----#
FORWARDER_HOST=slo.us.example.com
FORWARDER_CRED=ACD_NY:SYSCO
WORKING_DIRECTORY=/tmp
```

 Note:

Remove the Hybrid Gateway parameters if the SSH connection exists between the source and the destination hosts.

6. Enter the verb to import data in to the destination database: `emcli dbimport -input_file=data:/u01/files/dbimport.props.`

 Note:

To clone the destination to a pluggable database, ensure you enter `oracle_pdbin` the `DESTINATION_TARGET_TYPE` option in the properties file.

Sample properties file:

```
#-----#
#          DESTINATION                 #
#-----#
```

```

DESTINATION_TARGET=SS_OPC_DB
DESTINATION_TARGET_TYPE=oracle_database
DATABASE_CREDENTIAL=CRED_DB:sysman
HOST_NAMED_CREDENTIAL=CRED_HOST:sysman

#-----#
#          PROFILE                      #
#-----#
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/Export Dump
of Sample schemas10

#-----#
#          SCHEMA DETAILS                #
#-----#
REMAP_SCHEMA_LIST.0=HR:HR
REMAP_SCHEMA_LIST.1=OE:OE
REMAP_SCHEMA_LIST.2=PM:PM
REMAP_SCHEMA_LIST.3=IX:IX
REMAP_SCHEMA_LIST.4=SH:SH
REMAP_SCHEMA_LIST.5=BI:BI
REMAP_TABLESPACE_LIST.0=EXAMPLE:MYTBSP1
REMAP_TABLESPACE_LIST.1=USERS:MYTBSP1
REMAP_TABLESPACE_LIST.2=SYSTEM:MYTBSP1
DEGREE_OF_PARALLELISM=1
DUMP_FILE_LIST.0=/scratch/ae/dumpdir/sampschemas.dmp
IMPORT_LOG_FILE_DIRECTORY=DATA_PUMP_DIR

```

Cloning to Oracle Cloud Using Oracle Storage Service

You can clone an on-premise database to Oracle Cloud by using backups stored in the Oracle Storage Service.

Topics

- [Cloning a Database to Oracle Cloud Using the Clone Wizard \(OSS\)](#)

Cloning a Database to Oracle Cloud Using the Clone Wizard (OSS)

You can clone your on-premise database to Oracle Cloud using backups present on the Oracle Storage Service. These backups can be accessed both by the on-premise database as well as Oracle Cloud.

Cloning a database to Oracle Cloud using the Oracle Storage Service involves the following steps:

- Defining the Oracle Cloud Backup Service Settings.
- Configuring the Oracle Database Backup Cloud Service Account
- Creating a database profile that is backed up and stored on the OSS
- Cloning the Database

Defining the Oracle Cloud Backup Service Settings

The Oracle Public Cloud Database Backup Settings page allows you to define the storage account settings used to connect to the Oracle Cloud Database Backup Service. The global settings on this page are used when configuring individual databases to send backups to the Oracle Cloud Backup Service.

Specify the Service Name, Identity Domain, User ID, and Password. Optionally, if you are using a custom storage container for your backups, you can specify it in the Backup Container field. Click Save to save your settings.

Configuring the Oracle Database Backup Cloud Service Account


See [Defining the Database Backup Cloud Service Settings](#)

Creating a Database Profile (On OSS)

Create a database profile on the Oracle Storage Service. On the Database Profile page, click the **Oracle Cloud as Media** radio button, select the required option from the storage accounts drop down list and follow the steps to create the profile.

Cloning the Database

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Select the database to be cloned and right click the mouse button.
4. From the **Oracle Database** menu, select **Cloning**, then select **Clone to Oracle Cloud**.
5. On the Source and Destination: Clone to Oracle Cloud page, do the following:
 - In the Source Credentials section, specify or search and select the database and Oracle Home credentials. If the source database is encrypted, specify the Wallet Password. Click **OK**.



- In the Data Time Series section, the Prior Point in Time option is selected by default and a timeline of RMAN Backup on Oracle Cloud images appear. These backups are present in the Oracle Storage Service and can be directly used to clone the database on Oracle Cloud. Select the appropriate RMAN backup by clicking **Select** on the image.
- In the Source Backup Location field, the location of the source backup you have selected appears.
- In the Database Definition section, specify the **Display Name** for the database, the **Global Database Name**, **SID**.
- In the Hosts section, you can choose:

- **Clone Oracle Home:** If the Oracle Home is not pre-installed on the Oracle Public Cloud machine on which the database is being cloned, select the **Clone Oracle Home** option.

 Note:

If you are provisioning a new Oracle Home and if your source is not acting as a hybrid gateway, you must use the hybrid cloud agent to transfer the binaries.

- **Use an Existing Home:** If the Oracle Home is already cloned, select the **Use an Existing Home** option
 - If you select the **Clone Oracle Home** option, select the **Oracle Public Cloud** host machine on which the Oracle Home is to be provisioned. The Shape (size) of the Oracle Public Cloud machine and the Site (Oracle Cloud) is displayed. Specify the following:
 - **Database Host Credentials:** Specify the database host credentials.
 - **Run root scripts that require root privileges:** You can either specify the root credentials or choose to enter them later to run the root scripts.
 - If you select the **Use an Existing Home** option, specify the following:
 - **Oracle Home Location:** Select the location of the Oracle Home on an Oracle Public Cloud machine already managed in Enterprise Manager. The host name of the machine in which the Oracle Home is present is displayed.
 - **Database Host Credentials:** Specify the database host credentials.
6. Click **Next** to navigate to the Configuration page. Specify the following details:
- **Software Location:** Specify the Oracle Base and Oracle Home location where the database is to be cloned.
 - **Oracle Database User Groups:** Specify the name of the Oracle Database User Group performing this cloning operation.
 - **Database Files Location:** Specify the location where you want the data files, temp files, redo log files, and control files to be created.
 - **Recovery Files Location:** Specify the location where you want the recovery files, such as archived redo logs, RMAN backups, and other related files to be created. You can choose to use the fast recovery area. If you do, specify the fast recovery area size. In the **Database Mode** section, you can enable **Read Only**.
 - **Listener Configuration:** In the Hosts section on the previous page, if you have selected:
 - **Clone Oracle Home:** Specify the Listener Name and Port to create a new listener with which the database will be registered.
 - **Use an Existing Home:** Select the listener targets running under the new Oracle Home, to register the clone database, You can add new listeners by clicking Add, and specifying the listener name and port.

- **Database Credentials:** Specify passwords for the SYS, SYSTEM, and DBSNMP administrative users in the clone database. You can choose to have the same password for all the three users or a different password for each.
 - **Transparent Data Encryption:** If the source database is encrypted, specify the location in which the Wallet is to be placed on the Oracle Public Host machine. Select the Use same password as source checkbox to use the same password as the host database. Select the Reset/Regenerate the Master Encryption Key to reset the encryption key.
 - **Parallelism:** Setting a higher Degree of Parallelism will ensure faster recompilation of objects. The default setting is the number of CPU count of the host. Click Next.
7. On the Initialization Parameters page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click **Edit** to modify the value of the parameter. Click Next.
 8. On the Post Processing page, you can select masking rules, subsetting rules, or custom scripts such as pre script, post script, and post SQL script. Click **Next**.
 9. On the Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning deployment procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive e-mail notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

10. On the Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Submit** to launch the procedure to clone the database, transfer the selected backups to the destination host via the hybrid gateway. Once cloned, this database can be managed as an Enterprise Manager target.

To refresh the cloned database, from the **Oracle Database** menu, select **Cloning**, then select **Clone Management**. On the Clone Management Dashboard, select the database and click **Refresh**. Accept the default values in the wizard and optionally change the Initialization Parameters and Configuration details and click **Refresh**.

Cloning from Oracle Cloud

To clone a database, schema(s), or a PDB from Oracle Cloud, refer to the following use cases:

- [Cloning a PDB from Oracle Cloud](#)
- [Cloning Schema\(s\) from Oracle Cloud to a DB or PDB](#)
- [Cloning a DB from Oracle Cloud to an On-Premise PDB](#)

Cloning a PDB from Oracle Cloud

To clone a PDB from Oracle Cloud to an on-premise PDB, you can use either of the following solutions:

- [Cloning a PDB from Oracle Cloud Using the Clone Wizard](#)
- [Cloning a PDB from Oracle Cloud Using EM CLI](#)

Cloning a PDB from Oracle Cloud Using the Clone Wizard

To clone a PDB from Oracle Cloud to an On-Premise PDB, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source CDB (the CDB that the source PDB is a part of) in the list, then click the name of the PDB that you want to clone.
4. From the **Oracle Database** menu, select **Cloning**, then select **Clone from Oracle Cloud**.

Alternatively, in Step 3, you can right click the name of the PDB that you want to clone, select **Oracle Database**, select **Cloning**, then select **Clone from Oracle Cloud**.

5. On the Source and Destination: Clone from Oracle Cloud page, do the following:
 - In the Credentials section, specify the SYSDBA credentials for the source CDB, and the host credentials for the source CDB. You can choose to use the preferred credentials, use a saved set of named credentials, or specify a new set of credentials.
 - In the Pluggable Database Definition section, specify a name, and a display name for the PDB clone. Enterprise Manager uses the display name to identify the PDB clone target.
 - In the PDB Administrator Credentials section, specify the credentials of the Admin user account that you want to use to administer the PDB clone.
 - In the Container Database section, specify the destination CDB that is deployed in the public cloud setup (the CDB that the PDB clone must be a part of).
 - In the Credentials section, specify the SYSDBA credentials for the destination CDB, and the host credentials for the destination CDB.
6. If you do not need to specify anymore details, click **Clone**. This submits the deployment procedure to clone a PDB to a CDB that is deployed in a public cloud setup.

To specify other configuration details, mask data, as well as schedule the cloning process, click **Advanced**.

Follow the rest of the steps, if you have selected the Advanced option.

7. On the Clone from Oracle Cloud: Source and Destination page, verify the details, and then click **Next**.
8. On the Clone from Cloud: Configuration page, in the Database Files Location section, specify the storage location where the datafiles of the PDB clone must be stored.

In the Advanced Configuration section, specify the storage limits for the maximum size of the PDB clone, and the maximum size of a shared table space within the PDB clone. By default, no limits are placed on the values for these attributes.

In the Miscellaneous section, select the logging option that you want to use for the table spaces created within the PDB clone.

Click **Next**.

9. On the Clone from Cloud: Post Processing page, in the Data Masking section, specify the data masking definition that you want to apply after cloning the PDB. Data masking masks sensitive data in a database.

For information on how to create a data masking definition, see *Oracle Data Masking and Subsetting Guide*. Note that you can apply a data masking definition only if you have the Subset-Masking license pack.

In the Custom Scripts section, for **Pre Script** and **Post Script**, specify the Oracle Software Library components that contain the scripts that you want to run before cloning, and after cloning the PDB respectively. Also, for **SQL Script**, specify the SQL scripts that you want to run after cloning the PDB. For **Run As User**, select the user account that you want to use to run the SQL scripts.

Click **Next**.

10. On the Clone from Cloud: Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning deployment procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive email notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive email notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

11. On the Clone from Cloud: Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Clone** to submit the deployment procedure to clone the selected PDB to a CDB that is deployed in a public cloud setup.

Cloning a PDB from Oracle Cloud Using EM CLI

To clone a pluggable database on Oracle Cloud to an on-premise container database, enter the EM CLI verb `emcli pdb_clone_management -input_file=data:/xyz/sdf/pdb_clone.props`, where `pdb_clone.props` is the properties file which provides the cloning parameters and their values.

There are three methods in which you can clone a pluggable database. The difference between each of these methods is in the configuration of certain parameters in the properties file. The 3 methods and the details of the configuration parameters are explained below:

- Operating System (OS) Image backup

Takes a backup of the source PDB and creates a new PDB. The `BACKUP_TYPE` parameter should specify the type of backup. The allowed values for `BACKUP_TYPE` are `OSIMAGE`, `RMAN` and `TAR`. The `EXISTING_BACKUP` and `EXISTING_BACKUP_METADATA` parameters should not be provided.

Sample properties file:

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYCO
```

```

SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
BACKUP_TYPE=OSIMAGE

```

- Existing backup

Uses an existing backup of the source PDB and creates a new PDB. The BACKUP_TYPE parameter should specify the type of backup. The allowed values for BACKUP_TYPE are OSIMAGE, RMAN and TAR. The EXISTING_BACKUP parameter should specify the location with the backup name and EXISTING_BACKUP_METADATA should specify the location and the metadata file name for the backup.

Sample properties file:

```

SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SYC:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
EXISTING_BACKUP=/user1/pdbbackup/PDB1_Backup_14297779
EXISTING_BACKUP_METADATA=/user1/pdbbackup/PDB1_Backup_14297779/PDB1.xml
BACKUP_TYPE=RMAN

```

 Note:

To create a PDB backup, enter the verb `emcli pdb_backup -input-File="location of file containing properties required for taking backup of PDB"`, where the sample contents of the properties file is as follows:

```

TARGET_HOST_LIST=xyz.abccorp.com
HOST_NORMAL_NAMED_CRED=XYZ_CRED:CRED_OWNER
SRC_CDB_NAMED_CRED=CDB1_CRED:CRED_OWNER
SRC_CDB_TARGET_NAME=CDB1
SRC_CDB_TARGET_TYPE=oracle_database
SRC_PDB_TARGET_NAME=CDB1_PDB1
BACKUP_LOCATION=/user1/pdbbackup
WORK_DIR_LOCATION=/tmp
ORACLE_HOME_LOC=/scratch/d121hmcasm/product/12.1.0/dbhome_1

```

- Unplug/plug

Unplugs the source PDB and creates a new PDB at the destination using the unplugged source, and then plugs the source back. Both, EXISTING_BACKUP and BACKUP_TYPE parameters should not be provided.

Sample properties file:


```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
```



Note:

For all the 3 methods explained above, in case the destination PDB data files location is ASM then add the parameter `DEST_STAGE_DIR` whose value will be used as the destination while transferring the source PDB data files. This parameter is optional, if it is not provided a temporary directory will be used. For Linux systems the temporary directory is `/tmp`.

Cloning Schema(s) from Oracle Cloud to a DB or PDB

You can clone a schema that is on Oracle Cloud to on-premise either as a database or as a pluggable database using EM CLI verbs. To do so, follow the steps below:

1. Enter the EM CLI verb `emcli describe_dbprofile_input -data_mode=EXPORT`. The output provides all profile creation input variables.



Note:

Export is supported only for database and schema whereas import is supported for both database and PDBs. This indicates that the source should always be a database or schema and the destination can either be a database or a PDB.

2. Use the input variables to create a properties file with values for all the variables.
3. Create the directory object `SCHEMAS_DUMP_DIR`.
4. Export data from the source database by creating a database profile. To do so, enter the verb `emcli create_dbprofile - input_file=data:<properties file name along with path>`.



Note:

Use the properties file created in the previous step for this verb.

Sample properties file:

```
#-----#
# SOURCE                                     #
#-----#
```

```

REFERENCE_DATABASE=SS_TM_DB
REFERENCE_DATABASE_TYPE=oracle_database
REF_DB_CREDENTIALS=CRED_DB:sysman
REF_HOST_CREDENTIALS=CRED_HOST:sysman

#-----#
# DATA CONTENT DETAILS                                     #
#-----#
DATA_CONTENT_MODE=EXPORT
DATA_CONTENT=METADATA_AND_DATA

#-----#
# EXPORT DETAILS                                           #
#-----#
EXPORT.EXPORT_TYPE=SELECTED_SCHEMAS
EXPORT.SCHEMA_INCLUDE_LIST.0=HR
EXPORT.SCHEMA_INCLUDE_LIST.1=PM
EXPORT.SCHEMA_INCLUDE_LIST.2=OE
EXPORT.SCHEMA_INCLUDE_LIST.3=IX
EXPORT.SCHEMA_INCLUDE_LIST.4=SH
EXPORT.SCHEMA_INCLUDE_LIST.5=BI
EXPORT.DEGREE_OF_PARALLELISM=1
EXPORT.DUMP_DIRECTORY_LIST.0=directory=SCHEMAS_DUMP_DIR,file_name=samplsche-
mas.dmp,max_size=100
EXPORT.LOG_FILE_DIRECTORY=directory=SCHEMAS_DUMP_DIR,file_name=samplschemas.log

#-----#
# PROFILE DETAILS                                         #
#-----#
PROFILE_NAME=Export Dump of Sample schemas10
PROFILE_VERSION=11.2.0.4.0
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/
WORKING_DIRECTORY=/tmp

```

5. Enter the verb to transfer data: `emcli data_transfer -input_file=data:/u01/files/data_trans.props.`

Sample properties file:

```

#-----#
# SOURCE                                                  #
#-----#
SRC_HOST_CREDS=NC_HOST_SRAY
SOURCE_LOCATION=/tmp/newp/PDB_Backup_1428003803938
SRC_HOST=b12.idc.example.com

#-----#
# DESTINATION                                             #
#-----#
DEST_HOST_CREDS=NC_HOST_SRAY
DEST_LOCATION=/scratch/sray/app3/sray/oradata/migda
DEST_HOST=slo.us.example.com

#-----#
# HYBRID GATEWAY / FORWARDER                             #
#-----#
FORWARDER_HOST=slo.us.example.com
FORWARDER_CRED=ACD_NY:SYSCO
WORKING_DIRECTORY=/tmp

```

 Note:

Remove the Hybrid Gateway parameters if the SSH connection exists between the source and the destination hosts.

6. Enter the verb to import data in to the destination database: `emcli dbimport -input_file=data:/u01/files/dbimport.props.`

 Note:

To clone the destination to database or pluggable database, ensure you provide the required value in the `DESTINATION_TARGET_TYPE` option in the properties file. For database, enter `oracle_database`, and for PDB enter `oracle_pdb`.

Sample properties file:

```
#-----#
#          DESTINATION          #
#-----#
DESTINATION_TARGET=SS_OPC_DB
DESTINATION_TARGET_TYPE=oracle_database
DATABASE_CREDENTIAL=CRED_DB:sysman
HOST_NAMED_CREDENTIAL=CRED_HOST:sysman

#-----#
#          PROFILE              #
#-----#
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/Export Dump
of Sample schemas10

#-----#
#          SCHEMA DETAILS      #
#-----#
REMAP_SCHEMA_LIST.0=HR:HR
REMAP_SCHEMA_LIST.1=OE:OE
REMAP_SCHEMA_LIST.2=PM:PM
REMAP_SCHEMA_LIST.3=IX:IX
REMAP_SCHEMA_LIST.4=SH:SH
REMAP_SCHEMA_LIST.5=BI:BI
REMAP_TABLESPACE_LIST.0=EXAMPLE:MYTBSP1
REMAP_TABLESPACE_LIST.1=USERS:MYTBSP1
REMAP_TABLESPACE_LIST.2=SYSTEM:MYTBSP1
DEGREE_OF_PARALLELISM=1
DUMP_FILE_LIST.0=/scratch/ae/dumpdir/sampschemas.dmp
IMPORT_LOG_FILE_DIRECTORY=DATA_PUMP_DIR
```

Cloning a DB from Oracle Cloud to an On-Premise PDB

You can clone a database that is on Oracle Cloud to on-premise as a pluggable database using EM CLI verbs. To do so, follow the steps below:

1. Enter the EM CLI verb `emcli describe_dbprofile_input -data_mode=EXPORT`. The output provides all profile creation input variables.

**Note:**

Export is supported only for database and schema whereas import is supported for both database and PDBs. This indicates that the source should always be a database or schema and the destination can either be a database or a PDB.

2. Use the input variables to create a properties file with values for all the variables.
3. Create the directory object SCHEMAS_DUMP_DIR.
4. Export data from the source database by creating a database profile. To do so, enter the verb `emcli create_dbprofile - input_file=data:<properties file name along with path>`.

**Note:**

Use the properties file created in the previous step for this verb.

Sample properties file:

```
#-----#
# SOURCE                                     #
#-----#
REFERENCE_DATABASE=SS_TM_DB
REFERENCE_DATABASE_TYPE=oracle_database
REF_DB_CREDENTIALS=CRED_DB:sysman
REF_HOST_CREDENTIALS=CRED_HOST:sysman

#-----#
# DATA CONTENT DETAILS                     #
#-----#
DATA_CONTENT_MODE=EXPORT
DATA_CONTENT=METADATA_AND_DATA

#-----#
# EXPORT DETAILS                           #
#-----#
EXPORT.EXPORT_TYPE=FULL_DATABASE
EXPORT.DEGREE_OF_PARALLELISM=1
EXPORT.DUMP_DIRECTORY_LIST.0=directory=SCHEMAS_DUMP_DIR,file_name=samplsche-
mas.dmp,max_size=100
EXPORT.LOG_FILE_DIRECTORY=directory=SCHEMAS_DUMP_DIR,file_name=samplschemas.log

#-----#
#          PROFILE DETAILS                  #
#-----#
PROFILE_NAME=Export Dump of Sample schemas10
PROFILE_VERSION=11.2.0.4.0
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/
WORKING_DIRECTORY=/tmp
```

5. Enter the verb to transfer data: `emcli data_transfer -input_file=data:/u01/files/data_trans.props`.

Sample properties file:

```
#-----#
# SOURCE                                     #
#-----#
SRC_HOST_CREDS=NC_HOST_SRAY
SOURCE_LOCATION=/tmp/newp/PDB_Backup_1428003803938
SRC_HOST=b12.idc.example.com

#-----#
# DESTINATION                               #
#-----#
DEST_HOST_CREDS=NC_HOST_SRAY
DEST_LOCATION=/scratch/sray/app3/sray/oradata/migda
DEST_HOST=slo.us.example.com

#-----#
# HYBRID GATEWAY / FORWARDER               #
#-----#
FORWARDER_HOST=slo.us.example.com
FORWARDER_CRED=ACD_NY:SYSCO
WORKING_DIRECTORY=/tmp
```

 Note:

Remove the Hybrid Gateway parameters if the SSH connection exists between the source and the destination hosts.

6. Enter the verb to import data in to the destination database: `emcli dbimport -input_file=data:/u01/files/dbimport.props.`

 Note:

To clone the destination to database or pluggable database, ensure you provide the required value in the `DESTINATION_TARGET_TYPE` option in the properties file. For database, enter `oracle_database`, and for PDB enter `oracle_pdb`.

Sample properties file:

```
#-----#
#          DESTINATION                      #
#-----#
DESTINATION_TARGET=SS_OPC_DB
DESTINATION_TARGET_TYPE=oracle_database
DATABASE_CREDENTIAL=CRED_DB:sysman
HOST_NAMED_CREDENTIAL=CRED_HOST:sysman

#-----#
#          PROFILE                          #
#-----#
PROFILE_LOCATION=Database Provisioning Profiles/12.1.0.1.0/linux_x64/Export Dump
of Sample schemas10

#-----#
#          SCHEMA DETAILS                   #
#-----#
```

```
#-----#
REMAP_SCHEMA_LIST.0=HR:HR
REMAP_SCHEMA_LIST.1=OE:OE
REMAP_SCHEMA_LIST.2=PM:PM
REMAP_SCHEMA_LIST.3=IX:IX
REMAP_SCHEMA_LIST.4=SH:SH
REMAP_SCHEMA_LIST.5=BI:BI
REMAP_TABLESPACE_LIST.0=EXAMPLE:MYTBSP1
REMAP_TABLESPACE_LIST.1=USERS:MYTBSP1
REMAP_TABLESPACE_LIST.2=SYSTEM:MYTBSP1
DEGREE_OF_PARALLELISM=1
DUMP_FILE_LIST.0=/scratch/ae/dumpdir/sampleschemas.dmp
IMPORT_LOG_FILE_DIRECTORY=DATA_PUMP_DIR
```

Cloning a DB from Oracle Cloud (Oracle Compute Service)

To clone a database from Oracle Cloud, you can use one of the following solutions:

- [Cloning a DB from Oracle Cloud Using the Clone Wizard](#)
- [Cloning a DB from Oracle Cloud Using EMCLI](#)

Cloning a DB from Oracle Cloud Using the Clone Wizard

To clone a database from Oracle Cloud, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Select the database to be cloned and right click the mouse button.
4. From the **Oracle Database** menu, select **Cloning**, then select **Clone to Oracle Cloud**.
5. On the Source and Destination: Clone to Oracle Cloud page, do the following:
 - In the Source Credentials section, specify or search and select the Oracle Cloud database and Oracle Home credentials. Click OK.
 - In the Data Time Series section, the Prior Point in Time option is selected by default and a timeline of RMAN Backup images appear. Select the appropriate RMAN backup by clicking Select on the image.
 - In the Source Backup Location field, the location of the source backup you have selected appears.
 - In the Destination section, in the Clone to field, the Select Compute Cloud Service is selected by default. This indicates that you are requesting resources on Oracle Public Cloud to clone your on-premise database.
 - In the Database Definition section, specify the Display Name for the database, the Global Database Name, SID.
 - In the Hosts section, you can choose:
 - **Clone Oracle Home:** If the Oracle Home is not pre-installed on the on-premise host on which database is being cloned, select the **Clone Oracle Home** option.
 - **Use an Existing Home:** If the Oracle Home is already cloned, select the **Use an Existing Home** option.

- If you select the Clone Oracle Home option, select destination machine on which the Oracle Home is to be provisioned. Specify the following:
 - **Database Host Credentials:** Specify the database host credentials.
 - **Run root scripts that require root privileges:** You can either specify the root credentials or choose to enter them later to run the root scripts.

In the Stage Location field, the location on the destination host machine where the backups will be staged before cloning is displayed.
 - If you select the Use an Existing Home option, specify the following:
 - **Oracle Home Location:** Select the location of the Oracle Home on an Oracle Public Cloud machine already managed in Enterprise Manager. The host name of the machine in which the Oracle Home is present is displayed.
 - **Database Host Credentials:** Specify the database host credentials.

In the Stage Location field, the location on the Oracle Public Cloud machine where the backups will be staged before cloning is displayed.
 - **Hybrid Gateway Server:** Specify the host name and credentials for the Hybrid Gateway Server (configured Enterprise Manager Management Agent) which transfers data from Oracle Cloud to Enterprise Manager. Click **Next**. The Configuration page appears:
6. Specify the following details:
- **Software Location:** Specify the Oracle Base and Oracle Home location where the database is to be cloned.
 - **Oracle Database User Groups:** Specify the name of the Oracle Database User Group performing this cloning operation.
 - **Database Files Location:** Specify the location where you want the data files, temp files, redo log files, and control files to be created.
 - **Recovery Files Location:** Specify the location where you want the recovery files, such as archived redo logs, RMAN backups, and other related files to be created. You can choose to use the fast recovery area. If you do, specify the fast recovery area size. In the Database Mode section, you can enable Read Only.
 - **Listener Configuration:** In the Hosts section on the previous page, if you have selected:
 - **Clone Oracle Home:** Specify the Listener Name and Port to create a new listener with which the database will be registered.
 - **Use an Existing Home:** Select the listener targets running under the new Oracle Home, to register the clone database, You can add new listeners by clicking Add, and specifying the listener name and port.
 - **Database Credentials:** Specify passwords for the SYS, SYSTEM, and DBSNMP administrative users in the clone database. You can choose to have the same password for all the three users or a different password for each.
 - **Transparent Data Encryption:** If the source database is encrypted, specify the location in which the Wallet is to be placed on the Oracle Public Host machine. Select the Use same password as source checkbox to use the same password as the host database. Select the Reset/Regenerate the Master Encryption Key to reset the encryption key.

- **Parallelism:** Setting a higher Degree of Parallelism will ensure faster recompilation of objects. The default setting is the number of CPU count of the host. Click **Next**.
- 7. On the Initialization Parameters page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click Edit to modify the value of the parameter. Click **Next**.
- 8. On the Post Processing page, you can select masking rules, subsetting rules, or custom scripts such as pre script, post script, and post SQL script. Click **Next**.
- 9. On the Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning deployment procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive e-mail notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

- 10. On the Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Submit** to launch the procedure to clone the database, transfer the selected backups to the destination host via the hybrid gateway. Once cloned, this database can be managed as an Enterprise Manager target.

To refresh the cloned database, from the **Oracle Database** menu, select **Cloning**, then select **Clone Management**. On the Clone Management Dashboard, select the database and click **Refresh**. Accept the default values in the wizard and optionally change the Initialization Parameters and Configuration details and click **Refresh** to refresh the database

Cloning a DB from Oracle Cloud Using EMCLI

You can clone a database from Oracle Cloud using the following EMCLI command:

```
emcli db_clone_management -cloneFromCloud -input_file=data:/xyz/sdf/clone-
fromcloud.props
where clonefromcloud.props is the properties file which provides the cloning param-
eters and their values.
```

Sample Properties File

```
#-----#
#                               Source Details                               #
#-----#
SRC_DB_TARGET_NAME=hrdb.example.com
SRC_DB_TARGET_TYPE=oracle_database
SRC_HOST_NORMAL_NAMED_CRED=OPC_209:SYSMAN
SRC_DB_CRED=CLOUDDB_CRED:SYSMAN
SRC_WORK_DIR=/tmp
PROFILE_INTERNAL_PATH=Database Provisioning Profiles/11.2.0.4.0/linux_x64/Database
db11.example.com Profile 22-02-2016 05:13 PM
PROFILE_VERSION_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
nent:SUB_DatabaseConfigProfile:2C5B357D5FF67E2AE053F17CB10A53F1:0.1
SOURCE_BACKUP_LOCATION=/u01/hrdb/fast_recovery_area/hrdb/backup1458714912986

#-----#
```



```
#           Hybrid Gateway Details
#-----#
FORWARDER_HOST=slc04wim.example.com
FORWARDER_HOST_CREDS=HOST:SYSMAN
FORWARDER_STAGE_LOCATION=/scratch/fwd

#-----#
# Destination Details
#-----#
TARGET_HOST_LIST=slc01aia.example.com
HOST_NORMAL_NAMED_CRED=SLC01AIA_CRED:SYSMAN
HOST_PRIV_NAMED_CRED=SLC01AIA_CRED_ROOT:SYSMAN
DEST_LOCATION=/scratch/backupStage
DEST_WORK_DIR=/tmp
#-----#
# Database Definition
#-----#
COMMON_DB_SID=hrcln
COMMON_DOMAIN_NAME=xyz.example.com
COMMON_GLOBAL_DB_NAME=hrcln.example.com
DATABASE_TYPE=dbTypeSI
DB_ADMIN_PASSWORD_SAME=true
DATABASE_PASSWORDS=welcome1

#-----#
# Software Configuration
#-----#
#All these are only Applicable if a new oracle home provisioning needs to be done.
CLONE_HOME=Y
OINSTALL_GROUP=dba
OSBACKUPDBA_GROUP=dba
OSDBA_GROUP=dba
OSDGDBA_GROUP=opc
OSKMDBA_GROUP=opc
OSOPER_GROUP=opc
ORACLE_BASE_LOC=/scratch/app/dbhome11204
ORACLE_HOME_LOC=/scratch/app/dbhome11204/product/11.2.0/dbhome_1

#-----#
# Database Configuration
#-----#
ARCHIVE_LOG_MODE=YES
DB_FILE_LOC=/scratch/app/dbhome11204/oradata
DB_TARGET_DISPLAY_NAME=HRCLN
DB_TARGET_NAME=hrcln.example.com
FLASH_REC_AREA=/scratch/app/dbhome11204/fast_recovery_area
FLASH_REC_AREA_SIZE=4182
LISTENER_PORT=2887
//Only needed if we are provisioning a new home
LISTENER_NAME=LIST_2887

#-----#
# Data file transfer configuration parameters
#-----#
# This will be moved to emctl metadata. for now its a variable in DP
NUMBER_OF_THREADS=10
# This will be moved to emctl metadata. for now its a variable in DP
# The chunk size in bytes. For now default is 250MB
MIN_CHUNK_SIZE=262144000
#Transfer Unit can be FILE or CHUNK
```

```
TRANSFER_UNIT=FILE
#Maximum speed in Kbps
BANDWIDTH_LIMIT=2048
#Time out in seconds for data transfer
TIME_OUT=1200

#-----#
#   Wallet Configuration
#-----#
DEST_PATH_TO_WALLET=/u03/wallet
WALLET_PASSWORD=welcome123
NEW_WALLET_PASSWORD=welcome1
REKEY_REQUIRED=true
```

Cloning Within Oracle Cloud

To clone a database or a pluggable database within Oracle Cloud, refer to the following use cases:

- [Cloning a PDB Within Oracle Cloud](#)
- [Cloning a DB Within Oracle Cloud \(Oracle Compute Service\)](#)

Cloning a PDB Within Oracle Cloud

To clone a pluggable database within Oracle Cloud, you can use either of the following solutions:

- [Creating a Full Clone Pluggable Database Using the Clone Wizard](#)
- [Creating a Full Clone Pluggable Database Using EM CLI.](#)

Creating a Full Clone Pluggable Database Using the Clone Wizard

You can create a full clone of a PDB using the new Clone PDB Wizard. To create a full clone PDB, follow these steps:

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source CDB (the CDB that the source PDB is a part of) in the list, then click the name of the PDB that you want to clone.
4. From the **Oracle Database** menu, select **Cloning**, then select **Create Full Clone**.
Alternatively, in Step 3, you can right click the name of the PDB that you want to clone, select **Oracle Database**, select **Cloning**, then select **Create Full Clone**.
5. On the Source and Destination: Create Full Clone Pluggable Database page, do the following:
 - Specify the SYSDBA credentials for the source CDB. You can choose to use the preferred credentials, use a saved set of named credentials, or specify a new set of credentials.
 - In the Pluggable Database Destination section, specify a name, and a display name for the PDB clone. Enterprise Manager uses the display name to identify the PDB clone target.

- Select the Refreshable clone check box to create a manually Refreshable PDB.

 Note:

- Refreshable clone feature is only applicable to Oracle Database version 12.2 and above.
- To create a Refreshable clone you have to select a different Container Database (CDB). You will not be able to create one in the same CDB.
- The Refreshable PDB clone is created in read-only mode and can only be either in read-only or mounted state. The clone cannot be opened in read-write mode. Therefore, data masking and post SQL scripts cannot be run on the clone.

- In the PDB Administrator Credentials section, specify the credentials of the admin user account that you want to use to administer the PDB clone.
- To clone the PDB to a CDB different from the source CDB, select **Clone the Pluggable Database into a different Container Database**, then specify the destination CDB.
- In the Credentials section, specify the destination CDB host credentials. If you chose to clone the PDB to a CDB different from the source CDB, specify the SYSDBA credentials for the destination CDB. Also, if the destination CDB is using Automatic Storage Management (ASM) to manage disk storage, you must specify the ASM credentials.

 Note:

For Database versions 12.2 and above, if both the source and destination Container Databases are in Archive Log mode and have Local Undo configured, then a Hot clone is performed. Hot clone is the ability to clone a Pluggable Database while the source is open in read-write mode.

6. You can specify other configuration details, mask data, as well as schedule the cloning process, by selecting the **Advanced** option. Follow the rest of the steps, if you have selected the Advanced option. If you do not need to specify anymore details, click **Clone**. This submits the deployment procedure to clone a PDB to a CDB that is deployed in a public cloud setup.
7. On the Create Full Clone Pluggable Database: Source and Destination page, verify the details specified, and then click **Next**.
8. On the Create Full Clone Pluggable Database: Configuration page, do the following:
 - In the Database Files Location section, specify the storage location where the datafiles of the PDB clone must be stored. If the destination CDB is using ASM to manage disk storage, specify the disk group where the datafiles of the PDB clone must be stored.

- To ensure that only the source PDB data model definition is cloned (and the source PDB data is not cloned), select **Exclude User Data**.
- In the Advanced Configuration section, specify the storage limits for the maximum size of the PDB clone, and the maximum size of a shared tablespace within the PDB clone. By default, no limits are placed on the values for these attributes.
- In the Miscellaneous section, select the logging option that you want to use for the tablespaces created within the PDB clone.

Click **Next**.

9. On the Create Full Clone Pluggable Database: Post Processing page, do the following:

- In the Data Masking section, specify the data masking definition that you want to apply after cloning the PDB. Data masking masks sensitive data in a database.

For information on how to create a data masking definition, see *Oracle Data Masking and Subsetting Guide*. Note that you can apply a data masking definition only if you have the Subset-Masking license pack.

- In the Custom Scripts section, for **Pre Script** and **Post Script**, specify the Oracle Software Library components that contain the scripts that you want to run before cloning, and after cloning the PDB respectively. Also, for **SQL Script**, specify the SQL scripts that you want to run after cloning the PDB. For **Run As User**, select the user account that you want to use to run the SQL scripts.

Click **Next**.

10. On the Create Full Clone Pluggable Database: Schedule page, specify an instance name for the cloning deployment procedure. Also, specify the point in time when you want the cloning procedure to begin.

In the Notification section, select the deployment procedure states for which you want to receive e-mail notifications. For example, if you select **Scheduled** and **Succeeded** for **Status for Notification**, you will receive e-mail notifications when the cloning deployment procedure is scheduled, and when it succeeds.

Click **Next**.

11. On the Create Full Clone Pluggable Database: Review page, review all the details you provided. If you want to edit certain details, click **Back** to navigate to the required page.

Click **Clone** to submit the deployment procedure to create a full clone of the source PDB.

Creating a Full Clone Pluggable Database Using EM CLI

To create a full clone of a pluggable database, execute the verb `emcli pdb_clone_management -input_file=data:/xyz/sdf/pdb_clone.props`, where `pdb_clone.props` is the properties file.

Sample properties file (pdb_clone.props):

```
SRC_PDB_TARGET=cdb_prod_PDB
SRC_HOST_CREDS=NC_HOST_SCY:SYCO
SRC_CDB_CREDS=NC_HOST_SYC:SYCO
```

```
SRC_WORK_DIR=/tmp/source
DEST_HOST_CREDS=NC_SLC0_SSH:SYS
DEST_LOCATION=/scratch/sray/app/sray/cdb_tm/HR_TM_PDB6
DEST_CDB_TARGET=cdb_tm
DEST_CDB_TYPE=oracle_database
DEST_CDB_CREDS=NC_HOST_SYC:SYCO
DEST_PDB_NAME=HR_TM_PDB6
```

 Note:

- If you are creating a full clone PDB on Oracle cloud, wherein both the destination PDB and the source PDB are in different CDBs, while both the CDBs are on Oracle Cloud, then ensure that the source PDB is in read-write mode. This is necessary since a database link is created in the destination CDB for cloning the PDB, and a temporary user is created in the source PDB for using the database link. If there is an existing database link in the destination CDB that connects to the source PDB, then use the parameter `EXISTING_DB_LINK_NAME` to provide the database link name in the properties file.
- If the PDB full clone (on-premise) needs to be performed using a backup, then the following parameters should be added in the properties file.

```
BACKUP_TYPE=RMAN
DEST_STAGE_DIR=/tmp/bkp_common
```

If the temporary backup location is shared across source and destination, then add the `SKIP_DATA_TRANSFER` parameter and set it to 'Y'. For example, `SKIP_DATA_TRANSFER=Y`. This indicates that the data transfer step should be skipped in the clone procedure.

Cloning a DB Within Oracle Cloud (Oracle Compute Service)

To clone a database within Oracle Cloud, you can use either of the following solutions:

- [Cloning a DB Within Oracle Cloud Using the Clone Wizard](#)
- [Cloning a DB Within Oracle Cloud Using EMCLI](#)

Cloning a DB Within Oracle Cloud Using the Clone Wizard

To clone a database from Oracle Cloud using the Clone Wizard, follow the steps as described in [Cloning a DB to Oracle Cloud Using the Clone Wizard](#).

 Note:

The Hybrid Cloud Gateway server is not required as the database is being cloned within the Oracle Cloud.

Cloning a DB Within Oracle Cloud Using EMCLI

To clone a database within Oracle Cloud using EMCLI, enter the following command:

```
emcli db_clone_management -cloneToCloud -target_name="database target
name" -target_type="database target type - <oracle_database | rac_data-
base>" -input_file="path of the input file"
```

Sample Properties File

```
#-----#
#                               Source Details
#-----#
SRC_DB_TARGET_NAME=dbprod.example.com
SRC_DB_TARGET_TYPE=oracle_database
SRC_HOST_NORMAL_NAMED_CRED=SRCHOSTCRED:SYSMAN
SRC_DB_CRED=DBPROD_CRED:SYSMAN
SRC_WORK_DIR=/tmp
PROFILE_INTERNAL_PATH=Database Provisioning Profiles/11.2.0.4.0/linux_x64/Database
db11.example.com Profile 22-02-2016 05:13 PM
PROFILE_VERSION_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
nent:SUB_DatabaseConfigProfile:2C5B357D5FF67E2AE053F17CB10A53F1:0.1
SOURCE_BACKUP_LOCATION=/scratch/dbprod/app/dbprod/fast_recovery_area/back-
up1456141446021

#-----#
#                               Hybrid Gateway Details
#-----#
FORWARDER_HOST=slc04wim.example.com
FORWARDER_HOST_CREDS=HOST:SYSMAN
FORWARDER_STAGE_LOCATION=/scratch/fwd

#-----#
# Destination Details
#-----#
TARGET_HOST_LIST=129.124.22.34
HOST_NORMAL_NAMED_CRED=OPC_CRED:SYSMAN
HOST_PRIV_NAMED_CRED=HOST_CRED_ROOT:SYSMAN
DEST_LOCATION=/u03/backup
DEST_WORK_DIR=/tmp
#-----#
# Database Definition
#-----#
COMMON_DB_SID=Dcln1
COMMON_DOMAIN_NAME=xyz.example.com
COMMON_GLOBAL_DB_NAME=Dcln1.example.com
DATABASE_TYPE=dbTypeSI
DB_ADMIN_PASSWORD_SAME=true
DATABASE_PASSWORDS=welcome1

#-----#
# Software Configuration
#-----#
CLONE_HOME=Y # mention it as N if only database cloning is required
#All these below parameters are only required if a new oracle home provisioning
needs to be done.
OINSTALL_GROUP=dba
OSBACKUPDBA_GROUP=dba
OSDBA_GROUP=dba
OSGDBA_GROUP=opc
OSKMDBA_GROUP=opc
OSOPER_GROUP=opc
ORACLE_BASE_LOC=/u03/home/app/gaurav
ORACLE_HOME_LOC=/u03/home/app/gaurav/product/11.2.0/dbhome_1
```

```
#-----#
# Database Configuration
#-----#
ARCHIVE_LOG_MODE=YES
DB_FILE_LOC=/u03/home/oradata
DB_TARGET_DISPLAY_NAME=Dcln1
DB_TARGET_NAME=Dcln1.example.com
FLASH_REC_AREA=/u03/home/app/home/fast_recovery_area
FLASH_REC_AREA_SIZE=4182
LISTENER_PORT=2887
#Only needed if we are provisioning a new home
LISTENER_NAME=LIST_2887

#-----#
# Data file transfer configuration parameters
#-----#
# This will be moved to emctl metadata. for now its a variable in DP
NUMBER_OF_THREADS=10
# This will be moved to emctl metadata. for now its a variable in DP
# The chunk size in bytes. For now default is 250MB
MIN_CHUNK_SIZE=262144000
#Transfer Unit can be FILE or CHUNK
TRANSFER_UNIT=FILE
#Maximum speed in Kbps
BANDWIDTH_LIMIT=2048
#Time out in seconds for data transfer
TIME_OUT=1200

#-----#
# Wallet Configuration
#-----#
DEST_PATH_TO_WALLET=/u03/wallet
WALLET_PASSWORD=welcome123
NEW_WALLET_PASSWORD=welcome1
REKEY_REQUIRED=true
```


Hybrid Solution: Oracle Cloud Platform Management

This chapter describes how you can monitor and manage database cloud services from the Enterprise Manager Self Service Portal. It contains the following sections:

- [About Hybrid Cloud Management](#)
- [Viewing and Monitoring Database Cloud Services](#)
- [Typical Hybrid Cloud Management Workflow](#)
- [Setting Up Hybrid Cloud](#)
- [Managing Data Sources](#)
- [Defining the Oracle Database Cloud Service Settings](#)
- [Discovering the Database Cloud Service](#)
- [Monitoring the Discovered Database Cloud Service](#)
- [Refreshing the Cloud Database Service](#)
- [Clone Management](#)
- [Defining the Database Backup Cloud Service Settings](#)
- [Defining the Oracle Compute Service Settings](#)
- [Discovering the Oracle Compute Service](#)

About Hybrid Cloud Management

Oracle Enterprise Manager delivers comprehensive capabilities to manage, migrate, test and deploy databases across a hybrid cloud, providing complete reuse of existing IT operations management skills and practices. Self service users can now monitor and manage all their Oracle Cloud resources along with on-premises artifacts using the Enterprise Manager Self Service Portal.



Viewing and Monitoring Database Cloud Services

Enterprise Manager Hybrid Cloud allows you to on-board database cloud services to your existing management infrastructure and monitor as needed.


To view and monitor existing database cloud services in Enterprise Manager, follow these steps:

1. Open your web browser and enter <http://cloud.oracle.com>
2. Click **Sign In** and in the Cloud Account drop down, select a data center and click **My Services**.
3. Enter your identity domain, user name (user ID), and password.
You will see the Cloud Dashboard with a list of database cloud service instances.
4. Click on a database instance name to view the details.
Note the database cloud instances and their details listed here.
5. The next step is to on-board your database cloud services to Enterprise Manager. This involves registering the cloud account and discovering the database cloud services in Enterprise Manager.
 - See [Setting Up Hybrid Cloud Using the Hybrid Setup Console](#)
 - See [Discovering the Database Cloud Service](#)
6. You can then view and monitor the database cloud service. See [Monitoring the Discovered Database Cloud Service](#).

Typical Hybrid Cloud Management Workflow

The hybrid cloud solution allows you to refresh and clone databases in Oracle Cloud from the Enterprise Manager Self Service Portal.

A typical hybrid cloud management workflow is shown below:



To refresh or clone an Oracle Cloud database, follow these steps:

Task	Using Graphical User Interface	Using REST API
Create a clone of your production database and mask sensitive data.	See Managing Data Sources	

Task	Using Graphical User Interface	Using REST API
The cloned database can be automatically refreshed at specified intervals and backed up on the OSS. You can also create backups on demand.	See Creating Backups On the Oracle Storage Server	
View and monitor the discovered database cloud service	See Monitoring the Discovered Database Cloud Service	
Refresh the data from the backup.	See Refreshing the Cloud Database Service	
Clone the database cloud service instance.	See Clone Management	

Setting Up Hybrid Cloud

You can setup hybrid cloud either from the Hybrid Setup Console or by using a single REST API or EMCLI command. This section describes the different methods that can be used.

Topics

- [Setting Up Hybrid Cloud Using the Hybrid Setup Console](#)
- [Setting Up Hybrid Cloud Using REST API](#)
- [Setting Up Hybrid Cloud Using EMCLI](#)

Setting Up Hybrid Cloud Using the Hybrid Setup Console

The Hybrid Cloud Setup Console allows you to register your identity domains and set up your cloud accounts in Enterprise Manager. To access this console, log in to Enterprise Manager as a user with CLOUD_ADMINISTRATOR privileges. From the **Setup** menu, click **Hybrid Cloud** to navigate to the console pages. This section covers the following:

- Registering Identity Domains and Cloud Accounts. See [Registering the Identity Domain and the Cloud Account](#).
- Configuring Hybrid Cloud Gateways. See [Configuring the Hybrid Cloud Gateway](#).
- Configuring Proxy Settings. See [Configuring the Proxy Settings](#).
- Adding SSH Key Credentials. See [Configuring the SSH Keys](#).

Note:

You can perform all these operations using a single EMCLI command. See [Setting Up Hybrid Cloud Using EMCLI](#) for details.

Registering the Identity Domain and the Cloud Account

An identity domain controls the accounts of users who need access to service instances. It also controls the features that authorized users can access. A service instance belongs to an identity domain.

To register an identity domain and a cloud account, follow these steps:

1. From the **Setup** menu, click **Hybrid Cloud**, then click **Identity Domains**.
2. Click **Register**.
The Register Identity Domain dialog box appears.
3. Enter the name of the identity domain, select the data center and click **OK**. Once registered you will see a confirmation pop-up window click **OK**.
4. After registering the identity domain, you can now register the cloud account. This cloud account will be used for discovering and provisioning the database service instances.
5. Click **Accounts**. The Register Account dialog box appears.
6. Enter a name for the cloud account and select an identity domain that has been configured from the drop down list.
7. Enter the URL with which the cloud account can be tested. To obtain this URL, open your Web browser and enter <http://cloud.oracle.com>. On the Cloud Dashboard, click on the Overview tab on the left column, and copy the REST Endpoint. Add `/paas/service/dbcs/api/v1.1/instances/` and then add the Identity Domain. For example, `https://dbaas.oraclecloud.com/paas/service/dbcs/api/v1.1/instances/abcd/` where abcd is the Identity Domain.
8. Enter the user name and password for the Oracle Cloud Credentials and click **Register**.

You will see the registered account in the Accounts page.

Configuring the Hybrid Cloud Gateway

The hybrid cloud gateway is a specially configured management agent that provides a communication channel between the Oracle Cloud virtual hosts and the Enterprise Manager OMS. You can add a gateway that has already been registered in Enterprise Manager.



Note:

Configuring the hybrid gateway is an optional task. It is not required if you have configured Site-to-Site VPN and/or FastConnect from Oracle Network Cloud Services.

To add a hybrid cloud gateway, follow these steps:

1. Click **Add** on the Hybrid Gateway page .
2. In the Select Hybrid Cloud Gateways window, choose the gateway to be used for hybrid cloud and click **Select**. You can add multiple gateways here but the gate-

way that has least number of connections to the hosts on Oracle Cloud will be used during discovery.

The newly added gateway will appear on the Hybrid Cloud Gateway page and will be automatically enabled. If you disable the gateway, it will not be used for communicating with the Oracle Cloud virtual hosts.

Configuring the Proxy Settings

The proxy server acts as an intermediary for managing HTTP and HTTPS requests across network boundaries or firewalls. Using a proxy server, you can expose only certain ports for communication between the OMS and the hybrid cloud agent thus making the communication more secure and reliable. If you have configured site-to-site VPN or FastConnect from Oracle Network Cloud Services, this proxy server is not required.

You can select:

- No proxy: No proxy server is used.
- Use My Oracle Support proxy: The already configured My Oracle Support proxy server is used.
- Use a different proxy: Enter the following details to configure the proxy server:
 1. Enter the name of the host on which the proxy server is running.
 2. Enter the port number for the host.
 3. Select the HTTP or HTTPS protocol to be used for communication.

Click **Apply** to save the proxy server settings.

Configuring the SSH Keys

The hybrid cloud gateway agent communicates with the virtual hosts on Oracle Cloud over a secure SSH tunnel. Hence, SSH key based credentials are required for all monitoring and management operations.

To configure the SSH Keys, follow these steps:

1. Click **Add** in the SSH Keys page.

The Add SSH Key Credential dialog box appears.
2. Select a Named Credential from the list or create a New Credential and click **OK**.

Note:

If you are creating a New Credential, you must specify both the Public and Private key. If you are configuring the SSH Keys for a root user, select Sudo in the Run Privilege drop down and enter Root in the Run As field.

The configured SSH Key Credential is displayed on the SSH Keys page.

3. Selected a SSH Key Credential from the list and click **Grant** to grant specific privileges (View, Manage, or Full) to a user role.

Setting Up Hybrid Cloud Using REST API

You can set up all hybrid cloud operations such as creating identity domains, registering cloud accounts and hybrid gateways, configuring proxy settings, and so on using a single command.

Oracle Cloud Setup Data Model

This data model has the media type `application/json`. The following table describes the cloud setup data model.

Field	Type	Description	Occurs
domainName	String	Display name of identity domain.	1
dataCenterId	String	Datacenter identifier	1
userName	String	User name for cloud account	1
password	String	Password for cloud account	1
accountName	String	Displays the account name	1
urlToTest	String	Test the account details against this URL before saving.	1
skipTest	String	Provides an option to skip the URL test.	1
agentTargetGuid	String	Displays the GUID of available agent	1
proxyType	String	Type of the proxy. The possible values are NO_PROXY, MOS_PROXY, CUST_PROXY.	1
protocol	String	Protocol to be used for the proxy setting.	
hostname	String	Host name of the proxy server	
port	String	Port details of the proxy server	
runPrivilegeOption	String	Privilege option for creation of SSH Key. Possible values are NORMAL and SUDO	
credentialName	String	Name for the SSH key	
sshPrivateKey	String	Private key	
sshPublicKey	String	Public key	
runAsOption	String	Run as option, if privilege type is selected	
serviceType	String	Service type	
coudSettings	List <Cloud-SettingEntity>	Entity for the service settings	

Supported Request Methods

This table lists all the operations that are supported on the Oracle Cloud setup.

GET	POST	PUT	DELETE
No	Yes	No	No

Setting Up Oracle Cloud

This section shows an example of interaction with Enterprise Manager Cloud Control with Cloud Application plug-in 12.1.0.7 and higher. A POST request is issued on the Oracle Public Cloud setup.

Feature	Description
URI	/em/websvcs/restful/extws/cloudservices/admin/cfw/v2/opcsetup
Request Headers	Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI= Content-Type: application/json

Feature	Description
Body	<pre>{ "domainName": "mydomainname", "dataCenterId": "2", "serviceType": "OPC_DBCS_SERVICE_TYPE", "userName": "CloudUser2@oracle.com", "password": "password", "accountName": "CloudUser2", "proxyType": "CUST_PROXY", "protocol": "https", "hostname": "proxy.mycompany.com", "port": "8080", "credentialType": "ROOT", "credentialName": "cred-key1", "sshPrivateKey": "-----BEGIN RSA PRIVATE KEY-----\nMIIEowIBAAKCA- QEA2aVNjtMW+AQWdasdadsadsaj1Y00okoidX1Y8dxjtpAMed6SswRH1DC02Yu \nR168UvRx54v/7Y4d1LEj+33tiR9pJ0DGvs5ddWUrhBb5zbje4P1XP1twwzePmSzE\nR0- KUUhukAS68RKTFG9IK18eta0gwzGKmWgHrjKAU2tiI0W5t6baVpPg7SAHp9Yf2\nwlf +DvvoGuLLqBdFuWqoeGJa8WONkSes5UaHe1CIKCLBH4pfG0DEiNnALkv34C/V\nm +NXbudy1BAH9/iOT9hE0zShyUfe8IOc06/mSwIBIwKCA- QEA021hS0pQ1Sdo3kk8\n7b2m0Ov0NbHD7yn4o/XPInBD00bTa7IaJedTHTjxu- Zi56LdBtwHgpF/9QLT/ko5B\n8ReGd0TfwA0S4orfRNe4o653ChpWSFSh/ DAkqFss0k18RdqDY0aX+57mv7gdAY9h\nn+y7ecgqCC0mE87rnI7RGIEozBsZ2/ gPRwNXu0eswMDQ7ZVHLH4R57N0xIFFd68W\nn0M8WMh0ris0VfB/X0qYsYc- MY9Brfx1n6LgighpWYr48L2Wehuy+kyfJ/HwSV3CrE\n6e +gt7vh4JUJRp2B2ZEznsS5cd0cAnC8imf34FTgsT/9pzlhXR0tS0kpAx1PazLd \nwbXJ0wKBgQD6BEfbdlm5jmXk6FzWxgqXgQVRfEB5kK6GcfowaCv6rIFsnunnNNOA \n00fZx0K41i06IbfhE/uED3vu95Q05Eu1D5iWT8GLVg1eqMuTE2zH2oV20ZBThnd \nAMMDvkHEbGskAoH1q7pj9pcYJxoeDSy4B1TpMDDS+S7rjJBAmPDMTKBgQDe2rPN \nHp2SGMsrmdsHdTVrw+lqhayN6i5bnnfEbV4JQpolnXeUYK06sQMubn4lo2w7PdK \naP8ug0H9FVqjVpto4ld21sfrf+D6DjpZigmu6g5zp15JF85eURZ8mBJEY0QsJnY0 \nYxhHIFhZx+i/KBDE4SG0TJLZT0QwsdwVew5o9wKBgQCKS/S0yh20c4TJoAJ+giRI\nwnw/ WdjJ/bbSD9+2NhkjJTMoULEoqKyD97HELcsDf3S1rr8H6TvW9Lr8oAcfK/Gx\nnDyGuV +eWnKeuHWP0XTjbGpIksMy0qNIFnXlmkh5w5bdJjgLGQRVze5DEmGeuDNU \nXJbiYX8rjc5gRm1sR0NspwKBgQDLwKRjxDhLD1qrhzn8cwvT00GBJnsmmUn0G90j \nuGtmiEQH2gVdjFo/AodbuG+EMeChwY1Z0+1f1JGeex4FKtYLG05c+nkHSV26G/c \n1fo5hYmOT+Es4pDLQtK7Dre0+6wLG9mmwQBBB7gu+JpIXyVGSi105uypat9CdrSl \n612hywKBgEd1/aGglDCzUTC68lJrL89bG6J0ZfUdMKZaaD8P3hcZCE+t4R1k4NYq\nnyb- JexZZtbZSdvehk0D9q4ar6DITA68n/00urJZadpfxUZoy91AousarVmEtQ2aw \nXLq7C5rdpnIqKXglTKycg4Ej6PP7z00zZs40Je1wlvxJ3c2PTTeD\n-----END RSA PRIVATE KEY-----", "sshPublicKey": "AAAAB3NzaC1yc2EAAAABIwAAQEA2aVNjtMW+asdqwewq- dA@#@DSASASDSDSAMEd6SswRH1DC02YuR168UvRx54v/7Y4d1LEj +33tiR9pJ0DGvs5ddWUrhBb5zbje4P1XP1twwzePmSzEROKUUhU- kAS68RKTFG9IK18eta0gwzGKmWgHrjKAU2tiI0W5t6baVpPg7SAHp9Yf2wlf+Dvvo- GuLLqBdFuWqoeGJa8WONkSes5UaHe1CIKCLBH4pfG0DEiNnALkv34C/VM+NXbudy1BAH9/ iOT9hE0zShyUfe8IOc06/mSw== bpaLiwal@slc03ppq" }</pre>
Request Method	POST

A sample of the response received is shown below

```
{
  "statusMessages": [
    "Successfully registered Identity Domain with name mydomainname",
    "Successfully registered Cloud Account",
    "Successfully added hybrid Cloud Gateways",
  ]
}
```

```

    "Successfully added SSH key",
    "Successfully updated proxy setting",
    "Successfully saved service setting",
    "Successfully submitted discovery request with identifier
46D3C4EA4CAF6FC2E053057FB10A6EA1"
  ]
}

```

Setting Up Hybrid Cloud Using EMCLI

This section covers the emcli commands that can be used to setup hybrid cloud. The input parameters are specified in a property file and specified along with the emcli hybrid_cloud_setup verb on the command line. An example is shown for each operation along with the required input parameters and output.

1. Create Identity Domain

Registers an identity domain. If the DATACENTER_NAME is not specified, a data center is selected randomly from the existing data centers.

Usage:

```
emcli hybrid_cloud_setup -register_identity -input_file =data:/scratch/
tvmcfw_hybrid_opc_setup.prop
```

Input Text:

```
DOMAIN_NAME=gededicatedpoc
DATACENTER_NAME=US Commercial 1 (us1)
```

Response:

Successfully created identity domain.

2. Create Cloud Account

Registers a cloud account. If the property SAVE_ACCOUNT_WITHOUT_TEST is marked, the account is created without testing. If this property is not specified, the URL_TO_TEST must be specified to test the account before saving.

Usage:

```
emcli hybrid_cloud_setup -register_account -input_file=data:/scratch/
tvmcfw_hybrid_opc_setup.prop
```

Input Text:

```
DOMAIN_NAME=gededicatedpoc
USER_NAME=dummy@oracle.com
PASSWORD=Dbas2016
SERVICE_TYPE=OPC_DBCS_SERVICE_TYPE
ACCOUNT_NAME=dummy@oracle.com
URL_TO_TEST=https://dbcs.emea.oraclecloud.com:443/paas/service/dbcs/api/v1.1/instances/
SAVE_ACCOUNT_WITHOUT_TEST=false
```

Response:

Account name registered successfully.

3. Create Gateways

Registers an available agent as gateway. The GUID of the available agents can be retrieved using the [Listing All Available Agents](#) option.

Usage:

```
emcli hybrid_cloud_setup -add_hybrid_gateway -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Input Text:

```
GATEWAY_AGENT_GUID= 8B32C6AC9ASASSA4BF24B4266A6A38B01B
```

Response:

Hybrid gateway registered.

4. Setup Proxy

Updates proxy settings.

Usage:

```
emcli hybrid_cloud_setup -register_proxy_settings -input_file=data:/  
scratch/tvmcfw_hybrid_opc_setup.prop
```

Case 1 Input Text:

```
PROXY_HOSTNAME= proxy.mycompany.com  
PROXY_PORT=8080  
PROXY_PROTOCOL=https  
PROXY_TYPE=CUST_PROXY
```

Response:

Proxy settings updated.

Case 2 Input Text:

```
PROXY_TYPE=NO_PROXY
```

5. Create SSH Keys

Creates an SSH key. The required fields like private key, public key, user name must be specified in the property file.

Usage:

```
emcli hybrid_cloud_setup -add_ssh_keys -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Case 1 Input Text:

```
SSHKEY_USER_NAME=myuser  
CREDENTIAL_NAME=DUMMY_NORMAL  
PRIVILEGE_OPTION=NORMAL  
RUN_AS_OPTION=ROOT  
SSH_PRIVATE_KEY=<Private Key>  
SSH_PUBLIC_KEY=<Public Key>
```

Response:

SSH Keys updated.

Case 2 Input Text

```
SSHKEY_USER_NAME=myuser  
CREDENTIAL_NAME=DUMMY_ROOT  
PRIVILEGE_OPTION=SUDO  
RUN_AS_OPTION=ROOT  
SSH_PRIVATE_KEY=<Private Key>  
SSH_PUBLIC_KEY=<Public Key>
```

Response

SSH Keys updated.

6. Update Service Settings

Update global and account specific service settings. If the account name is not specified, the global settings are used.

Usage:

```
emcli hybrid_cloud_setup -register_service_settings -input_file=data:/  
scratch/tvmcfw_hybrid_opc_setup.prop
```

Case 1 Input Text:

```
SERVICE_TYPE=OPC_DBCS_SERVICE_TYPE  
ACCOUNT_NAME=dummy@oracle.com  
SVC_SETTINGS_DEFAULTSERVICEURL=https://dbaas.oraclecloud.com/  
SVC_SETTINGS_DEFAULTSSHNAMEDCREDENTIAL=43C5CA98SAASSAE053057FB10AB1A5  
SVC_SETTINGS_DEFAULTSSHROOTCREDENTIAL=43C5CA98126DSASFE053057FB10AB1A5  
SVC_SETTINGS_DEFAULTAGENTPORT=3872  
SVC_SETTINGS_DEFAULTAGENTBASEDDIRECTORY=/var/opt/oracle/emagents
```

Response:

Service settings updated successfully.

Case 2 Input Text:

```
SERVICE_TYPE=OPC_DBCS_SERVICE_TYPE  
SVC_SETTINGS_DEFAULTSSHNAMEDCREDENTIAL=43C5CA98126D013FE053057FB10AB1A5  
SVC_SETTINGS_DEFAULTSSHROOTCREDENTIAL=43C5CA98126DSASFE053057FB10AB1A5  
SVC_SETTINGS_DEFAULTAGENTPORT=3872  
SVC_SETTINGS_DEFAULTAGENTBASEDDIRECTORY=/var/opt/oracle/emagents
```

Response:

Service settings updated successfully.

7. Discover Service

This command is used to submit a discovery request. Multiple service names can be specified with comma separated values.

Usage:

```
emcli hybrid_cloud_setup -discover -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Input Text:

```
SERVICE_TYPE=OPC_DBCS_SERVICE_TYPE  
DISCOVER_TARGET_NAME=service1,service2
```

Response:

Discovery request submitted. Track progress using the <GUID_ID>

8. Listing All Service Setting Names**Usage:**

```
emcli hybrid_cloud_setup -list_service_settings -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Response: Listing service settings.

```
defaultAgentBaseDirectory  
defaultAgentImageName  
defaultAgentImageVersion  
defaultAgentGatewayProxyPort  
defaultAgentPort  
defaultSSHNamedCredential  
defaultSSHRootCredential  
defaultPrivDelSetting  
defaultPreScript  
defaultPostScript  
defaultPreScriptRunAsRoot  
defaultPreScriptRunAsOMS  
defaultPostScriptRunAsRoot  
defaultPostScriptRunAsOMS  
defaultAdditionalParams  
agentSoftwareStageLocation  
overrideAccessSettings
```

9. Listing All Available Agents**Usage:**

```
emcli hybrid_cloud_setup -list_available_agents -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Response:

```
Listing available agents  
xyz.example.com          8B32C6AC9ED924BF24B4266A6A38B01B
```

10. Listing All Data Centers**Usage:**

```
emcli hybrid_cloud_setup -list_datacenters -input_file=data:/scratch/  
tvmcfw_hybrid_opc_setup.prop
```

Response

```
Listing data centers  
1      US Commercial 1 (us1)  
2      US Commercial 2 (us2)  
3      US Commercial 6 (us6)  
4      US Commercial 8 (us8)  
5      APAC Commercial 1 (ap1)
```

6	APAC Commercial 2 (ap2)
7	CA Commercial 2 (ca2)

Defining the Oracle Database Cloud Service Settings

You can define the global and account specific settings for accessing your cloud database service here.

Agent Deployment

Specify the settings for installing the Hybrid Cloud Agent here:

- **Deploy using a gold image:** An Agent Gold Image represents the ideal state of a Management Agent in a data center managed by Enterprise Manager, having a customized configuration of the desired versions of the Management Agent software, the desired versions of the monitoring plug-ins, and the desired patches. Select this check box if you want to optionally use a gold image to deploy the hybrid cloud agent.

You must select an image name and a version from the drop down lists.

- **Installation Base Directory:** Enter the absolute path to the base directory on the Oracle Cloud virtual host where you want the software binaries, security files, and inventory files of the Hybrid Cloud Agent to be copied. If the path you enter does not exist, the application creates a directory at the specified path, and copies the Hybrid Cloud Agent software binaries, security files, and inventory files there.
- **Named Credential:** For Named Credential, select the named credential that you want to use to set up SSH connectivity between the on-premise OMS and the destination Oracle Cloud virtual hosts, and to install a Hybrid Cloud Agent on each of the Oracle Cloud virtual hosts. Ensure that you only specify a named credential that uses SSH public key authentication.
- **Root Credential:** The root credentials to be specified while discovering the database cloud service.
- **Privilege Delegation Setting:** Use the default value provided here. Privilege delegation providers and locked accounts are not supported for Hybrid Cloud Agent deployment. If the Hybrid Cloud Agent install user has root privileges, then root.sh is run automatically on the destination virtual hosts post deployment.
- **Port:** Accept the default port (3872) that is assigned for the Hybrid Cloud Agent to communicate, or enter a port of your choice. The custom port you enter must not be busy. If you are not sure, you can leave this field blank. Cloud Control automatically assigns the first available free port within the range of 1830 - 1849.
- **Hybrid Gateway Proxy Port:** This is a mandatory field. Enter a port on which the Hybrid Cloud Agent can communicate with the Hybrid Cloud Gateway Proxy. If you do not specify a value, port 1748 is used, and if port 1748 is not free, then a free port between 1830 and 1848 is used.

Advanced Options

- **Pre / Post Installation Script:** If you want to run certain scripts before or after deploying the Hybrid Cloud Agents, in the Optional Details section, enter the absolute path to the locations where the scripts that you want to run are available. Note that only shell scripts are supported, and only one pre installation or one post installation script can be specified.

If you want to run the script as root, then select **Run as Root**. If the script is on the host where the on-premise OMS is running and is not on the virtual host where you want to install the Hybrid Cloud Agent, then select **Script on OMS**. In this case, the script will be copied from the on-premise OMS host to the destination virtual hosts, and then run on the destination virtual hosts.

- **Additional Parameters:** If you want to specify certain additional parameters for the deployment, enter a white space-separated list of the additional parameters.

After entering all the details, click **Apply** to apply the Hybrid Cloud Agent settings.

Discovery

To configure the discovery settings for the database cloud service, specify the following details:

- **Service Instance Owner:** Select the Enterprise Manager user who will be able to manage monitor the services from the Self-Service Portal. If you do not select the user here, by default, the service instance owner is used to discover the database cloud service.
- **Service URL:** The REST endpoint URL for the Cloud Database Service. To obtain this URL, open your Web browser and enter `http://cloud.oracle.com`. On the Cloud Dashboard, click the **Overview** tab on the left column, and copy the REST Endpoint. Add `/paas/service/dbcs/api/v1.1/instances/` and then add the Identity Domain. For example, `https://dbaas.oraclecloud.com/paas/service/dbcs/api/v1.1/instances/abcd/` where `abcd` is the Identity Domain.
- **Default DBSNMP Monitoring Password:** Specify the default monitoring password for the cloud database service.

Click **Apply** to update the discovery settings.

Access

You can configure privileges for user roles who can manage the service instances. Click **Add**. In the Grant Privileges window, select one or more user roles and the privileges that are to be granted to these roles. Privileges can be:

- **View:** Users with this privilege, can view the service instance.
- **Manage:** Users with this privilege can view and manage the service instance.
- **Full:** Users with this privilege can view, manage, and delete the service instance

Click **Apply** to apply the access settings.

Target Properties

Specify the global target properties such as Contact, Cost Center, Department, and so on and click **Apply**.

Auto Discovery Schedule

Specify the discovery schedule here. Select the **Discover and Manage New Services Automatically** checkbox and specify the time, recurring schedule, and the frequency for auto discovery. The discovery job automatically discover new database service instances available on Oracle Cloud.

If any database cloud service instances have been decommissioned, you can select the **Cleanup obsolete instances** checkbox to delete any obsolete service instances during discovery.

Accounts

After defining the global settings, click the **Accounts** tab to define the Agent Deployment, Discovery, Access, Target Properties, and Auto Discovery Schedule settings for each user account. The Accounts table shows the list of registered user accounts and the configuration settings. Place your cursor on a configuration setting to view whether global or custom settings have been applied.

The global settings are displayed by default and you can select **Override Global Settings** to define account specific settings.

If you do not want to follow the auto discovery schedule but you want to discover the database cloud service now, click **Discover Now** and follow the steps in the Discover Cloud Services wizard:

1. In the Discover Service page, select a discovered database cloud service and click Next.
2. In the Hybrid Cloud Agent(s) page, the data in all the fields is pre-populated from the service settings page. Verify the details like Installation base directory, Port, Hybrid cloud gateway proxy port and so on and click Next.
3. In the Target Configuration page, the Default DBSNMP Monitoring Password and the global target properties are pre-populated. Click Next.
4. In the Review page, verify all the details and click Submit. You will see a confirmation message that deployment procedure for discover service operation is setup successfully. In the message, click **View Details** to view the Procedure Activity page to see the status.

Once the database service running on Oracle Cloud has been discovered, you can view the discovered database instance and monitor it in Enterprise Manager. To view and monitor your database cloud service, login into the Enterprise Manager as a user with EM_SSA_USER privileges and select Oracle Public Cloud Services and click the database cloud service instance. See [Monitoring the Discovered Database Cloud Service](#) for more details.

You can also click **Clean Up** to clean up and delete any obsolete instances.

Discovering the Database Cloud Service

You must discover the database cloud service instance so that it can be managed and monitored by a Self Service User in Enterprise Manager.

Before you discover the database cloud service instance, you must have done the following:

- Registered identity domains and created cloud accounts. See [Registering the Identity Domain and the Cloud Account](#).
- Defined the database cloud service settings. See [Defining the Oracle Database Cloud Service Settings](#).

To discover the database cloud service instance, follow these steps:

1. From the **Setup** menu, select **Hybrid Cloud**.

The Hybrid Cloud Setup: Console appears.

2. Click on the **Oracle Database Cloud Service** option.
The Service Settings: Oracle Database Cloud Service page appears.
3. Click on the **Accounts** tab. The list of accounts is displayed.
For each account, hover over the icon to view if global settings or account specific settings have been defined for agent deployment, access, and the target properties.
4. Select an account from the list and click **Discover Now** to launch the Discover Oracle Cloud Services wizard.
5. In the Services page, select the Cleanup Obsolete Instances checkbox to clean up and delete obsolete instances.
6. Select the database cloud service instance(s) that are to be discovered and click **Next**.
7. In the Hybrid Cloud Agents page, you can select:
 - Select the **Use the same hybrid cloud agents for all services** checkbox to associate the same hybrid cloud gateway agent with all services.] Specify the port number on which the hybrid cloud gateway agent will communicate hybrid gateway proxy.
 - Select the Use same private key credential for all virtual machines option to use the same key credentials.
8. Enter the Installation Base Directory path and in the Port field, accept the default port number or specify a custom port number. A custom port number between 1830 and 1849 is displayed. Specify the Privilege Delegation Setting and click **Next**.
9. Enter the target configuration details such as the default SNMP password, cost center, lifecycle status and so on and click **Next**.
10. Review the details and click **Submit**.

Managing Data Sources

This section covers the procedures to create and add data sources that are based on an on-premise database or a database cloud service.

Topics

- [Creating a Data Source](#)
- [Adding a Data Source](#)

Creating a Data Source

You can create a data source from an on-premises production database or a database cloud service. This data source is then used to refresh your database cloud service and create clones.

Before you create a data source, you must have done the following:

- Registered the identity domains and cloud accounts. See [Registering the Identity Domain and the Cloud Account](#).
- Configured the hybrid cloud gateway. See [Configuring the Hybrid Cloud Gateway](#).

- Configured the proxy settings (Optional). See [Configuring the Proxy Settings](#).
- Added the SSH Key Credentials. See [Configuring the SSH Keys](#).
- Discovered the database cloud service in Enterprise Manager. See [Defining the Oracle Database Cloud Service Settings](#).
- Defined the Global Oracle Database Backup settings. See

To create a data source, follow these steps:

1. Click **Create** in the Data Sources page.

The **Data Source** dialog box appears.

2. You can select one of the following:

- Using a Database: This option creates a test master database on-premise or on Compute Cloud that is automatically refreshed at a specified frequency. The database backups are uploaded to the Oracle Storage Server. This data source is then used to create clones in Oracle Cloud.
- Using a Database Cloud Service: This option creates a test master database on Oracle Cloud that is automatically refreshed from the source at a specified frequency. The backups are uploaded to Oracle Storage Server. This data source is then used to create full clones in Oracle cloud.

3. Depending on your choice, you can either select either the on-premise database or the cloud database service and click **Select**.

The **Create Data Source** wizard appears.

4. In the Source and Destination page, enter the following details:

- The name for the data source.
- The source database and host credentials.
- The Oracle Home Location on the destination host. The destination host on which the test master database is being created (on-premise or on Oracle Cloud) is displayed.
- The destination database definition including the global database name and the SID is displayed. Enter the database host credentials on the destination host.

5. Click **Next**. In the Configuration page, specify the following:

- Database Files Location
 - File System: The Oracle Database File System creates a standard file system interface on top of files and directories that are stored in database tables. If you select this option, you must specify or select the Location of the File System. You can specify a common location for all the files or you can select the Use Oracle Optimal Flexible Architecture-compliant directory structure (OFA) checkbox and specify different locations for data files, redo log files, and so on.
 - Automatic Storage Management: The Oracle Automatic Storage Management (ASM) is a volume manager and a file system for database files that supports single-instance and RAC configurations. ASM groups the disks in your storage system into one or more disk groups. If you select ASM, select a common location for the database files.
- Listener Configuration: Click **Add** to add one or more listener targets that are to be associated with the test master database.

- Database Credentials: Specify the passwords for the administrative users (SYS, SYSTEM and DBSNMP) of the new database being cloned. You can choose to use the same password for all the schemas or different passwords for each schema.
6. Click **Next**. In the Initialization Parameters page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click the Edit icon to modify the value of the parameter. Some values such as db_block_size cannot be modified.
 7. Click **Next**. In the Post Processing page, specify the following:
 - Masking Definition: Specify the masking definition that is to be applied after the database is cloned.
 - Custom Script: Specify the custom scripts that need to be executed before and after the database is created.
 - Database Mode: Select the database mode after it is created or refreshed. This can be Read-Only, Read-Write, or Shutdown.
 8. Click **Next**. The Data Management page appears:

- OSS Container Details: Specify the Oracle Storage Service (OSS) details on which this database is to be backed up. This backup is used to clone the database in the Oracle Cloud.
 - Select the Cloud Account from the drop down list. This account will be used to refresh and clone the database in Oracle Cloud. The Identity Domain, User Name, and Service Name of the Cloud Account is displayed. See [Registering the Identity Domain and the Cloud Account](#).
 - All backups are stored in the default storage container. If you have custom containers in which the backup is to be stored, specify the prefix here.
 - You can choose to take the OSS backup as soon as the database is created or take the backup later.
- Database Refresh Policy: The cloned database must be refreshed at regular intervals to ensure that the latest data is available. Whenever the cloned database is refreshed, the database must be backup up on OSS so that the new data is available to the cloned databases on Oracle Cloud. You may choose to take the backups automatically or perform the same operation as a separate activity from the Data Source Management page. Select the **Automatically refresh the database** checkbox and specify the refresh frequency.

- **Backup Management Policy:** You may also limit the number of backups that are available for cloning at any time. This can be determined by a policy. The policy is defined by the total number of backups as well as the maximum lifetime of these backups. This will ensure that older backups are periodically deleted, and sufficient storage is available for new backups.
 - **Encryption Details:** The backups available on cloud must be encrypted. Specify the encryption algorithm and the password to encrypt the OSS backup.
9. Click **Next**. Enter a name for the deployment procedure and specify the schedule.
 10. Click **Next**. Review the details entered so far and click **Create**.

The Procedure Activity page appears where you can monitor the progress of the data source creation. This procedure creates a clone of the on-premise database or cloud database service, then creates a backup profile of the clone created, and saves the backup in the Oracle Storage Service. When the procedure has been completed, you can view the newly created data source on the Data Sources page.

Select the Data Source that has been created to view the database details, backup policy, and so on. Click the **Backups** tab and click **Schedule Now** to create a new backup of the selected data source on demand. You can navigate to the Procedure Activity to monitor the progress. Once the procedure is completed, the newly created backup will appear under the Backups tab.

Adding a Data Source

You can add an existing on-premise or a database on Oracle Cloud that can be used to upload backups to the Oracle Storage Service at a specified frequency.

To add a data source, follow these steps:

1. Click **Add** in the Data Sources page.

The **Add Data Source** dialog box appears.

2. You can select:

- **Using an existing database:** Identifies an existing on-premise database or on Compute Cloud that can be used to upload the backups to Oracle Storage Service at a specified frequency. This data source can then be used to create clones in Oracle Cloud.
- **Using an existing database cloud service:** Identifies an existing database in Oracle Cloud that can be used to upload the backups to Oracle Storage Service at a specified frequency. This data source can then be used to create clones in Oracle Cloud.

The **Add Data Source** wizard appears.

3. In the Select Database page, enter the **Data Source Name**, **Description**, and depending on your choice in Step 2, select a database or a database cloud service from the list.
4. Enter the database and host credentials and click **Next**.
5. In the Data Management page, specify the OSS Container Details, Backup Schedule, Backup Management Policy, and Encryption Details. Click **Add**.

The Procedure Activity page appears where you can monitor the progress of the data source creation. This procedure creates a database profile and stores it in the Oracle Storage Service. This profile is used to create clones in Oracle Cloud. When the pro-

cedure has been completed, you can view the newly added data source on the Data Sources page.

Creating Backups On the Oracle Storage Server

While creating a cloned database to be used for Oracle Cloud, you can choose to refresh the database automatically and create backups. You can also choose to create backups on demand from the Data Sources page.

To create on demand backups, follow these steps:

1. From the **Hybrid Cloud Setup** menu, navigate to the Data Sources page.

The list of available data sources appears.

2. Select the data source to be backed up.

The data source details are displayed under the General tab.

3. Click the **Backups** tab and then click **Schedule Now**.

This launches the procedure to create a backup for the selected data source. You can view the status of the Procedure Activity page.

You will see the DB instance details such as Version, OCPUs, Memory, Storage, and so on. This backup will be stored on the OSS and can be used to create clones.

Defining the Database Backup Cloud Service Settings

The Oracle Database Backup Cloud Service is a secure, scalable, on-demand storage solution for backing up Oracle databases to Oracle Cloud. The service complements your existing backup strategy by providing an off-site storage location in the public cloud.

You can configure the global and account specific settings for the Oracle Database Backup Cloud Service. Enter the following details:

- **REST End Point:** Specify the REST End Point URL for the database backup service.
- **Container Name:** Specify the storage container for the backups.
- **Service Name:** Specify the database backup service name.

Monitoring the Discovered Database Cloud Service

After you have discovered the database cloud service in Enterprise Manager, you can view and monitor the database service instances from the Self Service Portal.

1. Login to Enterprise Manager as a user with the EM_SSA_USER role.

The discovered database services are displayed on the Oracle Public Cloud Services Home page.

2. Click on a database instance name to view the Database Cloud Service Instance Home page.

You will see the DB instance details like Version, OCPUs, Memory, Storage etc. The data that you see in this page is exactly the same that you will see on the Oracle Database Cloud Service page. In this page, next to Additional Information click on **Show More** to view the data in detail.

On the left hand side, you will see the following tabs:

- Summary: In this page, details such as Identity Domain, Subscription Type, Creation Date, and so on are displayed. Click **Show More** to view more detailed information.
- Clone Management: See [Clone Management](#)
- Data Management: See [Refreshing the Cloud Database Service](#)
- Load and Capacity: This page load the performance data of the database instance. Expand on Resources section to view the Host CPU, Active Sessions, Memory usage, and Data Storage details
- Patching
- High Availability

Refreshing the Cloud Database Service

To ensure that the latest data is available on your Oracle Cloud database, you must refresh your cloud database service at periodic intervals.

To refresh the cloud database service, follow these steps:

1. Login to Enterprise Manager as a user with EM_SSA_USER privileges.

The Oracle Public Cloud Services Home page appears and you will see a list of discovered database instances.

2. Click on the Database Instance Name to view the Database Home page.
3. Click the Data Management tab on the Home page. You will see a message indicating that no data source is associated with the database cloud service.
4. Click Choose Data Source. The Choose Dialog box appears. Choose a data source from the list and click Select.

The data source available for this database cloud service is displayed.

5. Select a backup and click **Associate**.

The Associate Data Source confirmation dialog box appears. Select the **Yes, perform this task** checkbox and click **Associate**. You will be returned to the Database Home page where you can see that the database instance has been associated with the selected backup and has been refreshed as part of this process. When a new backup is available, select the backup and click **Refresh** to refresh the contents of the database.

Clone Management

You can create full clones and snap clones of the database cloud services from the Self Service Portal.

Topics

The following topics are covered:

- [Using Snapshots to Create Snap Clones](#)
- [Using Database Backups to Create Full Clones](#)

Using Snapshots to Create Snap Clones

You can create snapshots of the database and use these snapshots to create snap clones.

To create a snap clone, follow these steps:

1. Login to Enterprise Manager as a user with the EM_SSA_USER privileges.
The Oracle Public Cloud Services Home page appears and you will see a list of discovered database instances.
2. Click on the Database Instance Name to view the Database Home page.
3. Click the Clone Management tab and click the Snapshots, then click Create.
The Create Snapshot dialog box appears.
4. Enter a name and description for the snapshot and click **Create** to create the snapshot.
The newly created snapshot appears under the Snapshots tab.
5. In the Snap Clone Databases header, click **Create**.
The Create Snap Clone dialog box appears.
6. Select a snapshot on which the clone is to be based. The Destination details are populated but can be modified if required. Click Clone to create a snap clone of the database.

Using Database Backups to Create Full Clones

You can use database backups to create full clones of the cloud database service

To create a full clone, follow these steps:

1. Login to Enterprise Manager as a user with the EM_SSA_USER privileges.
The Oracle Public Cloud Services Home page appears and you will see a list of discovered database instances.
2. Click on the Database Instance Name to view the Database Home page.
3. Click the Clone Management tab and then click **Create** in the Full Clone Databases header.
4. The Clone dialog box appears.
5. Select a backup on which the clone is to be based. The Destination details are populated but can be modified if required. Click Clone to create a full clone of the database.

Defining the Oracle Compute Service Settings

You can define the global and account specific settings for accessing your compute service here.

Agent Deployment

Specify the settings for installing the Hybrid Cloud Agent here:

- **Deploy using a gold image:** An Agent Gold Image represents the ideal state of a Management Agent in a data center managed by Enterprise Manager, having a customized configuration of the desired versions of the Management Agent software, the desired versions of the monitoring plug-ins, and the desired patches. Select this check box if you want to optionally use a gold image to deploy the hybrid cloud agent.

You must select an image name and a version from the drop down lists.

- **Installation Base Directory:** Enter the absolute path to the base directory on the Oracle Cloud virtual host where you want the software binaries, security files, and inventory files of the Hybrid Cloud Agent to be copied. If the path you enter does not exist, the application creates a directory at the specified path, and copies the Hybrid Cloud Agent software binaries, security files, and inventory files there.
- **Named Credential:** For Named Credential, select the named credential that you want to use to set up SSH connectivity between the on-premise OMS and the destination Oracle Cloud virtual hosts, and to install a Hybrid Cloud Agent on each of the Oracle Cloud virtual hosts. Ensure that you only specify a named credential that uses SSH public key authentication.
- **Root Credential:** The root credentials to be specified while discovering the compute cloud service.
- **Privilege Delegation Setting:** Use the default value provided here. Privilege delegation providers and locked accounts are not supported for Hybrid Cloud Agent deployment. If the Hybrid Cloud Agent install user has root privileges, then root.sh is run automatically on the destination virtual hosts post deployment.
- **Port:** Accept the default port (3872) that is assigned for the Hybrid Cloud Agent to communicate, or enter a port of your choice. The custom port you enter must not be busy. If you are not sure, you can leave this field blank. Cloud Control automatically assigns the first available free port within the range of 1830 - 1849.
- **Hybrid Gateway Proxy Port:** This is a mandatory field. Enter a port on which the Hybrid Cloud Agent can communicate with the Hybrid Cloud Gateway Proxy. If you do not specify a value, port 1748 is used, and if port 1748 is not free, then a free port between 1830 and 1848 is used.

Advanced Options

- **Pre / Post Installation Script:** If you want to run certain scripts before or after deploying the Hybrid Cloud Agents, in the Optional Details section, enter the absolute path to the locations where the scripts that you want to run are available. Note that only shell scripts are supported, and only one pre installation or one post installation script can be specified.

If you want to run the script as root, then select **Run as Root**. If the script is on the host where the on-premise OMS is running and is not on the virtual host where you want to install the Hybrid Cloud Agent, then select **Script on OMS**. In this case, the script will be copied from the on-premise OMS host to the destination virtual hosts, and then run on the destination virtual hosts.

- **Additional Parameters:** If you want to specify certain additional parameters for the deployment, enter a white space-separated list of the additional parameters.

Discovery

To configure the discovery settings for the compute cloud service, specify the following details:

- **Service Instance Owner:** Select the Enterprise Manager user who will be able to manage monitor the services from the Self-Service Portal. If you do not select the user here, by default, the service instance owner is used to discover the compute cloud service.
- **Service URL:** The REST endpoint URL for the Cloud Database Service. To obtain this URL, open your Web browser and enter `http://cloud.oracle.com`. On the Cloud Dashboard, click on the **Overview** tab on the left column, and copy the REST Endpoint. Add `/paas/service/dbcs/api/v1.1/instances/` and then add the Identity Domain. For example, `https://dbaas.oraclecloud.com/paas/service/dbcs/api/v1.1/instances/abcd/` where `abcd` is the Identity Domain.

Click **Apply** to update the discovery settings.

Access

You can configure privileges for user roles who can manage the service instances. Click **Add**. In the Grant Privileges window, select one or more user roles and the privileges that are to be granted to these roles. Privileges can be:

- **View:** Users with this privilege, can view the service instance.
- **Manage:** Users with this privilege can view and manage the service instance.
- **Full:** Users with this privilege can view, manage, and delete the service instance

Click **Apply** to apply the access settings.

Target Properties

Specify the global target properties such as Contact, Cost Center, Department, and so on and click **Apply**.

Auto Discovery Schedule

Specify the discovery schedule here. Select the **Discover and Manage New Services Automatically** checkbox and specify the time, recurring schedule, and the frequency for auto discovery. The discovery job automatically discover new database service instances available on Oracle Cloud.

If any compute cloud service instances have been decommissioned, you can select the **Cleanup obsolete instances** checkbox to delete any obsolete service instances during discovery.

Accounts

After defining the global settings, click on the **Accounts** tab to define the Agent Deployment, Discovery, Access, Target Properties, and Auto Discovery Schedule settings for each user account. The Accounts table shows the list of registered user accounts and the configuration settings. Place your cursor on a configuration setting to view whether global or custom settings have been applied.

The global settings are displayed by default and you can select **Override Global Settings** to define account specific settings.

If you do not want to follow the auto discovery schedule but you want to discover the compute cloud service now, click **Discover Now** and follow the steps in the Discover Cloud Services wizard:

1. In the Discover Service page, select a discovered compute service and click **Next**.

2. In the Hybrid Cloud Agent(s) page, the data in all the fields is pre-populated. Verify the details like Installation base directory, Port, Hybrid cloud gateway proxy port and so on and click Next.
3. In the Target Configuration page, the global target properties are pre-populated. Click **Next**.
4. In the Review page, verify all the details and click **Submit**. You will see a confirmation message that deployment procedure for discover service operation is setup successfully. In the message, click **View Details** to view the Procedure Activity page to see the status.

The Enterprise Management Agent is installed on the host and the host is registered as a cloud service instance. Once discovery is complete, login to Enterprise Manager as a user with the EM_SSA_USER privileges, select Oracle Public Cloud Services and click on the compute cloud service instance. You can view the host details but no monitoring capabilities are provided.

You can also click **Clean Up** to clean up and delete any obsolete instances.

Discovering the Oracle Compute Service

You must discover the compute cloud service instance so that it can be managed and monitored by a Self Service User in Enterprise Manager

Before you discover the compute cloud service instance, you must have done the following:

- Registered identity domains and created cloud accounts. See [Registering the Identity Domain and the Cloud Account](#).
- Defined the compute cloud service settings. See [Defining the Oracle Compute Service Settings](#).

To discover the compute cloud service instance, follow these steps:

1. From the **Setup** menu, select **Hybrid Cloud**.
The Hybrid Cloud Setup: Console appears.
2. Click on the **Oracle Compute Service** option.
The Service Settings: Compute Service page appears.
3. Click on the **Accounts** tab. The list of accounts is displayed.
For each account, hover over the icon to view if global settings or account specific settings have been defined for agent deployment, access, and the target properties.
4. Select an account from the list and click **Discover Now** to launch the Discover Oracle Cloud Services wizard.
5. In the Services page, select the Cleanup Obsolete Instances checkbox to clean up and delete obsolete instances.
6. Select the compute cloud service instance(s) that are to be discovered and click **Next**.
7. In the Hybrid Cloud Agents page, you can select:
 - Select the **Use the same hybrid cloud agents for all services** checkbox to associate the same hybrid cloud gateway agent with all services.] Specify the

port number on which the hybrid cloud gateway agent will communicate hybrid gateway proxy.

- Select the Use same private key credential for all virtual machines option to use the same key credentials.
8. Enter the Installation Base Directory path and in the Port field, accept the default port number or specify a custom port number. A custom port number between 1830 and 1849 is displayed. Specify the Privilege Delegation Setting and click Next.
 9. Enter the target configuration details such as the cost center, lifecycle status and so on and click **Next**.
 10. Review the details and click **Submit** to submit the discovery job.

Part V

Database Upgrade

This part contains the following chapter:

- [Upgrading Databases](#)

Upgrading Databases

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

- [Getting Started](#)
- [Supported Releases](#)
- [Database Upgradability Report](#)
- [Upgrading Databases Using Deployment Procedure](#)
- [Upgrading an Oracle Database or Oracle RAC Database Instance Using the Database Upgrade Wizard](#)
- [Data Guard Rolling Upgrade](#)

Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in mass upgrade of databases or when you want to install Oracle Home and upgrade database. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully upgrade Oracle database. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 19-1 Getting Started with Upgrading Database

Step	Description	Reference Links
Step 1	Selecting the Usecase Decide whether you want to upgrade single instance database, Oracle RAC database, or Oracle database clusterware, and then select the type of upgrade you want to perform on them. Oracle Database You can upgrade multiple database instances at a time using deployment procedure according the following usecases. : • Upgrading database instance only. • Upgrading cluster database,. • Upgrading clusterware. You can even upgrade one database instance at a time. For example, if you have already upgraded some of the databases in the Oracle home earlier, and you now want to upgrade the other databases in the same Oracle home.	To upgrade databases using deployment procedure, see Upgrading Databases Using Deployment Procedure. To upgrade one single instance database or one Oracle RAC database instance, see Upgrading an Oracle Database or Oracle RAC Database Instance Using the Database Upgrade Wizard.
Step 2	Knowing About the Supported Releases Know what releases of Oracle Database can be upgraded by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Supported Releases.

Table 19-1 (Cont.) Getting Started with Upgrading Database

Step	Description	Reference Links
Step 3	Generating the Database Upgradability Report Generate a Database upgradability report to know the database upgrade readiness.	To generate a report, see Database Upgradability Report .
Step 4	Understanding the Deployment Procedure Understand the Deployment Procedure you need to select, and its scope and coverage.	To learn about the Deployment Procedure offered for upgrading databases, see About Deployment Procedures .
Step 5	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, and setting up of Oracle Software Library.	- To learn about the prerequisites for upgrading using deployment procedure, see Meeting the Prerequisites. - To learn about the deployment phases involved in upgrading a database instance, see Meeting the Prerequisites.
Step 6	Running the Deployment Procedure Run the Deployment Procedure to successfully upgrade Oracle database.	- To learn about the procedure for upgrading database clusters, follow the steps explained in Upgrading Oracle Cluster Database Using Deployment Procedure. - To learn about the procedure for upgrading database clusterware, follow the steps explained in Upgrading Oracle Clusterware Using Deployment Procedure. - To learn about the procedure for upgrading database instance using deployment procedure, follow the steps explained in Upgrading Oracle Database Instance Using Deployment Procedure. - To learn about the procedure to upgrade a database instance using the database upgrade wizard, see Performing the Upgrade Procedure.

Supported Releases

Using this Deployment Procedure, you can mass upgrade the following releases of Oracle Database across multiple hosts:

Table 19-2 Supported Releases for Mass Upgrade of Oracle Databases

Supported Target	Supported Release to Upgrade to	Supported Platform
Oracle Database (single instance database)	11g Release 2 12c Release 2	All Platforms

For upgrading one database instance, the following releases are supported:

Table 19-3 Supported Releases for Upgrading a Database Instance

Supported Target	Supported Release to Upgrade to	Supported Platform
Oracle Database (single instance database)	11g Release 2 12c Release 2	All Platforms
Oracle Real Application Clusters (Oracle RAC)	11g Release 2 12c Release 2	All Platforms
Oracle RAC One	11g Release 2 12c Release 2	All Platforms


Database Upgradability Report

Database upgradability report is an upgrade readiness report that helps you understand if the RAC database is ready for an upgrade. When the report generation is triggered, ORAchk is run to verify if the Grid Infrastructure and RAC database is upgrade ready. It helps you plan the upgrade process for Oracle Clusters and Oracle RAC databases by automating many of the manual pre-checks for database upgrade. Running this check helps you ensure that you have enough time to correct potential issues before the upgrade.


Generating a Database Upgradability Report

To generate a Database upgradability report, follow the steps below:

1. From the **Enterprise** menu select **Provisioning and Patching**, and then select **Database Provisioning**.
2. In the Deployment Procedures section, click **Launch Upgradability Report**.




3. Click **Create New**.
4. In the Create a upgradability report window displayed, enter the **Report Name** and **Report Description**.
5. Select the **Upgrade to Version** and **Credential**.



6. Click **Add Target** and select the required target(s).
7. Click **OK**.

The upgrade checks are triggered and the **Status** column in the Database Upgradability Report page displays the status as 'executing'.

8. After the status changes to 'completed' click on the report name.
9. In the report page, select a row to view the report for each target.



Upgrading Databases Using Deployment Procedure

This section consists of the following:

- [About Deployment Procedures](#)

- [Meeting the Prerequisites](#)
- [Upgrading Oracle Cluster Database Using Deployment Procedure](#)
- [Upgrading Oracle Clusterware Using Deployment Procedure](#)
- [Upgrading Oracle Database Instance Using Deployment Procedure](#)

About Deployment Procedures

Upgrading to Oracle Database 12c enables you to access the latest technology, thereby increasing efficiency and providing secure data management for your applications. Ensure that you follow these steps to plan, prepare, and upgrade to make the upgrade process simpler, faster, and more predictable from start to finish.

Use the Upgrade Database deployment procedure to perform a mass upgrade of databases or when you want to install Oracle Home and upgrade database.

For mass upgrade of databases, source databases must be homogeneous, that is, they must have the same platform and same versions. The gold image to be used for the new database version must contain all patches required for these databases, effectively helping the user to get to a single standard gold image for the new release.

Cloud Control offers the *Upgrade Oracle Database* deployment procedure for mass upgrade of Oracle databases. The deployment procedure supports the following use cases:

- Upgrading Database Instance Only
- Upgrading Database Real Application Clusters
- Upgrade of Clusterware

 Note:

Upgrading Grid infrastructure automatically upgrades the underlying ASM targets.

For upgrading one Oracle database instance at a time or any Oracle RAC database instance, you must access the Oracle Database Upgrade wizard from the Home page of the database that you want to upgrade.

 Note:

The Upgrade Database deployment procedure does not support upgrade of databases with Oracle Database Vault.

The Deployment Procedure can be run by two types of administrators, mainly Designer and Operator. As a designer, you can set up a test database, deploy new software and patches, and test the upgrade process, and then create a gold image out of it. You can then access the Deployment Procedure, provide the required details, and lock one or more of the fields, such as platform, version, move to, and so on. Finally, you can save the procedure and then publish it to the operators.

As an operator, you can access the save procedure and perform only certain operations such as selecting a set of databases based on the locked criteria, providing any additional input that is specific to the runtime activity, and then scheduling the procedure. This way, the operators run a fully tested and certified Deployment Procedure in their production environments, and the entire operation tends to be less error prone.

For more information on these types of administrators, and to learn how you can use these locking feature, see [Creating Enterprise Manager User Accounts](#).

Meeting the Prerequisites

Before you upgrade an Oracle database, follow these prerequisites:

- The following table lists the minimum source version and destination version required for each database target type:

Table 19-4 Meeting the Prerequisites

Database Target Type	Minimum Source Version Supported	Destination Version
Clusterware	10.2.0.4.0 and higher in 10gR2 series 11.1.0.7.0 and higher in 11gR2 series 11.2.0.1 and higher in 11gR2 series 12.1.0.1 and higher in 12cR1 series	11.2 and 12cR2 series
	11.2.0.4 and higher in 11g series 12.1.0.1 and higher in 12cR2 series	18c and 19c
Cluster Database	10.2.0.4.0 and higher in 10gR2 series 11.1.0.7.0 and higher in 11gR2 series 11.2.0.1 and higher in 11gR2 series 12.1.0.1 and higher in 12cR1 series	11.2 and 12cR2 series
	11.2.0.4 and higher in 11g series 12.1.0.1 and higher in 12cR2 series	18c and 19c
Single Instance Database	10.2.0.4.0 and higher in 10gR2 series 11.1.0.7.0 and higher in 11gR2 series 11.2.0.1 and higher in 11gR2 series 12.1.0.1 and higher in 12cR1 series	11.2 and 12cR2 series

Table 19-4 (Cont.) Meeting the Prerequisites

Database Target Type	Minimum Source Version Supported	Destination Version
Single Instance High Availability	11.2.0.4 and higher in 11g series	18c and 19c
	12.1.0.1 and higher in 12cR2 series	
	11.2.0.1 and higher in 11gR2 series	11.2 and 12cR2 series
	12.1.0.1 and higher in 12cR1 series	
Oracle RAC One Node	11.2.0.4 and higher in 11g series	18c and 19c
	12.1.0.1 and higher in 12cR2 series	
	10.2.0.4.0 and higher in 10gR2 series 11.1.0.7.0 and higher in 10gR2 series	11.2 and 12cR2 series
	11.2.0.1 and higher in 11gR2 series	
Automatic Storage Management	12.1.0.1 and higher in 12cR1 series	
	11.2.0.4 and higher in 11g series	18c and 19c
	12.1.0.1 and higher in 12cR2 series	
	10.2.0.4 and higher	11.2 and 12cR2 series
	11.1.0.7 and higher (stand alone type)	
	12.1.0.1 and higher in 12cR1 series	
	11.2.0.4 and higher in 11g series	18c and 19c
	12.1.0.1 and higher in 12cR2 series	

- The database user must have SYSDBA privileges or the OS user must be part of the DBA group.
- The database to be upgraded, and all its node instances (in case of a cluster database) must be up and running.
- Ensure that you have created designer and operator roles with the required privileges. The designer must have EM_PROVISIONING_DESIGNER role and the operator must have EM_PROVISIONING_OPERATOR role.
- The designer must have the following privileges:
 - Edit access to Software Library to manage Software Library entities such as gold images.
 - Add/Edit Target privileges.
 - Create Named Credentials privileges.

- The operator must have the following privileges:
 - View access to Software Library to view Software Library entities such as gold images.
 - Add/View Target privileges.
 - Privileges to the named credentials granted by designer.

Upgrading Oracle Cluster Database Using Deployment Procedure

To upgrade an Oracle Cluster Database using the deployment procedure, follow these steps:

1. In the designer role, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning console, select the **Upgrade Oracle Database** Deployment Procedure and click **Launch**. The Database Upgrade wizard is launched.
3. On the Targets page, in the Select Targets section, select **Cluster Database** as the Target Type.

Note:

The Cluster Database option enables you to upgrade the Cluster Databases and optionally the underlying Automatic Storage Management target and the managing Clusterware as part of the same process.

4. Select the version that you want the database to be upgraded to from the **To** list.
5. Click **Add Targets**.
6. In the Select Targets for Upgrade dialog box, click on the Cluster search icon to search for the cluster database.

In the Search and Select: Cluster Target dialog box, select the cluster database that you want to upgrade, and then click **Select**.

Note:

If the database you want to upgrade does not appear in the table, verify that the database is available and there are no metrics collection errors for the target.

Note:

If the cluster database that you selected is a parent cluster, then all the child nodes such as HAS, ASM, and cluster databases are automatically selected.

- Select the Cluster Database Version. Click **OK**.
7. On the Targets page, click **Save**.
 8. In the Save Procedure dialog box, enter a description.
For example:
Upgrade Oracle Database Procedure 1

Click **Save**.
 9. An Information dialog box appears. Click **OK**.
Click **Log Out**.
 10. The designer role is completed. Now, log in with the Operator credentials.
 11. From the **Enterprise Menu**, select **Provisioning and Patching**, and then **Database Provisioning**.
 12. On the Database Provisioning page, all operations except for Launch are disabled for the Operator. The upgrade database deployment procedure that was saved is selected. Click **Launch**.
 13. On the Targets page, click **Next**.
 14. On the Software Details page, in the Grid Infrastructure section, click the Search icon to select a Grid Infrastructure software from the Software Library to create a new Oracle Home.



Note:

The software may be stored as a gold image or a zipped up installation media in the Software Library.

- From the dialog box that appears, select a Grid Infrastructure software, and then click **Select**.
15. In the Oracle Database section, specify the following details:
 - a. In the Oracle Database Software section, click the search icon to select the Oracle Database Software from the Software Library for creating a new Oracle Home.

The software may be stored as a gold image or a zipped up installation media in the Software Library. Ensure that the zipped up Oracle Home contains all critical patches for the new Oracle Home.

 Note:

To ensure that the gold image you create includes all the recommended patches, follow these steps:

- i. Click the 'Database Upgrade Planner' link and log in to My Oracle Support using Cloud Control in the online mode.
- ii. Select the following types of patches to be applied to the gold image:
 - Recommended patches on the release you want to upgrade to
 - Patches on top of the release that maintain the fixes in the base release
 - Patches resolving merge conflicts if any present between the patches chosen
- iii. Apply the patches to an Oracle Home of the release to upgrade to using Patch Plans or manually.
- iv. Create a gold image from the Oracle Home and use it for the upgrade.

- b. In the Software Location section, specify or check if the software locations of the Oracle base for database and the database Oracle home are correct.
- c. In the User Groups section, specify or check if the specified values for the following user groups are correct:
 - **Database Backup and Recovery (OSBACKUPDBA) group:** Create this group if you want a separate group of operating system users to have a limited set of database backup and recovery related administrative privileges (the SYSBACKUP privilege). The usual name for this group is backupdba.
 - **Data Guard administrative (OSDBDBA) group:** Create this group if you want a separate group of operating system users to have a limited set of privileges to administer and monitor Oracle Data Guard (the SYSDG privilege). The usual name for this group is dgdba.
 - **Encryption Key Management administrative(OSKMDBA) group:** Create this group if you want a separate group of operating system users to have a limited set of privileges for encryption key management such as Oracle Wallet Manager management (the SYSKM privilege). The usual name for this group is kmdba.
- d. In the Working Directory section, specify the working directory on each target host. Ensure that all the hosts have read write permission on the working directory specified.

 Note:

Click the **Lock** icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

- e. Click **Next**.

16. In the Credentials page, specify the Operating System credentials for Grid Infrastructure and Database, Privileged Operating System credentials (run as root), and Database credentials. If you choose to specify Preferred Credentials, select either Normal Host or Privileged Host credentials. For Named Credentials, you can specify the same or different credentials for Oracle homes.

If you have not set Named Credentials, click the plus sign (+) in the Credentials section. In the Add New Database Credentials popup, specify the **User name**, **Password**, **Role**, and specify the Save Details. Select **Run As** and specify **root**. Click **OK**.

Click the **Lock** icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Click **Next**.

17. In the Configuration Details page, the Backup and Restore settings option in the Restore Strategy is selected by default.

In the Restore Strategy section, you can select:

- **Backup and Restore Settings Only** to restore configuration changes made during database upgrade and not actual data, in case upgrade fails.
- **Create RMAN backup Before Upgrade** and enter the backup location.

 Note:

If you perform a backup using RMAN then you can restore the entire database irrespective of datafiles on filesystem or ASM. Backup using RMAN is available when upgrading to 12c only.

- **Use Existing RMAN Backup** where the latest RMAN backup will be used.
- **Use Flashback and Guaranteed Restore Point.**

 Note:

The source database version should be from 11g onwards. The target home version should be 12c

- **Full Backup** to backup the database and restore configuration and oratab settings if upgrade fails. The backup location is, by default, \$ORACLE_BASE/admin/\$GDB/backup where '\$GDB' is the global database name.
- **Ignore** if you have your own backup options and do not want Cloud Control to perform a backup of your database.

 Note:

Click the **Lock** icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Select **Upgrade Options** from the left panel.

18. In the Upgrade Options section, you can do the following:
 - Select **Recompile invalid objects at the end of upgrade** to make valid database objects for the new database version.
 - If you are upgrading to database version 11.2.0.2 or higher, you will be able to set the time zone upgrade option. You can select **Upgrade Time Zone Version and Timestamp with Time Zone data**.
 - Select **Gather Statistics Before Upgrade**.
 - If archive logging has been turned on for the database, then you have the option to disable Archiving and flashback logging for each database.
19. In the Pre and Post Upgrade Script section, specify custom scripts to run on the database before or after upgrade. Select a component from the software library that contains the SQL script to be executed before upgrading the database for each of the following scripts:
 - Pre Upgrade SQL Script
 - Post Upgrade SQL Script
 - Post Upgrade Perl Script

Click **Next**.

20. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the upgrade, if any. Click **Next**.
21. In the Schedule page, specify a **Deployment Instance Name** and schedule for the upgrade job. If you want to run the procedure immediately, then retain the default selection, that is, Immediately. If you want to run the job later, then select Later and provide time zone, start date, and start time details. Specify a Grace Period, a duration after the start period for Cloud Control to attempt to start the upgrade if it cannot start at the specified time.

In the Breakpoint section, you can set the breakpoint by selecting **Set Breakpoint**, and then selecting which step you want the breakpoint to be after, from the **Set Breakpoint After** list.

Setting these breakpoints appropriately can reduce the actual downtime of the upgrade process.

For example, you can set the breakpoint before the Upgrade Cluster Database step as downtime is application only during the actual upgrade process and not during the software installation.

You can also change the breakpoint and resume the execution from the Database Upgrade Instance Tracking page.

In the Set Notification Details section, select the events for which you want to be notified.

Click **Next**.

22. In the Review page, verify that the details you have selected are correctly displayed. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

To save the deployment procedure for future use, click **Save**.

To submit the deployment procedure, click **Submit**. When the deployment procedure is submitted for execution, the database upgrade instance tracking page is displayed. You can also navigate to this page by clicking the procedure instance in the Job Activity page.

23. Submit the configured Database Upgrade procedure after providing values for the editable fields. After you have submitted the procedure, the summary for the running procedure is displayed.
24. In the Upgrade Oracle Database procedure execution page, in the Current Run tab, view the upgrade job steps and status.
25. If you have specified a breakpoint, the procedure execution will pause at the step specified. Once the execution pauses, you can do either of the following using the **Run to step** list.
 - Set a new breakpoint by selecting an appropriate step
 - Select **All remaining steps**.

You can also perform the following actions: The possible actions are Stop, Resume, Suspend, Cleanup, Resubmit, and Skip Step. Click **Resubmit** to resubmit the current instance for execution.

- **Suspend:** to suspend the upgrade deployment procedure.
- **Resume:** to resume the upgrade deployment procedure after you have suspended it.
- **Stop:** to stop the upgrade deployment procedure.
- **Restore:** to do a rollback of the Grid Infrastructure. This submits another deployment procedure.

 Note:

The rollback instance can be tracked in the same page as that of the upgrade execution run. This can be used where there is a fatal failure during the GI rollback.

- **Retry:** to retry the failed step. This option can be used to retry the prerequisite step that might have failed in the previous run (The fixups were performed outside the flow),
 - **Ignore:** to ignore a failed step.
 - **Resubmit:** to resubmit the upgrade deployment procedure (usually after fixing errors found). In general, this should be used when you are unable to perform a retry on the previous run.
26. If a step has status Failed, click **View Log**. The Job Run for the step is listed. Click **Show** in the Details column to view the entire log. Fix the error and click **Retry**.
 27. After the procedure execution is completed, click on the **Database** menu and verify that the newly upgraded databases appear as Cloud Control targets

Upgrading Oracle Clusterware Using Deployment Procedure

To upgrade an Oracle Clusterware using the deployment procedure, follow these steps:

1. In the designer role, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning console, select the **Upgrade Oracle Database** Deployment Procedure and click **Launch**. The Database Upgrade wizard is launched.
3. On the Targets page, in the Select Targets section, select **Clusterware** as the Target Type.

 Note:

The Clusterware option enables you to upgrade the Clusterware and optionally the underlying Automatic Storage Management target as part of the same process.

4. Select the version that you want the database to be upgraded to from the **To** list.
5. Click **Add Targets**.
6. In the Select Targets for Upgrade dialog box, click on the Cluster search icon to search for the cluster.

In the Search and Select: Cluster Target dialog box, select the cluster that you want to upgrade, and then click **Select**.

 Note:

If the database you want to upgrade does not appear in the table, verify that the database is available and there are no metrics collection errors for the target.

 Note:

If the cluster database that you selected is a parent cluster, then all the child nodes such as HAS, ASM, and cluster databases are automatically selected.

Click **OK**.

7. On the Targets page, click **Save**.
8. In the Save Procedure dialog box, enter a description.

For example:

Upgrade Oracle Database Procedure 1

Click **Save**.

9. An Information dialog box appears. Click **OK**.

Click **Log Out**.

10. The designer role is completed. Now, log in with the Operator credentials.

11. From the **Enterprise Menu**, select **Provisioning and Patching**, and then **Database Provisioning**.
12. On the Database Provisioning page, all operations except for Launch are disabled for the Operator. The upgrade database deployment procedure that was saved is selected. Click **Launch**.
13. On the Targets page, click **Next**.
14. On the Software Details page, in the Grid Infrastructure section, click the Search icon to select a Grid Infrastructure software from the Software Library to create a new Oracle Home.

**Note:**

The software may be stored as a gold image or a zipped up installation media in the Software Library.

From the dialog box that appears, select a Grid Infrastructure software, and then click **Select**.

15. In the Working Directory section, specify the working directory on each target host. A working directory on the host target is required to stage files during Grid Infrastructure or Database installation or upgrade. Ensure that all the hosts have read write permission on the working directory specified.

**Note:**

Click on the Lock icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

16. In the Credentials page, specify the Operating System credentials for Grid Infrastructure and the Privileged Operating System credentials (run as root). If you choose to specify Preferred Credentials, select either Normal Host or Privileged Host credentials. For Named Credentials, you can specify the same or different credentials for Oracle homes.

If you have not set Named Credentials, click the plus sign (+) in the Credentials section. In the Add New Database Credentials popup, specify the **User name**, **Password**, **Role**, and specify the Save Details. Select **Run As** and specify **root**. Click **OK**.

Click on the Lock icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Click **Next**.

17. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the upgrade, if any. Click **Next**.
18. In the Schedule page, specify a **Deployment Instance Name** and schedule for the upgrade job. If you want to run the procedure immediately, then retain the default selection, that is, Immediately. If you want to run the job later, then select Lat-

er and provide time zone, start date, and start time details. Specify a Grace Period, a duration after the start period for Cloud Control to attempt to start the upgrade if it cannot start at the specified time.

In the Breakpoint section, you can set the breakpoint by selecting **Set Breakpoint**, and then selecting which step you want the breakpoint to be after, from the Set Breakpoint After list .

Setting these breakpoints appropriately can reduce the actual downtime of the upgrade process.

For example, you can set the breakpoint before the Upgrade Clusterware step as downtime is application only during the actual upgrade process and not during the software installation,

Downtime is applicable only during the actual upgrade and not during the software installation. The downtime is only for the actual upgrade process itself. There is a step in the DP which says Upgrade Clusterware and thats the step that will require a downtime.

You can also change the breakpoint and resume the execution from the Database Upgrade Instance Tracking page.

In the Set Notification Details section, select the events for which you want to be notified.

Click **Next**.

19. In the Review page, verify that the details you have selected are correctly displayed. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

To save the deployment procedure for future use, click **Save**.

To submit the deployment procedure, click **Submit**. When the deployment procedure is submitted for execution, the database upgrade instance tracking page is displayed. You can also navigate to this page by clicking the procedure instance in the Job Activity page.

20. Submit the configured Database Upgrade procedure after providing values for the editable fields. After you have submitted the procedure, the summary for the running procedure is displayed.
21. In the Upgrade Oracle Database procedure execution page, in the Current Run tab, view the upgrade job steps and status.
22. If you have specified a breakpoint, the procedure execution will pause at the step specified. Once the execution pauses, you can do either of the following using the **Run to step** list.

- Set a new breakpoint by selecting an appropriate step
- Select **All remaining steps**.

You can also perform the following actions: The possible actions are Stop, Resume, Suspend, Cleanup, Resubmit, and Skip Step. Click **Resubmit** to resubmit the current instance for execution.

- **Suspend:** to suspend the upgrade deployment procedure.
- **Resume:** to resume the upgrade deployment procedure after you have suspended it.
- **Stop:** to stop the upgrade deployment procedure.

- **Restore:** to do a rollback of the Grid Infrastructure. This submits another deployment procedure.

 Note:

The rollback instance can be tracked in the same page as that of the upgrade execution run. This can be used where there is a fatal failure during the GI rollback.

- **Retry:** to retry the failed step. This option can be used to retry the prerequisite step that might have failed in the previous run (The fixups were performed outside the flow),
 - **Ignore:** to ignore a failed step.
 - **Resubmit:** to resubmit the upgrade deployment procedure (usually after fixing errors found). In general, this should only be used when you are unable to perform a retry on the previous run.
23. If a step has status Failed, click **View Log**. The Job Run for the step is listed. Click **Show** in the Details column to view the entire log. Fix the error and click **Retry**.
24. After the procedure execution is completed, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly upgraded databases appear as Cloud Control targets.

Upgrading Oracle Database Instance Using Deployment Procedure

To upgrade an Oracle database instance using the deployment procedure, follow these steps:

1. In the designer role, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. In the Database Provisioning console, select the **Upgrade Oracle Database** Deployment Procedure and click **Launch**. The Database Upgrade wizard is launched.
3. On the Targets page, in the Select Targets section, select **Single Instance Database** as the Target Type.

 Note:

The Single Instance Database option enables you to upgrade the Single Instance Database and optionally the underlying Automatic Storage Management target and the managing High Availability Service as part of the same process.

4. Select the version that you want the database to be upgraded to from the **To** list.
5. Click **Add Targets**.
6. In the Select Targets for Upgrade dialog box, select the Source Version and the Platform of the database instance that you want to upgrade. Click **Search**.

In the Search and Select: Cluster Target dialog box, select the database instance target that you want to upgrade.

 Note:

If the database you want to upgrade does not appear in the table, verify that the database is available and there are no metrics collection errors for the target.

You can choose to upgrade the listeners by selecting **Upgrade Listeners**. The selected listeners in the source Oracle Home will be migrated and restarted in the destination Oracle Home of the database.

Click **OK**.

7. On the Targets page, click **Save**.
8. In the Save Procedure dialog box, enter a description.

For example:

Upgrade Oracle Database Procedure 1

Click **Save**.

9. An Information dialog box appears. Click **OK**.
Click **Log Out**.
10. The designer role is completed. Now, log in with the Operator credentials.
11. From the **Enterprise Menu**, select **Provisioning and Patching**, and then **Database Provisioning**.
12. On the Database Provisioning page, all operations except for Launch are disabled for the Operator. The upgrade database deployment procedure that was saved is selected. Click **Launch**.
13. On the Targets page, click **Next**.
14. In the Oracle Database section, select **Upgrade Database Instance only**. Specify the Database Oracle Home location.

In the Working Directory section, specify the working directory on each target host. A working directory on the host target is required to stage files during Grid Infrastructure or Database installation or upgrade. Ensure that all the hosts have read write permission on the working directory specified.

 Note:

Click on the Lock icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Click **Next**.

15. In the Credentials page, specify the Operating System credentials for Database, Privileged Operating System credentials (run as root), and Database credentials. If

you choose to specify Preferred Credentials, select either Normal Host or Privileged Host credentials. For Named Credentials, you can specify the same or different credentials for Oracle homes.

If you have not set Named Credentials, click the plus sign (+) in the Credentials section. In the Add New Database Credentials popup, specify the **User name**, **Password**, **Role**, and specify the Save Details. Select **Run As** and specify **root**. Click **OK**.

Click on the Lock icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Click **Next**.

16. In the Configuration Details page, the Backup and Restore settings option in the Restore Strategy is selected by default.

In the Restore Strategy section, you can select:

- **Backup and Restore Settings Only** to restore configuration changes made during database upgrade and not actual data, in case upgrade fails.
- **Create RMAN backup Before Upgrade** and enter the backup location.

 Note:

If you perform a backup using RMAN then you can restore the entire database irrespective of datafiles on filesystem or ASM. Backup using RMAN is available when upgrading to 12c only.

- **Use Existing RMAN Backup** where the latest RMAN backup will be used.
- **Use Flashback and Guaranteed Restore Point**

 Note:

The source database version should be from 11g onwards. The target home version should be 12c

- **Full Backup** to backup the database and restore configuration and oratab settings if upgrade fails. The backup location is, by default, \$ORACLE_BASE/admin/\$GDB/backup where '\$GDB' is the global database name.
- **Ignore** if you have your own backup options and do not want Cloud Control to perform a backup of your database.

 Note:

Click on the Lock icon to lock the fields that you do not want to be editable. These fields will not be available for editing in the operator role.

Select **Upgrade Options** from the left panel.

17. In the Upgrade Options section, you can do the following:

- Select **Recompile invalid objects at the end of upgrade** to make valid database objects for the new database version.
 - If you are upgrading to database version 11.2.0.2 or higher, you will be able to set the time zone upgrade option. You can select **Upgrade Time Zone Version and Timestamp with Time Zone data**.
 - **Select Gather Statistics Before Upgrade**
 - If archive logging has been turned on for the database, then you have the option to disable Archiving and flashback logging for each database.
18. In the Pre and Post Upgrade Script section, specify custom scripts to run on the database before or after upgrade. Select a component from the software library that contains the SQL script to be executed before upgrading the database for each of the following scripts:
- Pre Upgrade SQL Script
 - Post Upgrade SQL Script
 - Post Upgrade Perl Script

Click **Next**.

19. The Custom Properties page will be displayed only for user customized deployment procedures that require custom parameters. Specify custom properties for the upgrade, if any. Click **Next**.
20. In the Schedule page, specify a **Deployment Instance Name** and schedule for the upgrade job. If you want to run the procedure immediately, then retain the default selection, that is, Immediately. If you want to run the job later, then select Later and provide time zone, start date, and start time details. Specify a Grace Period, a duration after the start period for Cloud Control to attempt to start the upgrade if it cannot start at the specified time.

In the Breakpoint section, you can set the breakpoint by selecting **Set Breakpoint**, and then selecting which step you want the breakpoint to be after, from the Set Breakpoint After list .

Setting these breakpoints appropriately can reduce the actual downtime of the upgrade process.

For example, you can set the breakpoint before the Upgrade Database Instance step as downtime is application only during the actual upgrade process and not during the software installation,

You can also change the breakpoint and resume the execution from the Database Upgrade Instance Tracking page.

In the Set Notification Details section, select the events for which you want to be notified.

Click **Next**.

21. In the Review page, verify that the details you have selected are correctly displayed. If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

To save the deployment procedure for future use, click **Save**.

To submit the deployment procedure, click **Submit**. When the deployment procedure is submitted for execution, the database upgrade instance tracking page is displayed. You can also navigate to this page by clicking the procedure instance in the Job Activity page.

22. Submit the configured Database Upgrade procedure after providing values for the editable fields. After you have submitted the procedure, the summary for the running procedure is displayed.
23. In the Upgrade Oracle Database procedure execution page, in the Current Run tab, view the upgrade job steps and status.
24. If you have specified a breakpoint, the procedure execution will pause at the step specified. After the execution pauses, you can do either of the following using the **Run to step** list.
 - Set a new breakpoint by selecting an appropriate step
 - Select **All remaining steps**.

You can also perform the following actions: The possible actions are Stop, Resume, Suspend, Cleanup, Resubmit, and Skip Step. Click **Resubmit** to resubmit the current instance for execution.

- **Suspend:** to suspend the upgrade deployment procedure.
- **Resume:** to resume the upgrade deployment procedure after you have suspended it.
- **Stop:** to stop the upgrade deployment procedure.
- **Restore:** to do a rollback of the Grid Infrastructure. This submits another deployment procedure.

 Note:

The rollback instance can be tracked in the same page as that of the upgrade execution run. This can be used where there is a fatal failure during the GI rollback.

- **Retry:** to retry the failed step. This option can be used to retry the prerequisite step that might have failed in the previous run (The fixups were performed outside the flow),
 - **Ignore:** to ignore a failed step.
 - **Resubmit:** to resubmit the upgrade deployment procedure (usually after fixing errors found). In general, this should only be used when you are unable to perform a retry on the previous run.
25. If a step has status Failed, click **View Log**. The Job Run for the step is listed. Click **Show** in the Details column to view the entire log. Fix the error and click **Retry**.
 26. After the procedure execution is completed, click on the **Database** menu and verify that the newly upgraded databases appear as Cloud Control targets

Upgrading an Oracle Database or Oracle RAC Database Instance Using the Database Upgrade Wizard

This section describes how you can use the wizard to upgrade one single instance database or one Oracle RAC database instance at a time.

Upgrading enables you to access the latest technology, thereby increasing efficiency and providing secure data management for your applications. Ensure that you follow these steps to plan, prepare, and upgrade to make the upgrade process simpler, faster, and more predictable from start to finish.

The Database Upgrade feature allows you to submit a database upgrade job remotely from Enterprise Manager or schedule it for later execution.

Use this wizard to upgrade a standalone database where the Oracle Home has been installed or upgraded. For database upgrade automation or when Oracle Homes are not installed or upgraded, use the Upgrade Database deployment procedure explained in [Upgrading Oracle Databases Using Deployment Procedures](#).

**Note:**

Since mass upgrade of Oracle RAC database is not supported at the moment, Oracle recommends that you use the wizard described in this section to upgrade one Oracle RAC database instance at a time.

This section covers the following:

- [Meeting the Prerequisites](#)
- [Performing the Upgrade Procedure](#)

Meeting the Prerequisites

Before you upgrade an Oracle Database Instance, follow these prerequisites:

- The database version must be 10.2.0.4 or above for upgrade to 11g or 12c.
- For Oracle Real Application Clusters databases, if you select an Oracle RAC database instance and start the database upgrade process, it will upgrade the entire cluster database.
- If OS authentication is not turned on, SYSDBA credentials are required for the upgrade.
- Database to be upgraded must be up and running.
- Ensure that you have DBA privileges to run this procedure.

Performing the Upgrade Procedure

To upgrade an Oracle Database Instance, follow these steps:

1. From the **Enterprise** menu, select **Targets**, then select **Database**. In the Databases page, select the source database to be upgraded.
2. In the Database Instance home page, from the **Oracle Database** menu, select **Provisioning**, then select **Upgrade Database**.

**Note:**

For single instance database instances, you will see another menu option to upgrade the Oracle home and the instance. If you select that option, you will be taken to the wizard described in [Upgrading an Oracle Database or Oracle RAC Database Instance Using the Database Upgrade Wizard](#), however, only the database instance, from where you navigated to the wizard, will be pre-selected for upgraded.

3. Specify the Database user and password credentials and click **Continue**. The Database Upgrade wizard is launched.
4. In the Oracle Home page, select the **New Oracle Home** where you want the new Oracle Home for the upgrade to be installed, based on the version of the database to be upgraded.

If the Oracle Home is not a discovered target in Cloud Control, either discover the Oracle Home using the Cloud Control Discovery feature and then initiate the upgrade process or type the path of the Oracle Home manually. For Oracle Real Application Clusters databases, specify the Oracle RAC home.

For information about discovering targets in Cloud Control, see [Discovering Hosts and Software Deployments](#).

When specifying the new Oracle Home, you must have DBA permissions on both the source and destination Oracle Homes and these Oracle Homes must reside on the same host.

5. In the Oracle Home Credentials section, specify the host credentials. Host credentials must have DBA privileges and can be **Preferred Credentials**, or **Named Credentials**, or, you can select **Enter Credentials** and specify the user name and password and save it. Click **More Details** to view details about the host credentials you have selected. The specified Oracle Home credentials should have privileges on both the source database Oracle Home and the new Oracle Home. Click **Test** to verify that the credentials have the required privileges. If you are using Named Credentials, ensure that these are user and password credentials, else they will not be supported in Cloud Control.
6. Click **Next**. The errors and warnings from the prerequisite checks are displayed. Fix all errors and warnings and revalidate. Click **OK** to proceed to next step.
7. In the Options page, the Diagnostics destination field is displayed only for database upgrade from version 10.2.x to 11.1.0.6. The diagnostic destination is defaulted to Oracle Base and all diagnostic and trace files are stored at this location.

If you are upgrading from version 11.1.0.7 or higher to 11.2.x, the diagnostic destination field does not appear.

8. The **Upgrade parallelism** is set at a default value of 4. The maximum value is 8. This feature is available only when you are upgrading the database to 12c.

If archive logging has been turned on for the database, then you have the option to disable or **Keep archive logging enabled during upgrade**.

If flash recovery area has been configured for the database, then the Flash Recovery section will be displayed. Specify **Flash Recovery Area Location** and provide an adequate space for **Size**.

9. If you are upgrading to database version 11.2.0.2 or higher, you will be able to set the time zone upgrade option. You can select to **Upgrade Time Zone Version and Timestamp with Time Zone data**.

You can select **Gather statistics** if you want to gather optimizer statistics on fixed tables. This helps the optimizer to generate good execution plans. It is recommended that you gather fixed object statistics before you upgrade the database.

You can also select **Make user tablespaces readonly**. This makes the tablespaces read-only during the database upgrade, and it reverts back to read-write after the upgrade is done.

 Note:

The **Gather statistics** and **Make user tablespaces readonly** options are available only when upgrading to 12c.

 Note:

If the database has ASM configured with it, the Backup section will not be displayed.

10. In the Backup section, you can select:

- **Restore Settings Only** to restore configuration changes made during database upgrade and not actual data, in case upgrade fails.
- **Perform full backup before upgrade and restore upon failure** to restore oratab configuration. Specify a file system location for **Backup Location**. The credentials that you have specified earlier must have read-write permissions to this location.

 Note:

If you perform a backup using RMAN then you can restore the entire database irrespective of datafiles on filesystem or ASM. Backup using RMAN is available when upgrading to 12c only.

- **Use an existing RMAN backup to restore database** enables you to restore a database using the latest RMAN backup. This option is specific to upgrading a database to 12c.
- **Use guaranteed restore point to flashback the database.**

 Note:

The source database version should be from 11g onwards. The target home version should be 12c.

- **None** if you do not want to specify a database backup.



Note:

Starting with Oracle Database 12c Release 1 (12.1), Oracle Database supports the use of Oracle Home User, specified at the time of Oracle Database installation. This Oracle Home user is the owner of Oracle services that run from Oracle home and cannot be changed post installation. Different Oracle homes on a system can share the same Oracle Home User or use different Oracle Home User names.

An Oracle Home User is different from an Oracle Installation User. An Oracle Installation User is the user who needs administrative privileges to install Oracle products. An Oracle Home User is a low-privileged Windows User Account specified during installation that runs most of the Windows services required by Oracle for the Oracle home. For more information about Oracle Home User, see *Oracle Database Platform Guide*.

For database version 12.1 or higher, for Microsoft Windows operating systems, the database services will be configured for the Microsoft Windows user specified during Oracle home installation. This user will own all services run by Oracle software.

In the Oracle Home Windows User Credentials section, specify the host credentials for the Microsoft Windows user account to configure database services. Select existing named credentials or specify new credentials. To specify new credentials, provide the user name and password. You can also save these credentials and set them as preferred credentials.

In the Advanced section, specify the custom SQL scripts you want to run before and after the database upgrade. Copy these scripts to the host file system and select them. If your custom scripts are stored as a component in the Software Library, select **Select these scripts from the Software Library** and then browse the Software Library for these scripts. During execution, the main file specified in the Software Library component will be run. So, if you want to run a set of scripts, organize them in the main script file and specify the main script in the Software Library component.

11. Select **Recompile invalid objects at the end of upgrade** to make valid database objects for the new database version. Setting a higher **Degree of Parallelism** will ensure faster recompilation of objects. The default setting is the number of CPU count of the host.

Click **Next**.

12. The Listeners page is displayed only for single instance database upgrade. In the Listeners page, listeners that are registered with Oracle Restart and those that are running in the new Oracle Home are displayed. You can create a new listener or migrate your existing listener manually and then upgrade the database. If you create a new listener, the listener will then be an Cloud Control target and will be monitored. If you migrate your existing listener, the upgrade job will register the database with the listener.

If you have listeners running in the source Oracle Home and need to maintain the same listener port after upgrade, migrate your listener manually to the new Oracle Home first.

For Oracle Real Application Clusters database, the upgraded database will be registered with the Clusterware listener automatically and the Listeners page will not appear.

To add a new listener, specify the **Name** and **Port Number**.

Click **Next**.

13. In the Schedule page, edit or retain the Job **Name** and **Description** for the database upgrade. If you want to run the job immediately, then retain the default selection, that is, Immediately. If you want to run the job later, then select Later and provide time zone, start date, and start time details. Specify a Grace Period, a duration after the start period for Cloud Control to attempt to start the upgrade job if it cannot start at the specified time.

Select **Blackout the database target in Enterprise Manager during upgrade** if you do not want the database to be monitored and alerts to be raised by Cloud Control during the upgrade.

Click **Next**.

14. In the Review page, ensure that you review all warnings generated in the Validation Summary. Click the Validation Summary icon to view validation results and severity and action taken for any warnings. Verify that the details you have provided for the upgrade job appear correctly and then click **Submit Job** to run the job 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. Click **Save** to save the deployment procedure for future deployment.
15. After you have submitted the job, the Database Upgrade Job page with the summary for the running job will be displayed. In the Jobs page, view the job summary and the list of steps and view their status.

 Note:

If the database upgrade fails for any reason and you have not specified a backup option in the Database Upgrade wizard, restore the database manually and perform the upgrade again. If the database upgrade succeeded, but post upgrade custom scripts did not run, then the database will not be restored since upgrade has succeeded.

16. After the upgrade job is completed successfully, click on the **Targets** menu and select **All Targets** to navigate to the All Targets page and verify that the newly upgraded database is displayed as an Cloud Control target with the correct database version.

Data Guard Rolling Upgrade

This section talks about Data Guard Rolling Upgrade. The details are covered in the following subsections:

- [About Rolling Upgrades](#)
- [Prerequisites to Rolling Upgrades](#)
- [Submitting a Rolling Upgrade Procedure for a Primary Database With One Physical Standby Database](#)

- [Viewing a Running or Completed Rolling Upgrade Procedure in Enterprise Manager](#)

About Rolling Upgrades

A rolling upgrade is an alternative to a standard database upgrade that involves upgrading the databases in a Data Guard configuration. The Data Guard broker is shut-down before starting the rolling upgrade. By utilizing the capability of Data Guard to run the primary and standby databases at different releases, a designated physical standby database can first be converted to a transient logical standby database and upgraded without impacting the availability of the primary database.

A transient logical standby database is a physical standby database that can be transiently converted into a logical standby database for the duration of a rolling upgrade. After the upgrade, the logical standby database is converted back to its original role as a physical standby database.

A switchover to the upgraded standby database is then performed, followed by the upgrade of the original primary database. Applications can continue to run against the new primary database, with database downtime limited to the time taken by the switchover, which is likely to be significantly less than the downtime incurred by a conventional upgrade. Finally, a second switchover is performed back to the original primary database.

Prerequisites to Rolling Upgrades

The following items should be considered before attempting a rolling upgrade:

- Enterprise Manager Cloud Control supports rolling upgrade for database versions 11.2 and higher. An Active Data Guard license is required on the primary database and the physical standby database selected to perform the rolling upgrade.
- The software for the upgrade version must already be installed and the new Oracle Home must already exist prior to performing the rolling upgrade. The Enterprise Manager Cloud Control rolling upgrade process does not include these steps.
- For databases using Automatic Storage Management (ASM), ASM must be upgraded separately prior to performing a rolling upgrade, unless it is already at the required version. For 11.2 Real Application Cluster (RAC) or single-instance Oracle Restart databases, ASM and Grid Infrastructure share the same Oracle Home, so ASM is upgraded when the Grid Infrastructure home is upgraded. (The Grid Infrastructure version must be equal to or greater than the target database upgrade version.)
- The databases that are in the Data Guard configuration must be discovered as Enterprise Manager targets.
- The rolling upgrade workflow recommends performing a backup of the primary database, but the backup is not performed during the rolling upgrade process.
- Automatic rollback of failed rolling upgrades is not supported. However, guaranteed restore points are created on all the databases such that the respective databases can be restored back to these restore points in case of failures.
- Guaranteed restore points are created at appropriate points during the upgrade process, and are dropped once the rolling upgrade process has been successfully completed.

- In order to proceed with a rolling upgrade of an existing Data Guard configuration, at least one physical standby database must already exist in the Data Guard configuration and be discovered as an Enterprise Manager target.
- Unless explicitly specified, all required role changes are automatically performed during the rolling upgrade process to ensure that the original roles of the databases are restored after the successful completion of the rolling upgrade. An option is available to pause before the 1st switchover, and if there are no bystander standby databases in the configuration, there is an option to stop after the first switchover if you specify not to return the databases to their original roles.
- Changes to unsupported data type objects are temporarily suspended for the transient logical standby database for the duration of the rolling upgrade process.
- In order to proceed with a rolling upgrade to a particular database version, all the bystander logical standby databases in the configuration must first be upgraded to that version.

Submitting a Rolling Upgrade Procedure for a Primary Database With One Physical Standby Database

To perform a rolling upgrade of a primary database in an existing Data Guard configuration with one physical standby database, follow these steps:

1. From the home page of the primary database, choose **Provisioning** from the **Oracle Database** or **Cluster Database** menu, then select **Upgrade Database**.

If you are submitting a Rolling Upgrade procedure for databases with a primary cluster database, the menu option shows Cluster Database, otherwise it appears as Oracle Database.

If you have not yet provided the credentials, the credentials page displays where you can enter them. Enterprise Manager then displays the first page of the Database Rolling Upgrade wizard, the Select Standby page, which displays the list of all the standby databases associated with the primary database that can be selected for the rolling upgrade process. The Select Standby page also displays an overall summary of the Data Guard configuration such as the name of the primary database, Protection Mode, Fast-start Failover status, Data Guard Status, and the availability status of all the physical standby databases and their respective host names.

Note:

Oracle recommends that you perform a full backup of the primary database before you advance in the wizard. An Active Data Guard license is required on the primary database and the physical standby database selected to perform the rolling upgrade.

2. On the Select Standby page, select the physical standby database to use for the rolling upgrade from the physical standby database table, then click **Next**. All other physical standby databases will be upgraded individually after the rolling upgrade is complete.

Enterprise Manager performs Data Guard-related pre-requisite checks on both the primary database and the selected physical standby database. It then displays the Primary Oracle Home page which shows the current Oracle Home location and

version, and allows you to enter the new Oracle Home location and to provide the relevant host credentials.

Database Rolling Upgrade: Select Standby

Primary Database

Name: prim
Version: 11.2.0.3.0
Host: [dbhost.us.oracle.com](#)
Data Guard Status: Normal
Protection Mode: Maximum Performance
Fast-Start Follower: Disabled

Physical Standby Databases

Select the physical standby database that will be used for the rolling upgrade. All other physical standby databases will be upgraded via redo apply during the rolling upgrade procedure.

Name	Role	Version	Data Guard Status	Host
dgsr2	Physical Standby Database	11.2.0.3.0		dbhost.us.oracle.com
dgsr1	Physical Standby Database	11.2.0.3.0		dbhost.us.oracle.com

3. On the Primary Oracle Home page, enter the **New Oracle Home** location for the primary database. Click **Next**.

Enterprise Manager displays the Standby Oracle Home page of the Database Rolling Upgrade wizard.

Database Rolling Upgrade: Primary Oracle Home

Primary Database: prim
Current Oracle Home: /scratch/orabase11203/product11.2.0/dbhome_1
Version: 11.2.0.3.0
Host: [dbhost.us.oracle.com](#)

New Oracle Home:

Oracle Home Credentials

Specify credentials for the operating system user who owns the primary database Oracle Home.

Credential: ☒ Named ☐ New

Credential Name:


Attribute	Value
Username	dbhost
Password	*****

[More Details](#)

4. On the Standby Oracle Home page, enter the new Oracle Home location for the selected physical standby database. Click **Next**.

Enterprise Manager performs the pre-requisite checks and storage checks for the physical standby database and then opens the Options page.


If you are submitting a Rolling Upgrade Procedure for a primary database with multiple physical standby databases, in this step Enterprise Manager performs the pre-requisite checks and storage checks for the physical standby database and then opens the Bystander Physical Standby Databases page. On the Bystander Physical Standbys page, you view the list of bystander physical standby databases in the configuration and specify the destination Oracle Homes for them. Enterprise Manager then displays the Options page.



- On the Options page, review or change the default options, then click **Next**.

The Options page displays options that allow you to pause before the first switchover, stop after the first switchover (this option displays only when there are no by-stander standby databases), recompile invalid objects, and specify either a common diagnostic destination locations for all databases or specify individual locations for each of them.


Enterprise Manager displays the Listeners page of the Database Rolling Upgrade wizard.



6. On the Listeners page, choose the default Common Listener listed or select from a list of common listeners on the primary and the standby database hosts. The upgraded databases will be registered with the selected listener. You can also specify a new listener name and port to create and register the respective databases with the new listener.

The Listeners page provides an option to select from a list of common listeners on the primary and the standby database hosts (common listeners refer to those with the same name and port) for selection along with an option to create a new common listener.

Click **Next**. Enterprise Manager displays the Procedure Details page.




- On the Procedure Details page, use the default procedure **Name** or enter the name and **Description** of the procedure to be submitted.

Click **Next**. Enterprise Manager displays the Review page of the Database Rolling Upgrade wizard.

- On the Review page, review the summary of selections you have made and click **Submit**.

The Review page displays a **Procedure Summary** and an **Option Summary** of the selections you chose, and then shows a **General Summary** that provides information about the listener along with a **Components** and **Parameters** section.



Viewing a Running or Completed Rolling Upgrade Procedure in Enterprise Manager

To view a running, failed or completed rolling upgrade procedure in Enterprise Manager, follow these steps:

- From the **Enterprise** menu, choose **Provisioning and Patching**, then select **Procedure Activity** from the **Activity** menu.

Enterprise Manager displays the Deployment Procedure Manager page.

- Choose the **Procedure Activity** tab.

You can use this page to view all deployment procedures in various stages of their lifecycle. You can set the refresh settings to update the page automatically. Click the link in the Run column to view more details on that run.

Part VI

Database Security

This part contains the following chapters:

- [Managing Oracle Audit Vault and Database Firewall](#)
- [Using Oracle Data Redaction](#)
- [Managing Oracle Database Vault and Privilege Analysis](#)

Managing Oracle Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall (AVDF) secures databases and other critical components of IT infrastructure. It provides a database firewall that can monitor database activity and block SQL statements on the network based on a firewall policy. It also collects audit data, and ensures that the data is available in audit reports.

You can manage and monitor Oracle AVDF components in Enterprise Manager Cloud Control (Cloud Control) using the Oracle AVDF plug-in. For information on how to install this plug-in and manage Oracle AVDF using Cloud Control, see *Oracle Enterprise Manager System Monitoring Plug-in User's Guide for Audit Vault and Database Firewall*.

Using Oracle Data Redaction

Oracle Data Redaction is an Oracle Database security feature that enables you to mask (redact) the data that is returned from queries issued by applications. Enterprise Manager Cloud Control (Cloud Control) provides a user interface for creating and managing Oracle Data Redaction policies and formats. You can perform these tasks using the Data Redaction page.

Figure 21-1 Data Redaction Page

The screenshot displays the 'Data Redaction' page in the Oracle Enterprise Manager Cloud Control interface. At the top, there is a navigation bar with tabs for 'Policies' and 'Formats'. Below the navigation bar, a search section titled 'Search Data Redaction Policies' contains three input fields: 'Schema' with a placeholder '%', 'Table/View' with a placeholder '%', and 'Policy Name' with a placeholder '%'. A 'Go' button is located below these fields. Below the search section, there is a 'Policies' section with a table listing existing data redaction policies. The table has columns for 'Schema', 'Table/View', 'Policy Name', 'Enabled', and 'Redacted Columns'. Two policies are listed: 'test_policy_1' and 'test_policy_2', both of which are enabled and have one redacted column.

Schema	Table/View	Policy Name	Enabled	Redacted Columns
APEX_040200	WWW_FLOW_POP...	test_policy_1	✓	1
DBSNMP	MGMT_BSLN_METRIC	test_policy_2	✓	1

For detailed information on using Oracle Data Redaction, see *Oracle Database Advanced Security Guide*.

Managing Oracle Database Vault and Privilege Analysis

Oracle Database Vault provides powerful security controls to help protect database application data from unauthorized access, and comply with privacy and regulatory requirements. Using Oracle Database Vault, you can deploy controls to block privileged account access to database application data, and control sensitive operations within the database. Oracle Database Vault with Oracle Database 12c includes a feature called privilege analysis that helps you increase the security of your database applications and operations. Privilege analysis policies reduce the attack surface of database applications and increase operational security by identifying used and unused privileges.

You can manage Oracle Database Vault and privilege analysis policies using Enterprise Manager Cloud Control. For detailed information on how to do this, see *Oracle Database Vault Administrator's Guide*.

Managing Oracle Transparent Data Encryption

Transparent Data Encryption (TDE) encrypts sensitive data stored in data files. Encrypted data is transparently decrypted for a database user or application that has access to data.

Oracle Database uses authentication, authorization, and auditing mechanisms to secure data in the database, but not in the operating system data files where data is stored. To protect these data files, Oracle Database provides Transparent Data Encryption (TDE). TDE encrypts sensitive data stored in data files. Encrypted data is transparently decrypted for a database user or application that has access to data. TDE helps protect data stored on media in the event that the storage media or data file gets stolen.

Online tablespace encryption is an encryption process wherein a tablespace which is online and being used, is encrypted by creating alternate data files of the tablespace thereby causing no data loss. Offline tablespace encryption is an encryption process wherein a tablespace is encrypted after placing the tablespace in offline mode.

For more information on database encryption and TDE, see *Oracle Database Advanced Security Administrator's Guide*.

Figure 23-1 Transparent Data Encryption Home Page

Oracle Advanced Security - Transparent Data Encryption

Frequently Used Configuration Option

Current Master Key in use (days) 0

Keystore Status OPEN

[Rekey](#) [Close](#)

Overview

Oracle Advanced Security Transparent Data Encryption (TDE) enables you to encrypt data to prevent unauthorized access from outside of the database environment. TDE transparently encrypts data at the user or application level. The solution is transparent to applications because data is encrypted without the need to be modified to take advantage of this feature. TDE is licensed as part of the Oracle Advanced Security Option.

TDE provides a two-tier encryption key management architecture consisting of data encryption keys, and the data encryption keys are used to encrypt the data. The master encryption keys are stored in a keystore that is outside of the database. The keystore can be Oracle Wallet (a standards-based PKCS12 file) or Oracle Key Vault.

Keystore and Master Keys

Encrypted Objects

Encrypted Tablespaces

Online Operation Offline Operation Refresh

Tablespace Name	Status	Encryption Algorithm
TEST_ONLINE_ENC	ONLINE	AES256
TEST_OFFLINE_ENC...	ONLINE	AES128

Tables with Encrypted Columns

[Detach](#) [Encrypt Table Columns](#)

Table Name	Table Owner	Encrypted Column Count
No data to display.		

Jobs

Jobs in progress Jobs completed

This chapter contains the following sections:

- [Prerequisites for Transparent Data Encryption Operations](#)
- [Encrypting a Tablespace in Online Mode](#)
- [Decrypting a Tablespace in Online Mode](#)
- [Changing the Encryption Algorithm of an Encrypted Tablespace - Rekey](#)
- [Encrypting a Tablespace in Offline Mode](#)
- [Decrypting a Tablespace in Offline Mode](#)

Prerequisites for Transparent Data Encryption Operations


To encrypt or decrypt a tablespace using TDE, meet the following prerequisites:

- Setup the wallet.
- Ensure the Keystore Status is Open in the Transparent Data Encryption home page.
- Ensure that the space available in the database is equivalent or more than the tablespace size that is being encrypted, as auxiliary data files need to be created for online encryption.
- Only Database versions 12.2 and above are supported for online encryption.

Encrypting a Tablespace in Online Mode

To encrypt a tablespace, follow the steps below:

1. From the Database home page, click **Security** and then select **Transparent Data Encryption**.
2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.
3. Click **Online Operation** menu in the Encrypted Tablespaces pane, and select **Encrypt**.
4. In the Encrypt pop-up window, select the tablespace and the encryption algorithm.



5. The **Target path** column indicates the new data files that will be created post-encryption for online encryption. To retain the old data files post the encryption operation, select the relevant check box.
6. You may schedule the encryption for a future time. Select the relevant radio button and configure the scheduled time, if required.
7. Click **OK**.

A job is scheduled to complete the encryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.




Note:

You may need to refresh the view to see the latest status of tablespaces in the Encrypted Objects section and the jobs.

Decrypting a Tablespace in Online Mode

To decrypt an encrypted tablespace, follow the steps below:

1. From the Database home page, click **Security** and then select **Transparent Data Encryption**.
2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.
3. Select the encrypted tablespace from the Encrypted Tablespaces table.
4. Click **Online Operation** menu in the **Encrypted Tablespaces** pane, and select **Decrypt**.




5. In the Decrypt pop-up window, the **Target path** column under Datafiles section indicates the new data files that will be created post-encryption for online encryption. To retain the old data files post the encryption operation select the relevant check box.
6. You may schedule the encryption for a future time. Select the relevant radio button and configure the scheduled time, if required.
7. Click **OK**.

A job is scheduled to complete the encryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.

Changing the Encryption Algorithm of an Encrypted Tablespace - Rekey

To change the encryption algorithm of a tablespace which is also known as Rekey, follow the steps below:

1. From the Database home page, click **Security** and then select **Transparent Data Encryption**.
2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.
3. Click **Online Operation** menu in the Encrypted Tablespaces pane, and select **Rekey**.
4. In the Rekey pop-up window select the new encryption algorithm.




5. The Target path column indicates the new data files that will be created post-encryption for online encryption. To retain the old data files post the encryption operation select the relevant check box.
6. You may schedule the encryption for a future time. Select the relevant radio button and configure the scheduled time, if required.
7. Click **OK**.

A job is scheduled to complete the encryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.

Encrypting a Tablespace in Offline Mode

To encrypt a tablespace in the offline mode, you will have to first switch the tablespace to offline mode. To do so, follow the steps below:

1. From the Database home page, click **Security** and then select **Transparent Data Encryption**.
2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.
3. Click **Offline Operation** menu in the Encrypted Tablespaces pane, and select **Offline**.
4. In the Offline pop-up window, select the tablespace and schedule, and click **OK**.



A job is scheduled to switch the tablespace to offline mode.

After the tablespace is in offline mode, follow the steps below to encrypt the tablespace:

1. Click **Offline Operation** menu, and select **Encrypt**.
2. Select the tablespace and the schedule.

 Note:

In the **Datafiles** section you have an option to selectively choose the data files that need to be encrypted. However, the tablespace cannot be placed in the online mode if all the data files are not encrypted. Moreover, if you do not select the check box all the data files are encrypted and the tablespace is automatically switched to the online mode post-encryption. Do note that in case you select the check box and then select all the data files under the tablespace, all the data files will be encrypted but, the tablespace will not be switched to the online mode.

3. Click **OK**.

Decrypting a Tablespace in Offline Mode


To decrypt a tablespace in the offline mode, you will have to first switch the tablespace to offline mode. To do so, follow the steps below:

1. From the Database home page, click **Security** and then select **Transparent Data Encryption**.
2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.
3. Click **Offline Operation** menu in the Encrypted Tablespaces pane, and select **Offline**.
4. In the Offline pop-up window, select the tablespace and schedule, and click **OK**.

A job is scheduled to switch the tablespace to offline mode.

After the tablespace is in offline mode, follow the steps below to decrypt the tablespace:

1. Select the encrypted tablespace from the **Encrypted Tablespaces** table.
2. Click **Offline Operation** menu in the Encrypted Tablespaces pane, and select **Decrypt**.



3. In the **Datafiles** section you have an option to selectively choose the data files that need to be decrypted. However, the tablespace cannot be placed in the online mode if only some of the data files are decrypted. Moreover, if you do not select the check box all the data files are decrypted and the tablespace is automatically switched to the online mode post-decryption. Do note that in case you select the check box and then select all the data files under the tablespace, all the data files will be decrypted but, the tablespace will not be switched to the online mode automatically.
4. You may schedule the decryption for a future time. Select the relevant radio button and configure the scheduled time, if required.
5. Click **OK**.

A job is scheduled to complete the decryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.

Part VII

Middleware Provisioning

This part contains the following chapters:

- [Overview of Middleware Provisioning](#)
- [Provisioning Fusion Middleware Domain and Oracle Homes](#)
- [Provisioning the SOA Domain and Oracle Homes](#)
- [Provisioning the Service Bus Domain and Oracle Homes](#)
- [Provisioning the Oracle WebCenter Domain and Oracle Homes](#)
- [Middleware Provisioning using the EM CLI](#)
- [Middleware Profiles Using REST APIs](#)
- [Scaling Up / Scaling Out Fusion Middleware Domains](#)
- [Export and Import of Domain Partitions](#)
- [Deploying / Redeploying / Undeploying Java EE Applications](#)
- [Provisioning Coherence Nodes and Clusters](#)
- [Provisioning SOA Artifacts and Composites](#)
- [Provisioning Service Bus Resources](#)

Overview of Middleware Provisioning

Provisioning involves repeatable, reliable, automated, unattended, and scheduled mass-deployment of software, applications, or servers across different platforms, environments, and locations.

Middleware Provisioning involves remotely installing Oracle homes and configuring domains using automated deployment procedures. A deployment procedure is a pre-defined sequence of steps that is meant to perform a set of operations on one or more hosts monitored by Cloud Control.


In particular, this chapter covers the following:

- [Introduction to Middleware Provisioning](#)
- [Oracle Fusion Middleware Provisioning Terminology](#)
- [Supported Use Cases for Middleware Provisioning Procedures](#)

Introduction to Middleware Provisioning

Provisioning is an important solution offered as a part of Lifecycle Management that enables you to provision middleware artifacts such as WebLogic Domain, Java EE Applications, Coherence Nodes and Clusters, SOA Artifacts and Composites, Service Bus Resources, and Oracle WebCenter.

The following figure shows the Middleware Provisioning solutions offered in Cloud Control:



To manage these Deployment Procedures and Profiles effectively in Cloud Control there is a centralised middleware provisioning page which exposes all the features relevant to Middleware Provisioning like creating, viewing, and provisioning profiles and deployment procedures. To access this page, from **Enterprise** menu, select **Provisioning and Patching**, then click **Middleware Provisioning**.

ORACLE Enterprise Manager Cloud Control 13c

Enterprise Targets Favorites Hodge Setup SYSMAN

Page Refreshed Jan 23, 2017 2:22:57 AM PST

Middleware Provisioning

Profiles

View Create Edit Delete Provision Import Partition

Name	Products	Type	Platform	Version	Owner	Description
Domain Profile	Oracle WebLogic Server	WebLogic Domain	Linux x86-64	12.2.1.2.0	SYSMAN	WebLogic domain profile that includes Oracle home
SOA Domain Profile	Oracle SOA	WebLogic Domain	Linux x86-64	12.2.1.2.0	SYSMAN	SOA domain profile that includes Oracle home and metadata
Partition Export	Oracle WebLogic Server	WebLogic Domain Partition	Linux x86-64	12.2.1.2.0	SYSMAN	A profile created by exporting a WebLogic domain partition
Domain Export	Oracle WebLogic Server	WebLogic Domain Partition	Linux x86-64	12.2.1.2.0	SYSMAN	A profile created by exporting a WebLogic domain. It can be imported as a partition.
JRF Installer	Oracle WebLogic Server	Installation Media	Generic Platform	12.2.1.1.0	SYSMAN	Fusion Middleware Infrastructure Installation Media

Deployment Procedures

View Launch Create Like Edit Edit Permission Delete

Name	Description	Version	Owner	Type
Provision Fusion Middleware	This procedure clones and configures a Fusion Middleware Home and/or a Fusion Middleware Domain from the Software Library.	10.0	ORACLE	Middleware Provisioning
Scale up/Scale out Fusion Middleware Domain	This procedure expands a Fusion Middleware Domain.	10.0	ORACLE	Middleware Provisioning
Create a Business Intelligence Provisioning Profile	Create a Business Intelligence Provisioning Profile	2.0	ORACLE	Middleware Provisioning
Business Intelligence Cloning	Cloning Business Intelligence environment based on source Business Intelligence environment profile and user provided destination environment details	2.0	ORACLE	Middleware Provisioning
SOA Artifacts Provisioning	Deployment procedure to provision SOA artifacts	3.8	ORACLE	Middleware Provisioning
Deploy SOA composites	Procedure for provisioning SOA Composites on a selected soainfra domain.	2.6	ORACLE	Middleware Provisioning
Scale Out Business Intelligence	This procedure scales out the existing Business Intelligence domain.	2.0	ORACLE	Middleware Provisioning
Deploy/Undeploy Java EE Applications	One or more Java EE applications can be deployed to, redeployed to, or undeployed from one or more WebLogic Domains.	13.220	ORACLE	Java EE Application Provisioning
Service Bus Resource Provisioning	Procedure for deploying Service Bus resources to an Service Bus Domain	3.0	ORACLE	Service Bus Resource Provisioning
Coherence Node Provisioning	Procedure for adding and updating Coherence Nodes.	1.1	ORACLE	Coherence Node Provisioning

In addition to accessing all the relevant provisioning related features from **Middleware Provisioning** page, you can also launch some of them from the **target** menu. To do so, from **Targets** menu, select **Middleware**. Select a target of type domain, and select

Provisioning, and the options in this menu allow you to **create provisioning profile** based on WebLogic Domain or Oracle Home.

You can create the profile for multiple purposes. Based on the details configured in the selected target type domain, the appropriate profile types are displayed under the provisioning sub-menu.

The screenshot displays the Oracle Enterprise Manager Cloud Control interface. On the left is a search bar and a navigation tree. The main area shows a 'Target Type' pie chart and a 'Status' pie chart. Below these is a table of targets. A context menu is open over the 'iEMGC_EMGC_DOMAIN' target, showing options like 'Home', 'Monitoring', 'Diagnostics', 'Control', 'Logs', 'Information Publisher Reports', 'Dashboard', 'JDBC Data Sources', 'Messaging', 'Web Services', 'Environment', 'Administration', 'System MBean Browser', 'Routing Topology', 'Members', 'Refresh WebLogic Domain', 'Target Management Best Practices', 'WebLogic Server Administration Console', 'Configuration', 'Compliance', 'Provisioning', and 'Target Setup'. The 'Provisioning' option is highlighted, and a sub-menu is visible with options: 'Create Provisioning Profile', 'Clone WebLogic Domain', 'Import WebLogic Partition', 'Scale Up/Scale Out WebLogic Domain', 'Deploy/Undeploy Java EE Applications', 'SOA Artifacts Provisioning', 'Deploy SOA composites', and 'Procedure Activity'.

Target Name	Target Type	Status	Member	Status Summary	N/A
iEMGC_EMGC_DOMAIN	Oracle WebLogic Domain	-	3	-	2
iarm01_base_domain	Oracle WebLogic Domain	-	3	-	-
iarm02_base_domain	Oracle WebLogic Domain	-	3	-	-
iarm03_base_domain	Oracle WebLogic Domain	-	3	-	-
iarm04_SOADomain	Oracle WebLogic Domain	-	13	-	2

The below figure shows the different profile types:

The screenshot shows the 'Create Provisioning Profile: Type' dialog box in Oracle Enterprise Manager Cloud Control. It contains three sections: 'From a WebLogic Domain', 'From an Oracle Home', and 'By Exporting a WebLogic Domain'. Each section provides a brief description of the profile type and the steps to create it.

From a WebLogic Domain
This option enables you to create a provisioning profile from an existing, reference Oracle WebLogic Domain target. Future domains - as well as its underlying Oracle Homes - can be cloned from this type of profile.

From an Oracle Home
This option enables you to create a provisioning profile from an existing, reference Oracle Home in which WebLogic Server software has already been installed and patched. Future WebLogic Server Oracle Homes can be cloned from this type of profile.

By Exporting a WebLogic Domain
This option enables you to create a provisioning profile from an existing domain in order to convert the entire domain into a single partition.

The Middleware Provisioning page is categorized into the following sections:

- [Profiles](#)
- [Deployment Procedures](#)

Profiles

The profiles section lists all the provisioning profiles that you have created and the profiles on which you have been granted access. No profiles are shipped out of the box. You have to create your own profiles based on the provisioning needs. You can:

- Click the profile to view the profile details.
- Filter the profile based on what you want to display in the Profiles table. To do so, from **View** menu, select **Show Profiles**, then click the option that you want to display. For example, if you click **All**, then all the profiles are displayed.
- To clone a WebLogic Domain or an Oracle Home, select a profile and click **Provision**.
- To create a new profile, select an option from the **Create** menu.
- To delete an existing profile, select the profile name, and click **Delete**.
- To create a new partition in an existing domain, select the appropriate profile (i.e. a profile that was created either by exporting an older domain or by exporting a partition) and click **Import Partition**.

Deployment Procedures

The deployment procedures section lists all the Oracle-provided deployment procedures, the Custom Deployment Procedures (CDP) that you have created, and the procedures on which you (the administrator you have logged in as) have been granted access. Select a deployment procedure from the list, and perform any of the following action on it:

- To run a deployment procedure, select the procedure, and click **Launch**.
- To create a copy of an Oracle-supplied deployment procedure, click **Create Like**.
- To delete a User-owned deployment procedure, select the procedure, and click **Delete**. Note that, you can not delete Oracle-provided procedures.
- To edit a deployment procedure, select the procedure and click **Edit**.

Note: You can not edit Oracle-supplied procedures. If you want to edit a procedure, you must first clone an Oracle-supplied procedure using the Create Like option. You can now edit the newly cloned procedure that you own. To do so, click **Edit**.

- To grant other administrators, for example: EM_PROVISIONING_OPERATOR role, access on a particular deployment procedure, click **Edit Permissions**.

Oracle Fusion Middleware Provisioning Terminology

The definitions of terms listed in this section are useful in helping to understand the Middleware concepts presented in this part:

- **WebLogic Domain:** A WebLogic Domain is a logically related group of Java components. A domain includes a special WebLogic Server instance called the Administration Server, which is the central point from which you configure and manage all resources in the domain. Usually, you configure a domain to include additional WebLogic Server instances called Managed Servers. You deploy Java components, such as Web applications, EJBs, and Web services, and other resources, to

the Managed Servers and use the Administration Server for configuration and management purposes only.

- **Administration Server:** The Administration Server operates as the central control entity for the configuration of the entire domain. It maintains the domain's configuration documents and distributes changes in the configuration documents to managed servers. The Administration Server is the central location from where you can monitor all the resources in a domain. Each WebLogic Server domain must have one server instance that acts as the Administration Server.
- **Managed Server:** Managed servers host business applications, application components, Web services, and their associated resources. To optimize performance, managed servers maintain a read-only copy of the domain's configuration and security document. When a managed server starts up, it connects to the domain's Administration Server to synchronize its configuration document with the document that the Administration Server maintains.
- **Node Manager:** Node Manager is a Java utility that runs as a separate process from Oracle WebLogic Server and allows you to perform common operations for a Managed Server, regardless of its location with respect to its Administration Server. While use of Node Manager is optional, it provides valuable benefits if your Oracle WebLogic Server environment hosts applications with high-availability requirements.

If you run Node Manager on a computer that hosts Managed Servers, you can start and stop the Managed Servers remotely using the Administration Console, Fusion Middleware Control, or the command line. Node Manager can also automatically restart a Managed Server after an unexpected failure.

- **WebLogic Server Home:** A WebLogic Server home contains installed files necessary to host a WebLogic Server. The WebLogic Server home directory is a peer of Oracle home directories and resides within the directory structure of the Middleware home.
- **Middleware Home:** A Middleware home is a container for the Oracle WebLogic Server home, and, optionally, one Oracle Common home and one or more Oracle homes. A Middleware home can reside on a local file system or on a remote shared disk that is accessible through NFS.
- **Oracle Home:** An Oracle home contains installed files necessary to host a specific product. For example, the SOA Oracle home contains a directory that contains binary and library files for Oracle SOA Suite. An Oracle home resides within the directory structure of the Middleware home.
- **Cloning:** The process of creating a copy of the WebLogic Domain and the Oracle home binaries present within the domain is referred to as cloning. Typically, cloning is performed at the WebLogic Domain-level. Fusion Middleware Domain cloning can be performed from an existing target or using provisioning profiles.
- **Provisioning Profiles:** A profile is a snapshot of a live WebLogic Domain or Oracle Home, or it can simply contain a set of installation media archives pertaining to the product suite. A profile can be created from the Cloud Control or using Enterprise Manager Command Line Interface or through REST APIs.
- **Gold Image:** The gold image is a single image that includes the binary and library files for an Oracle home.

For Oracle Fusion Middleware 11g, the Middleware Home was the top-level directory that comprised of multiple product-specific Oracle Homes. For example:

```
[user1@<agent host name> middhome]$ ls
Oracle_OSB1 Oracle_SOA1 coherence_3.7 domain-registry.xml logs modules ora-
cle_common osb patch_ocp371 patch_wls1036 registry.dat registry.xml utils
wlserver_10.3
```

For Oracle Fusion Middleware 12c, there is no concept of Middleware Home. In-
fact, the Middleware Home itself functions as the Oracle Home, and middleware
products like SOA, Service Bus, WebCenter are installed within this folder directly.
For example:

```
[user1@<agent host name> OH12JRF]$ ls
OPatch cfgtoollogs crs em install jdeveloper ldap mft nls oep
oraInst.loc oracore oui plsql rdbms root.sh.old root.sh.old.2
soa srvm wlserver
bin coherence css has inventory jlib lib network odi ohs
oracle_common osb perl precomp root.sh root.sh.old.1 slax
sqlplus webgate xdk
```

- **Scaling Up:** When a managed server is added or cloned to a host that already ex-
ists in the domain or cluster.
- **Scaling Out:** When a managed server is added or cloned to a host that is not
present in the domain or cluster.
- **Scaling Down:** The managed server cannot be removed from host that is not
present in the domain or cluster.

Supported Use Cases for Middleware Provisioning Proce- dures

This section lists the supported use cases in the middleware space, provisioning capa-
bilities like Exporting/Importing partitions and Exporting/Importing a weblogic domain,
and the corresponding targets that get provisioned.

In particular, this section covers the following:

- [Provisioning Middleware Domains and Oracle Homes](#)
- [Scaling Up/Out/ WebLogic Server, SOA, Service Bus, and WebCenter Domains](#)
- [Deploying / Redeploying / Undeploying Java EE Applications](#)
- [Provisioning Coherence Nodes and Clusters](#)
- [Provisioning SOA Artifacts](#)
- [Provisioning Service Bus Resources](#)
- [Exporting WebLogic Domain Partition](#)
- [Importing WebLogic Domain Partition](#)
- [Exporting WebLogic Domain](#)

 Note:

To configure a domain with a TCPS enabled database, you must first disable TCPS in the database, provision the domain with TCP, enable TCPS, and re-configure the domain for TCPS mode manually.

For more information, see SSL Configuration in Oracle Fusion Middleware of the *Fusion Middleware Administrator's Guide*.

Provisioning Middleware Domains and Oracle Homes

This table covers the use cases for deploying SOA domain, Service Bus domain, WebLogic domain, WebCenter domain and Oracle homes.

Table 24-1 Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
Provision Fusion Middleware	• Install media based profile • Domain based profile • Oracle home based Profile	WLS 12.2.1.x, 12.1.x, 10.3.x.	Provisioning Fusion Middleware Domain and Oracle Homes
Note: Top end for month the use cases again installed release 12.2.			

Table 24-1 (Cont.) Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
			1 . 3 , y o u m u s t f i r s t a p p l y P a t c h I D 2 6 8 4 9 0 4 4 t o t h e W e b L o g i c S e r

Table 24-1 (Cont.) Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
			version of Fusion Middleware Oracle Platform for Oracle Home · Download the patch for

Table 24-1 (Cont.) Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
			 om M y O r a c l e S u p p o r t W e b s i t e a t h t t p : / / s u p p o r t . o r a c l e . c o m

Table 24-1 (Cont.) Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
			This space contains information not necessary for earliler release of WebLogi

Table 24-1 (Cont.) Provisioning Middleware Domains and Oracle Homes

Deployment Procedure	Use Case	Targets Provisioned	Link
			 Provisioning the SOA Domain and Oracle Homes
Provision Fusion Middleware	- SOA Installation Media based profile - SOA Gold Image based profile - SOA Domain Profile - Existing SOA Middleware Home	SOA Domain 12.2.1.x, 12.1.x, 11.1.1.x	Provisioning the SOA Domain and Oracle Homes
Provision Fusion Middleware	- Service Bus Installation Media based profile - Service Bus Gold Image based profile - Service Bus Domain Profile - Existing Service Bus Middleware Home	Service Bus Domain 12.2.1.x, 12.1.x, 11.1.1.x	Provisioning the Service Bus Domain and Oracle Homes
Provision Fusion Middleware	- WebCenter Installation Media based profile - Webcenter Server Gold Image based profile - WebCenter Domain Profile - Existing WebCenter Home Profile	Oracle WebCenter Portal - 11.1.1.x Oracle WebCenter Content - 11.1.1.x	Provisioning the Oracle WebCenter Domain and Oracle Homes

Steps to access the Certification Matrix

Use the steps below to view the correct certified versions of various components.

1. Access <https://support.oracle.com>

2. Click **Certifications** tab.
3. In the **Certification Search** tab, enter the following details:

Field	Details
Product	Enterprise Manager Base Platform – OMS
Release	13.2.0.0.0

4. Click **Search**.
5. The search result displays the Enterprise Manager Base Platform - OMS 13.2.0.0.0 Certifications.
6. Expand **Middleware** in the search results.
7. The Certification Matrix of various components and provisioning targets along with certified versions is displayed.

Java Prerequisites

When using deployment procedures to provision Oracle Fusion Middleware software, Oracle Enterprise Manager does not automatically provision Java. You must ensure that Java gets installed prior to performing provisioning-related operations.

When provisioning install-media based profiles, consult the certification information available on My Oracle Support in order to determine the vendor and version of Java that the middleware software supports.

Figure 24-1 Install-media based profile provisioning use case

The screenshot shows the Oracle Enterprise Manager provisioning interface. The 'General' tab is selected. Below it, the 'Hosts' section is visible, with a message 'Add the destination hosts where the domain/Oracle Home will be provisioned.' and buttons for '+ Add', 'Delete...', and 'Test'. Below the 'Hosts' section, the 'Middleware' tab is selected. A message reads 'Enter the WebLogic Administrator username and password that will be used for the newly provisioned domain.' The 'Middleware' section contains the following fields:

- * Middleware Base: /u01/oracle/mwBase
- * Java Home: /usr/local/packages/jdk7
- * Administrator Username: weblogic
- * Administrator Password: (empty field)
- * Confirm Administrator Password: (empty field)

At the bottom, the 'Deploy JVM agent' checkbox is checked.

In addition, when provisioning middleware-related profiles that include an Oracle Home, ensure that the destination has the proper vendor and version of Java. The Java version at the destination should preferably be identical or at least be of the same

major version as the Java version used at the original source. The deployment procedure user interface informs you of the Java version used at the original source so that you are aware of which version to use at the destination.

Figure 24-2 Provisioning use case with profile which includes Oracle Home

General

Hosts

Add the destination hosts where the domain/Oracle Home will be provisioned.

+ Add X Delete... Test

Host	Credential
No hosts selected	

Middleware

Enter the WebLogic Administrator username and password that will be used for the newly provisioned domain.

* Java Home /ade_autofs/gd29_3rdparty/htfsdo_generic/JDK7/MAIN/LINL

The selected profile indicates the source Java version as 1.7.0_141. Make sure to select a matching Java version at the destination.

* Administrator Username weblogic

* Administrator Password

* Confirm Administrator Password

Deploy JVMMD agent ☐

Scaling Up/Out/ WebLogic Server, SOA, Service Bus, and WebCenter Domains

This table covers the use cases for scaling an existing SOA Domain, Service Bus Domain, and WebLogic Domain:

Table 24-2 Scaling SOA, Service Bus, WLS, and WebCenter Domains

De- ploy- ment Pro- ce- dure	Use Cas e	Targets Provisioned	Link
Scal- ing WLS Do- main	Scali ng WLS Do- main	WLS 12.2.1.x, 12.1.x, 10.3.x.	Running the Scale Up / Scale Out Middleware Deploy- ment Procedure
Note: To perform the upgrade as a single instance, use the 1			

Table 24-2 (Cont.) Scaling SOA, Service Bus, WLS, and WebCenter Domains

De- ploy- ment Pro- ce- dure	Use Cas e	Targets Provisioned	Link
			2
			.
			2
			.
			1
			.
			3
			,
			y
			o
			u
			m
			u
			s
			t
			f
			i
			r
			s
			t
			a
			p
			p
			l
			y
			p
			a
			t
			c
			h
			2
			6
			8
			4
			9
			0
			4
			4
			t
			o
			t
			h
			e
			W
			e
			b
			L

Table 24-2 (Cont.) Scaling SOA, Service Bus, WLS, and WebCenter Domains

De- ploy- ment Pro- ce- dure	Use Cas e	Targets Provisioned	Link
			o g i c S e r v e r o r F u s i o n M i d d l e w a r e P l a t f o r m h o m e .T h i s p a t

Table 24-2 (Cont.) Scaling SOA, Service Bus, WLS, and WebCenter Domains

De- ploy- ment Pro- ce- dure	Use Cas e	Targets Provisioned	Link
			ch i s n o t n e c e s s a r y a g a i n s t e a r l i e r r e l e a s e s o f W e b L o g i c

Table 24-2 (Cont.) Scaling SOA, Service Bus, WLS, and WebCenter Domains

De- ploy- ment Pro- ce- dure	Use Case	Targets Provisioned	Link
			o r F u s i o n M i d d l e w a r e .
Scal- ing up/ Scale out Mid- dle- ware	Scali ng SOA Do- main	SOA Domain 12c, 11g	Running the Scale Up / Scale Out Middleware Deployment Procedure
Scal- ing Serv- ice Bus Do- main	Scali ng Serv ice Bus Do- main	Service Bus Domain 12c, 11g	Running the Scale Up / Scale Out Middleware Deployment Procedure
Scal- ing Web- Cen- ter Do- main	Scali ng Web Cen- ter Do- main	Oracle WebCenter Portal - 11g Oracle WebCenter Content - 11g	Running the Scale Up / Scale Out Middleware Deployment Procedure

Steps to access the Certification Matrix

Use the steps below to view the correct certified versions of various components.

1. Access <https://support.oracle.com>
2. Click **Certifications** tab.
3. In the **Certification Search** tab, enter the following details:

Field	Details
Product	Enterprise Manager Base Platform – OMS
Release	13.2.0.0.0

4. Click **Search**.
5. The search result displays the Enterprise Manager Base Platform - OMS 13.2.0.0.0 Certifications.
6. Expand **Middleware** in the search results.
7. The Certification Matrix of various components and provisioning targets along with certified versions is displayed.

Java Prerequisites

When using deployment procedures to scale out domains, Oracle Enterprise Manager does not automatically provision Java. You must ensure that Java gets installed prior to performing the scale out operation.

When scaling out a domain, ensure that the destination has the proper vendor and version of Java. The Java version at the destination should preferably be identical or at least be of the same major version as the Java version used at the source.

Deploying / Redeploying / Undeploying Java EE Applications

This table covers the use cases for deploying, undeploying, and redeploying Java EE Application.

Table 24-3 Deploying, Undeploying, or Redeploying Java EE Applications

Deployment Procedure	Use Case	Targets Provisioned	Link
Deploy/Undeploy Java EE Applications	Deploying a Java EE Application	Deploy Java EE Applications to and from WebLogic versions 10.3, 10.3.x and later, including 12.1, 12.1.x, 12.1.2. and 12.2.1.x.	Deploying a Java EE Application
Deploy/Undeploy Java EE Applications	Undeploying a Java EE Application	Undeploy Java EE Applications to and from WebLogic versions 10.3, 10.3.x and later, including 12.1, 12.1.x, 12.1.2. and 12.2.1.x.	Undeploying a Java EE Application

Table 24-3 (Cont.) Deploying, Undeploying, or Redeploying Java EE Applications

Deployment Procedure	Use Case	Targets Provisioned	Link
Redeploy Java EE Applications	Redeploying a Java EE Application	Redeploy Java EE Applications to and from WebLogic versions 10.3, 10.3.x and later, including 12.1, 12.1.x, 12.1.2. and 12.2.1.x.	Redeploying a Java EE Application

Provisioning Coherence Nodes and Clusters

This table covers the use case to successfully deploy a Coherence node.

Table 24-4 Provisioning Coherence Nodes and Clusters

Deployment Procedure	Use Case	Targets Provisioned	Link
Coherence Node Provisioning	Deploying Coherence Nodes and Clusters	Oracle Coherence 3.7 and Oracle Coherence 12.1.2 Standalone Version.	Deployment Procedure

Provisioning SOA Artifacts

This table covers the use case to successfully deploy the various SOA artifacts:

Table 24-5 Provisioning SOA Artifacts

Deployment Procedure	Use Case	Targets Provisioned	Link
SOA Artifacts Provisioning	Provisioning SOA Artifacts from a Reference Installation	Oracle SOA Suite 11gR1 (11.1.1.2.0 to 11.1.1.7.0): • SOA Composites • Oracle WebLogic Server Policies • Assertion Templates • JPS Policy and Credential Stores • Human Workflow • Oracle B2B	Provisioning SOA Artifacts from a Reference Installation
SOA Artifacts Provisioning	Provisioning SOA Artifacts from a Gold Image	Oracle SOA Suite 11gR1 (11.1.1.2.0 to 11.1.1.7.0): • SOA Composites • Oracle WebLogic Server Policies • Assertion Templates • JPS Policy and Credential Stores • Human Workflow • Oracle B2B	Provisioning SOA Artifacts from Gold Image

Table 24-5 (Cont.) Provisioning SOA Artifacts

Deployment Procedure	Use Case	Targets Provisioned	Link
Deploy SOA Composites	Provisioning SOA Composites	Oracle SOA Suite 11gR1 (11.1.1.2.0 to 11.1.1.7.0) SOA Composites	Deploying SOA Composites

Provisioning Service Bus Resources

This table covers the use case to successfully deploy Service Bus resources:

Table 24-6 Provisioning Service Bus Resources

Deployment Procedure	Use Case	Targets Provisioned	Link
Service Bus Resource Provisioning	Provision Service Bus resources from Service Bus Domain	Service Bus 2.6.0 - 2.6.1, 3.0.0, 10.3.0.0 - 10.3.1.0, 11.1.1.3.0 - 11.1.1.7.0, 12.1.3.0.0	Provisioning Service Bus Resources from Service Bus Domain
Service Bus Resource Provisioning	Provision Service Bus resources from Software Library	Service Bus 2.6.0 - 2.6.1, 3.0.0, 10.3.0.0 - 10.3.1.0, 11.1.1.3.0 - 11.1.1.7.0, 12.1.3.0.0	Provisioning Service Bus Resources from Oracle Software Library

Provisioning Fusion Middleware Domain and Oracle Homes

Provisioning is a solution offered as a part of Lifecycle Management by Enterprise Manager Cloud Control. As a part of Middleware Provisioning solution, Cloud Control enables you to provision Oracle WebLogic Domain (WebLogic Domain), Service-Oriented Architecture (SOA), Service Bus, and Oracle WebCenter. In addition to provisioning a domain, you can extend an existing domain to include new Oracle Homes. For provisioning any of the middleware entities, you can create profile, save it, and then use the saved profile as the source for provisioning. Doing this, ensures that the future installations follow a standard, consistent configuration.

 Note:

The term Middleware Home is applicable only for WebLogic Server versions 10.3.x and 12.1.x. For WebLogic Server version 12.1.2.0.0 and higher, Middleware Home is referred to as Oracle Home.

To provision a new Middleware Domain or an Oracle Home, you must have configured JDK version 1.6 or higher. Oracle recommends using the latest and the most updated JDK version of the product. To do so, log in to support.oracle.com, then click **Certifications** tab. In the Certification Search section, enter **Oracle WebLogic Server** in the Product field. From the **Release** menu, select a valid WebLogic Server version and from the **Platform** menu, select a valid operating system, and click **Search**.

To clone an existing domain, you must ensure that the JDK version configured on the cloned destination host is equal to or higher than the JDK version available on the source host. Oracle recommends using the latest and the most updated JDK version of the product. To do so, log in to support.oracle.com, then click **Certifications** tab. In the Certification Search section, enter **Oracle WebLogic Server** in the Product field. From the **Release** menu, select a valid WebLogic Server version and from the **Platform** menu, select a valid operating system, and click **Search**.

This chapter explains how you can automate common provisioning operations for Middleware Homes and WebLogic Domains using Oracle Enterprise Manager Cloud Control. In particular, this chapter covers the following:

- [Getting Started with Fusion Middleware Provisioning](#)
- [Different Approaches to Launch the Provision Fusion Middleware Deployment Procedure](#)
- [High-Level Steps for Middleware Provisioning](#)
- [Prerequisites for Provisioning from the Middleware Provisioning Profiles](#)
- [Creating Middleware Provisioning Profiles](#)

- [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#)
- [Provisioning a Fusion Middleware Domain from an Existing Oracle Home](#)
- [Cloning from an Existing WebLogic Domain Based-Profile](#)

Getting Started with Fusion Middleware Provisioning

This section helps you get started by providing an overview of the steps involved in provisioning WebLogic Domain and Middleware Home using the Fusion Middleware Deployment procedure.

Table 25-1 Getting Started with Fusion Middleware Provisioning

Step	Description	Reference Links
Step 1	Different approaches to launch the Fusion Middleware Deployment Procedure. Understanding the various approaches to launch the Fusion Middleware Deployment Procedure, which is used to provision the Middleware entities	To learn about the approaches to launch the Deployment Procedure, see Different Approaches to Launch the Provision Fusion Middleware Deployment Procedure Export and Import of Domain Partitions
Step 2	Creating the Middleware Provisioning Profiles. This chapter covers three types of provisioning profiles. Select the profile that best matches your requirement	To learn about the various Provisioning Profiles, see Creating Middleware Provisioning Profiles
Step 3	Meeting Prerequisites to Provision a Middleware Profile Before you run the Fusion Middleware Deployment Procedure, there are a few prerequisites that you must meet.	To learn about the prerequisites for provisioning an Installation Media/Oracle Home profile, see Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile To learn about the Prerequisites for provisioning a WebLogic Domain profile, see Prerequisites for Provisioning the WebLogic Domain Profile
Step 4	Running the Fusion Middleware Deployment Procedure Run this deployment procedure to successfully provision a Weblogic Domain and/or an Oracle Home.	To learn about provisioning from an Installation Media Profile or an Oracle Home Profile, see Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile To learn about provisioning from a WebLogic Domain Profile, see Provisioning a Fusion Middleware Domain from an Existing Oracle Home To provision from an existing home, see Cloning from an Existing WebLogic Domain Based-Profile

Different Approaches to Launch the Provision Fusion Middleware Deployment Procedure

You can use any of the following approaches to launch the Middleware Provisioning deployment procedure:

- *(Recommended Option)* For all the out of the box deployment procedures, launch the Provision Fusion Middleware procedure from the Profiles table. For this, from

Enterprise menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**. On the Middleware Provisioning page, from the Profiles table select a profile, and click **Provision**.

- If you are provisioning a new domain from an existing Middleware home, or if you are using a customized deployment procedure, then you can directly run the Provision Fusion Middleware deployment procedure. For this, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**. On the Middleware Provisioning page, from the Deployment Procedures table, click **Provision Fusion Middleware**.
- To automate the process of provisioning using the command line, submit your procedure using the Enterprise Manager Command Line Interface (EMCLI) utility. The EMCLI enables you to access Enterprise Manager Cloud Control functionality from text-based consoles (shells and command windows) for a variety of operating systems. You can call Enterprise Manager functionality using custom scripts, such as SQL*Plus, OS shell, Perl, or Tcl, thus easily integrating Enterprise Manager functionality with your company's business process.

**Note:**

For more information about related EM CLI verbs, see EM CLI verbs in the *Oracle Enterprise Manager Command Line Interface*.

High-Level Steps for Middleware Provisioning

Typically, for provisioning, you will need to create a profile and use the profile as the source for creating new WebLogic Domains and Oracle Homes. Use of profiles ensures that the future installations follow a standard, consistent configuration. Oracle additionally supports fresh install, which means that you can create a profile from install media, and then provision this profile by running the **Provision Fusion Middleware** deployment procedures. For this, you do not need any Oracle home or WebLogic Domain to have already been manually installed/configured and discovered by Enterprise Manager Cloud Console.

Middleware Provisioning is a two-step process as follows:

- [Step1: Creating a Profile](#)
- [Step2: Running Provision Fusion Middleware Procedure to Provision the Profile](#)

Step1: Creating a Profile

Profiles are templates that you create and internally store in the Software Library. Once a profile is created, it can be launched numerous times to provision WebLogic Domain and/or Oracle Home with the same characteristics.

Before going ahead with the profile creation, you must identify the profile that suits your requirements:

- If you do not have an existing WebLogic Domain, or an Oracle Home, use the Installation Media profile. This profile allows you to create a domain and a home using the shiphome (installation media files) that you have downloaded from the

OTN or eDelivery, see [Creating a Provisioning Profile Based on an Installation Media](#).

- If you have an Oracle Home that needs to be patched to fix some bugs, use the Oracle Home profile. This profile allows you to apply the relevant patches to the Oracle home, and then create a profile out of the patched Oracle home which can be your source of truth. Use the patched profile to provision all the Oracle Homes in your data center. For further details, see [Creating a Provisioning Profile Based on an Oracle Home](#).
- If you have an existing Weblogic Domain that you want to clone, use the WebLogic Domain Profile. This profile allows you to create a copy of the source domain. For further details, see [Creating a Provisioning Profile Based on a WebLogic Domain](#).
- Oracle WebLogic Server supports WebLogic domain partition functionality. It consolidates applications into fewer domains with partition within domains and maintains secured isolation. For further details, see [Export and Import of Domain Partitions](#).
- Enterprise Manager Cloud Control provides an user friendly user interface that helps you to convert domains into partitions. For further details, see [Exporting WebLogic Domain](#).

Step2: Running Provision Fusion Middleware Procedure to Provision the Profile

A deployment procedure is a pre-defined sequence of steps that is meant to perform a set of operations on one or most hosts monitored by Cloud Control. You can provision a domain or an Oracle home using the profiles you have already created. This section describes the various methods to provision a profile. In particular, it covers the following:



Note:

Meet the Prerequisites before going ahead with the provisioning procedure. For a detailed list of prerequisites, see [Prerequisites for Provisioning from the Middleware Provisioning Profiles](#).

For information about the other approaches to launch the Provision Fusion Middleware procedure, see [Different Approaches to Launch the Provision Fusion Middleware Deployment Procedure](#).

In particular, it covers the following:

- [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
- [Provisioning a Fusion Middleware Domain from an Existing Oracle Home](#)
- [Cloning from an Existing WebLogic Domain Based-Profile](#).

Prerequisites for Provisioning from the Middleware Provisioning Profiles

This section describes all the prerequisites to be met before actually launching the Provision Fusion Middleware deployment procedure. In particular, it covers the following:

Note:

Before provisioning using a domain profile or an Oracle Home profile, you must download the following patches by signing in to My Oracle Support Web site at <https://support.oracle.com>:

- For SOA: 22912570
- For Service Bus: 23003815

If you are provisioning from a domain profile using an existing Middleware Home, then ensure that the Middleware Home is patched appropriately.

- [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#)
- [Prerequisites for Provisioning the WebLogic Domain Profile](#)

Note:

Meet the basic prerequisites mentioned in [Setting Up Your Infrastructure](#). This chapter describes the prerequisites based on the privileges and access you have been granted.

If you are performing a provisioning operation on a domain that was created in Development Mode, ensure that you disable the active configuration locks. To do so, perform the following steps:

- Log in to the Middleware Administrator Console.
- Click **Preferences**.
- Deselect **Automatically Acquire Lock and Activate Changes** option.
- Click **Save**.
- Ensure that the **Lock and Edit** button is enabled in the Change Center.

Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile

Meet the following prerequisites:

- Write permission on the Working Directory, which is a temporary directory used for staging and provisioning entities across Cloud Control.
If the working directory is not specified, then the Management Agent's working directory is used, and write permission on that is required.
- Following permissions on the Oracle Home directory:
 1. Write, if a new Oracle Home has to be created. For an Installation Media profile, creating a new Oracle Home is mandatory.
 2. Read, if you are using an existing Oracle Home. Note that this is applicable only for an Oracle Home profile, and not for an Installation Media profile.
- If you are using a shared storage, then mount the domain home and inventory directories on all the hosts beforehand. Typically, for a two-node setup, with two SOA Managed Servers running on two different hosts, for example host 1 and host 2; you can choose to create the shared storage on host 1. Effectively, this means that the Middleware Home location, the domain details, Inventory details, and all other information are mounted on host 1 for easy access from host 2.

 Note:

Support for provisioning using a Shared Storage is not available on a Window machine.

- All the hosts involved in the provisioning operation should be monitored as targets in Enterprise Manager.
- Server and Node Manager Ports should be free.
- If the domain uses a database or LDAP or Oracle HTTP Server, then ensure that the respective servers are monitored as targets in Enterprise Manager.
- Write permission on Domain /Application/Node Manager Directories.

Prerequisites for Provisioning the WebLogic Domain Profile

In addition to the requirements mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#), you need to meet the following for cloning using a WebLogic Domain profile:

- Before cloning an existing Fusion Middleware domain, you must have cloned the source database, so that the data in the schema is in sync with the source database. If you haven't already cloned your source database, you can do so using the Cloning Database feature available in Enterprise Manager Cloud Control, see [Cloning Oracle Databases and Pluggable Databases](#).

**Note:**

Before cloning the domain, run the following data scrubbing SQL scripts on the database.

For SOA

Create a Generic Component and upload `truncate_soa_oracle.sql` script to Software Library. For more information on creating generic components, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*. Note that the truncate script (`truncate_soa_oracle.sql`) is located in the following directory under SOA Installation: `/MW_HOME/SOA_ORACLE_HOME/rcu/integration/soainfra/sql/truncate`

For Service Bus

Create a Generic Component and upload `llr_table.sql` script to Software Library. To create this script, for each server present in the Service Bus domain, you need to add the following statement in the SQL script:

```
TRUNCATE table WL_LLRL_SERVER_NAME
```

For example, if the Service Bus domain has administrator server and two managed servers with name `OSB_SERVER1` and `OSB_SERVER2`, then the content of the sql script would look like:

```
TRUNCATE table WL_LLRL_ADMINSERVER
```

```
TRUNCATE table WL_LLRL_OSB_SERVER1
```

```
TRUNCATE table WL_LLRL_OSB_SERVER2
```

- If the source domain was wired with LDAP, then before cloning an existing Fusion Middleware domain, ensure that the data (users, roles and policies) has been migrated from the source LDAP to a new LDAP and the new LDAP has been discovered in Enterprise Manager as a target.

**Note:**

Middleware provisioning fails for JRF 12.1.2.0.0 enabled WebLogic domain, in case the source database is either down or not reachable. While unpacking the installable file, an attempt is made for connecting to the source database. In case the source database is down or unresponsive the provisioning task fails. The database must be up and running. This is not applicable for WebLogic 12.1.3.0.0 and onwards.

Using Custom Scripts Stored in the Software Library

This section describes how to use the custom scripts available in the Software Library to customize your deployment procedure. In particular, this section covers the following:

- [Using Custom Scripts with Input Parameters](#)
- [Using Custom Scripts Without Inputs Parameters](#)

Using Custom Scripts with Input Parameters

To use the custom scripts as a part of the Provision Fusion Middleware procedure, you must ensure that you have stored these scripts on the Software Library as a directive.

To store custom scripts with input parameters on the Software Library, follow these steps:

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library**.
2. Create a folder called **Directive**.
3. On the Software Library page, from the **Actions** menu, select **Create Entity**, then click **Directives**.
4. On the Describe page, provide a unique name for the parent folder. For example, **My Custom Script With Parameters**. Click **Next**.
5. On the Configure page, in the Command Line Arguments section, click **Add**.
6. In the Add Command Line Argument dialog box, enter the property name **INPUT_FILE**, then click **OK**. Verify that the Command Line contains the value: "\$ {INPUT_FILE}"
7. On the Select Files page, select **Upload Files**. In the Specify Destination section, select an upload location. In the Specify Source section, select the scripts to be uploaded. Ensure the directly executable file is in the Main File menu.

For example, `myscript.pl` is the main executable file, and `listMyServers.py` is another file that is uploaded. Following are example scripts that you can refer:

Contents of the Perl script: `myscript.pl`

```
#!/usr/local/bin/perl

print "*****\n";
print "*      This is a      *\n";
print "*    test script      *\n";
print "*****\n";

my $inputFile = $ARGV[0];
my %properties;
open (FILE, "<$inputFile") or die "can't open $inputFile for reading: $!";
print "Input properties:\n";
while (<FILE>)
{
    chomp;
    my ($key, $val) = split /=/;
    $properties{$key} = $val;
    print "\t$key=$val\n";
}
close FILE;

my $mwHome = $properties{MIDDLEWARE_HOME};
my $protocol = $properties{ADMIN_PROTOCOL};
my $host = $properties{ADMIN_SERVER_LISTEN_ADDRESS};
my $port = $properties{ADMIN_SERVER_LISTEN_PORT};
my $cmd = $mwHome."/wlserver_10.3/common/bin/wlst.sh listMyServers.py $protocol $host $port";

print "\nExecuting:\n\t$cmd\n";
```

```
print "\nOutput is:\n\n";

system($cmd);

exit 0;
```

Contents of the Python script: listMyServers.py

```
#!/usr/bin/python

protocol = sys.argv[1];
host = sys.argv[2];
port = sys.argv[3];
username = 'weblogic';
password = 'welcome1';

connectUrl = protocol + '://' + host + ':' + port;

connect(username, password, connectUrl);
cd('/Servers');
ls();
disconnect();
exit();
```

Click **Next**.

8. On the Review page, review the details, and click **Save and Upload**.

Using Custom Scripts Without Inputs Parameters

To use the custom scripts as a part of the Provision Fusion Middleware procedure, you must ensure that you have stored these scripts on the Software Library as a directive.

To store custom scripts without input parameters on the Software Library, follow these steps:

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library**.
2. Create a folder called **Directive**.
3. On the Software Library page, from the **Actions** menu, select **Create Entity**, then click **Directives**.
4. On the Describe page, provide a unique name for the parent folder. For example, **My Custom Script Without Parameters**. Click **Next**.
5. On the Configure page, in the Command Line Arguments section, click **Next**.
6. On the Select Files page, select **Upload Files**. In the Specify Destination section, select an upload location. In the Specify Source section, select the scripts to be uploaded. Ensure the directly executable file is in the Main File menu.

For example, myscript1.pl is the main executable file, and listMyServers.py is another file that is uploaded. Following are example scripts that you can refer:

Contents of the Perl script: myscript1.pl

```
#!/usr/local/bin/perl

print "*****\n";
print "*      This is a      *\n";
print "*      test script      *\n";
```

```

print "*****\n";

my $mwHome = "/scratch/bbanthia/soa/middleware";
my $protocol = "t3";
my $host = "slc01mpj.us.example.com";
my $port = "7001";
my $cmd = $mwHome."/wlserver_10.3/common/bin/wlst.sh listMyServers.py $proto-
col $host $port";

print "Executing:\n\t$cmd\n";
print "\nOutput is:\n\n";

system($cmd);

exit 0;

```

Contents of the Python script: listMyServers.py

```

#!/usr/bin/python

protocol = sys.argv[1];
host = sys.argv[2];
port = sys.argv[3];
username = 'weblogic';
password = 'welcome1';

connectUrl = protocol + '://' + host + ':' + port;

connect(username, password, connectUrl);
cd('/Servers');
ls();
disconnect();
exit();

```

Click **Next**.

7. On the Review page, review the details, and click **Save and Upload**.


Creating Middleware Provisioning Profiles

Profiles are like templates that you can create and store in Software Library. Once a profile is created, it can be launched numerous times to provision WebLogic Domain and/or Oracle Home. The advantage of using a profile is that you can ensure that future WebLogic installations follow a standard, consistent configuration.

- [Creating a Provisioning Profile Based on an Installation Media](#)
- [Creating a Provisioning Profile Based on an Oracle Home](#)
- [Creating a Provisioning Profile Based on a WebLogic Domain](#)

Creating a Provisioning Profile Based on an Installation Media

You can upload Installation Media from a remote host, create a profile based on the Installation Media selected, and provision this profile from the Middleware Provisioning page.



Before you begin creating the middleware provisioning profile, ensure that you meet the following prerequisites:

- Access the Oracle Fusion Middleware Download, Installation and Configuration Read-me file:

http://docs.oracle.com/cd/E23104_01/download_readme.htm

Download the required installation media file. The Fusion Middleware documentation, supported system configurations, system requirements, and specifications are available on this page.

The installation media files available on Oracle Software Delivery Cloud (via edelivery.oracle.com) is certified.

- Create one directory for each product like SOA, Service Bus, WebCenter, RCU and WLS, and ensure that you add the necessary files under the respective directory.
- Ensure that you have read permissions on all the files and sub-directories inside the domain home, applications home, and the Oracle home.

To create an Installation Media profile, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. On the Middleware Provisioning home page, in the Profiles section, click **Create** and select **From Installation Media**.
3. Click **Next**.
4. On the Create Fusion Middleware Provisioning Profile page, enter a unique name and description for your profile.

By default, all the profiles are centrally located in Software Library under Fusion Middleware Provisioning/Profiles directory.

5. In the Product Details section, from the **Product** menu, select the product for which the installation media is being uploaded. Depending on the option selected, the **Version** menu gets updated. The Enterprise Manager does not determine the version of the installation media. Select the right version of the installation media file that you have downloaded.

 Note:

When you select Oracle WebLogic Server as the option in the **Product** field, a checkbox **Oracle Fusion Middleware Infrastructure** appears. The user must ensure to check this box while uploading the Fusion Middleware Infrastructure installer. The user must not check this box for WebLogic server installer.

6. In the Files section, do the following:
 - a. Click the search icon to search for the host. In the Select Target dialog box, search and select the target where the Installation Media files reside, then click **Select**.
 - b. To access the files on a remote host, you need to provide the host credentials. To do so, click search, and in the Select Credential dialog box, enter the necessary credentials, and click **OK**. Click **Test** to validate these credentials against the selected target.
 - c. Based on the product selected, the Files table gets updated. One of the following options are possible:


If you select a product from the Product menu, then you can upload Oracle WebLogic Server, Oracle SOA, Service Bus, and Oracle RCU.

Before actually uploading the files, as a prerequisite, download the software Installation Media files as stated above.

To add the files, select the product name from the files table, and click **Select Folder**. Navigate to the directory where the files are present, and click **OK**. To remove files, select the product type, and click **Remove**.
7. In the Storage section, select the Software Library storage details. Ensure that you provide a valid storage type, and upload location to upload the Installation Media profile.
8. Click **Create** to submit the profile creation job.
9. After the job has successfully run, a new entry is available in the Profiles table. You can click the profile name to view the details.

Creating a Provisioning Profile Based on an Oracle Home

Use this page to create a new Oracle home profile, save it in Software Library, then use the saved profile as the source for provisioning new Oracle homes.



Before you begin creating the middleware provisioning profile, ensure that you meet the following prerequisites:

- Oracle Home should be a managed target that has been discovered in Enterprise Manager Cloud Control.
- Ensure that you have read permissions on all the files and sub-directories inside the Oracle home.
- Ensure that you have write permissions on Management Agent's working directory. Working directory is a temporary directory used for staging and provisioning entities across Cloud Control.
- The disk space required to create a provisioning profile is calculated as follows:
$$\text{Disk Space} = \text{Middleware Home Size} + \text{Space for Temporary Scripts}$$

To create a Middleware Home profile, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. On the Middleware Provisioning home page, in the Profiles section, click **Create** and select **From an Oracle Home**.
3. Click **Next**.
4. On the Create Fusion Middleware Provisioning Profile page, enter a unique name and description for your profile.

By default, all the profiles are centrally located in Software Library under Fusion Middleware Provisioning/Profiles directory.


5. In the Reference Target section, click the search icon. In the Select Target dialog box, select an Oracle Home, and click **Select**. The corresponding host details are populated.

You can also launch the Create Middleware Home Profile from the Middleware targets page. If you do so, the context of the target is maintained, and the fields such as **Type**, **Name**, and **Host** appear pre-populated.

6. Click the search icon to provide the credentials. In the Select Credentials dialog box, provide the necessary credentials for your target that is already set, and click **OK**. Click **Test** to validate the credentials against the selected target.
7. In the Storage section, select the Software Library storage details. Ensure that you provide a valid storage type, and upload location details to update the Oracle Home profile.
8. Click **Create Profile** to submit the profile creation job.
9. After the job has successfully run, a new entry is available in the Profiles table. You can click the profile name to view the details.

Creating a Provisioning Profile Based on a WebLogic Domain

Use this page to create a profile, save it in Software Library, then use the saved profile as a source for creating new WebLogic domains. This way, you can ensure that future WebLogic installations follow a standard and consistent configuration.



 Note:

- Profile creation and provisioning of domains with partitions is not supported.
- In case the domain is cloned the partition may not function as expected once provisioned.
- To clone, set up a new domain, and then import or export partitions. This is supported from release 13.2 only.

Before you begin creating the Middleware provisioning profile, ensure that you meet the following prerequisites:

- The Management Agent must be running on the Administration Server.
- Ensure that you have read permissions on all the files and sub-directories inside the domain home, applications home, and the Oracle home.
- Host credentials must be set for the source machine on which Administration Server is running.
- The WebLogic domain for which the profile is being created must be a monitored target in Cloud Control.
- Ensure to disable host name verification on all the hosts. This can be accomplished by following a simple procedure. In **SSL**, click **Advanced**, then set the **Host Name Verification** to *None*.
- The disk space required to create a profile is calculated as follows:

Disk Space = Middleware Home Size + WebLogic Domain Size + Space for Temporary Scripts (about one GB)

 Note:

While creating a provisioning profile for SOA or Oracle Service Bus domains, there is an option **Include metadata**. When you check this option it saves some information of the profile. It creates new schema and inserts the saved metadata. Hence database cloning is not required. This is supported from release 13.2 only.

To create a WebLogic Domain profile, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. On the Middleware Provisioning home page, in the Profiles section, click **Create** and select **From a WebLogic Domain**.
3. Click **Next**.
4. On the Create Fusion Middleware Provisioning Profile page, in the Profile Details section, enter a unique name and description for your profile. By default, all the profiles are centrally located in Software Library under the Fusion Middleware Provisioning/Profiles directory.
5. In the Reference Target section, search and select the WebLogic Domain target. Once you select the target, the Domain Home, the Host, and the Oracle home details for the host get populated.

You must ensure that a database instance is associated with the WebLogic Domain except in the case of a plain WebLogic Domain where it may not be necessary. To associate a database with the domain, create a database profile using the Create Database Provisioning Profile wizard.

Enter the Oracle Home credentials required to log into the host target selected.

Enter an existing named credential of weblogic administrator for the domain or create a new named credential by entering the details.

 Note:

While creating Middleware Domain Provisioning profile, you must select a WebLogic named credential with admin privileges. You must explicitly select the option WebLogic Administrator Credentials as the Credential Type. If this option is not selected then a named credential of the same admin user without administrator privileges will be created and profile creation will be failed later.

 Note:

You can also launch the WebLogic Domain Profile from the [Middleware target page](#). If this is done, then the context of the of the target is maintained, and the fields like **Type**, **Name**, **Host**, and **Oracle Home** details appear pre-populated.

6. Based on the target selection, you may create one of the following profiles:
 - a. **Domain-Based Profile with Oracle Home:** If you do not wish to upload the Oracle Home files, de-select the **Include Oracle Home** checkbox available in the Reference Target section. If you do so, then while provisioning from this domain profile, you will need to have a pre-deployed Oracle Home that already exists at the destination, and the content of the Oracle Home should match with the one expected by the domain profile.

- b. **Domain-Based Profile without Oracle Home:** If you retain the default selection of **Include Oracle Home**, then a WebLogic Domain with the domain configurations, Middleware Home, and binaries is created.

Click the search icon to provide credentials. In the Select Credentials dialog box, select the credentials that you have already set, and click **OK**. Click **Test** to validate the credentials against the selected target.

7. In the Software Library Upload Location section, select the Software Library storage details. Ensure that you provide a valid storage type, and upload location to update the profile.
8. Click **Create** to submit the profile creation job. After the job has successfully run, a new entry is available in the Profiles table. You can click the profile name to view the details.

Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile

To provision an Oracle Fusion Middleware domain or Oracle Home using an Installation Media-based profile, follow the steps below. The steps are same whether it involves SOA, OSB, WebCenter, or WebLogic.



Note:


Starting with Enterprise Manager for Oracle Fusion Middleware 12.1.0.7, you can provision SOA and Service Bus in a single operation. For this, you need to create an Installation Media profile with SOA and Service Bus installation files, and provision this profile. Until the previous release, you could only provision one domain (for example, SOA), and extend that domain to include the other product (for example, Service Bus). This release allows the flexibility of provisioning both SOA and Service Bus at the same time.

Middleware Provisioning supports RAC only with GridLink Data Sources.



Note:

To provide inputs and further customize the destination environment, click **Advance**. To understand the settings and configuration parameters that can be customized, see [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).



1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**.
2. On the Middleware Provisioning home page, from the Profiles table select an Installation Media profile, and click **Provision**.
3. On the Provision Fusion Middleware page, in the General section, the Installation Media profile appears pre-selected.

 Note:

If you are provisioning a WebCenter profile, you can choose to use Development Topology or Production Topology. If you do not specify, by default, the Production Topology is selected. For more information about the topologies, see [About WebCenter Topologies Supported in Enterprise Manager](#).

4. In the Hosts section, search and select the destination hosts on which the Middleware Home and WebLogic Domain needs to be cloned. Click **+Add** to add the target hosts, and provide the login credentials for them. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all**.
5. In the Middleware section, the Middleware Base and Java Home values appear pre-populated; you can customize these values if required. Provide credentials for the domain Administrator.
6. In the Database section, depending on the profile being provisioned the following options are possible:
 - **From an Installation media Profile:**

If you are provisioning from an Installation Media profile, you can create a new schema using a profile contains an RCU shiphome. To do so, select the **Create Schema** option. A default value for the Schema Prefix is populated; you may change this if required. Provide credentials for the database target and the new schema. If you deselect the default schema creation option, then an existing schema on the database target is used.
 - **From an Oracle Home Profile:**

If you are provisioning from an Oracle Home profile, you cannot create a new schema. You must, choose an existing schema and provide the schema password.
7. *(optional)* In the Identity and Security section, you must enter the OID target name and the OID credentials. These are mandatory fields for creating an LDAP Authenticator and/or for reassociating the Domain Credential Store.

 Note:

If you are provisioning a WebCenter profile, this section is displayed only for the Production Topology which is the default option.

In addition to this, you must provide the following sets of inputs in the OID section:

- **Configure LDAP Authenticator Inputs.** This section describes how to create the LDAP authenticator:
 - Authenticator Name: Enter a name for the OIDAuthenticator provider. For example: MyOIDAuthenticator.
 - User Base DN: Specify the DN under which your Users start. For example, cn=users, dc=us, dc=mycompany, dc=com

- Group Base DN: Specify the DN that points to your Groups node. For example: `cn=groups,dc=us,dc=mycompany,dc=com`

 Note:

As a prerequisite, you must have already provisioned the users and groups in the LDAP.

- **Configure Security Store Inputs:** In this section, provide the JPS Root Node information. The JPS root node is the target LDAP repository under which all data is migrated. The format is `cn=nodeName`.

 Note:

As a prerequisite, you must have already created the root node in the LDAP.

8. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (`http|https://host-name:port`)

For example: `http://host.example.com:80`

 Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (`http|https|tcp:// host-name:port`)

For example: `http://host.example.com:80`

 Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

9. *(For a WebCenter Production Domain only)* To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.

- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, mycrawladmin.
- The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is SEARCHSYS.
- The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, we-search.
- In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: `http://search_server_listenAddress:search_server_listenPort`. For example, `http://host.example.com:80`.

Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy `webcenter_portal_ses_admin.zip` to the host where SES Domain is present.

Unzip the `webcenter_portal_ses_admin.zip` to find the files: `facet.xml` and `searchAttrSortable.xml`.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -http://host.example.com:80/search/api/admin/Admin-
Service createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://http://host.example.com:80/search/api/admin/AdminSer-
vice -p welcome1 updateAll searchAttr -a overwrite -i search-
AttrSortable.xml
```

Take a snapshot of the WebCenter Content server manually.

10. Click **Next** to schedule the procedure. If you click **Submit**, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, you can login with designer role and create a template with lockdowns, and assign the desired privileges to other administrator/operators to run this template. Click **Cancel** to exit the procedure configuration.
11. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page, see [About Deployment Procedures](#).
12. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile

To customize the destination environment, click Advance option available on the Fusion Middleware Provisioning wizard.

Note:

The Advance option is enabled only after you select the destination hosts.

Follow these steps:

Note:

Oracle recommends that you allow the default option **Typical** to remain selected, and provide all the details. Following which, you can click **Advance** to customize your destination environment. This way, the default values for most of the parameters appear pre-populated, and you will need to enter only the remaining (delta) details. Also, note that if you switch from Advance mode to typical mode, you will lose all the changes that you have made so far.

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**.
2. On the Middleware Provisioning home page, from the Profiles table select an Installation Media profile, and click **Provision**.
3. On the Provision Fusion Middleware Configuration page, in the General section, the Installation Media profile appears pre-selected.

Note:

If you are provisioning a WebCenter profile, you can choose to use Development Topology or Production Topology. If you do not specify, by default, the Production Topology is selected, see [About WebCenter Topologies Supported in Enterprise Manager](#).

By default, the middleware home is selected, you cannot deselect this option. However, you are allowed to choose whether or not you want to create a new domain for your middleware home. If you select **Provision Domain**, a middleware home is created on the destination host, and a new domain is set up for this middleware home on the destination host.

The Enterprise Manager template is added when the **Oracle Enterprise Manager - JRF (database required)** option is selected. If the **Basic WebLogic Server Domain** option is selected, no additional template is added. If the profile is of version 12.2.1.0.0 or above, then you can select the **Oracle Enterprise Manager - Restricted JRF** option to add the 'Enterprise Manager - Restricted' template to the domain. The database details are not asked as they are not required. If you select the **Oracle Enterprise Manager - JRF (database required)** option to configure the domain, then you must enter the database details for the data source connections.

 Note:

If you are provisioning a SOA or a Service Bus profile, you can extend the domain to include the other product. For example, if you are provisioning a SOA domain, and you select the **Extend an existing Domain** option, then you will be allowed to extend your SOA domain to include Service Bus. To extend a domain, you will need to select the domain and provide the administrator credentials. The extend domain feature is different from Scaling a domain where you add additional managed servers to an existing domain. In extend domain, you are able to create a hybrid domain that includes two products. As of now, only a SOA or an Service Bus domain can be extended. Support for other products like WebCenter is not available.

If you have added more than one destination host, select **Use Shared Storage** option to use a shared location for these hosts. The central inventory attached to that is executed locally in all hosts sharing the Oracle Home.

To clone only the Middleware Home, deselect the **Provision Domain**.

4. In the Hosts section, search and select the destination hosts on which the Middleware Home and WebLogic Domain needs to be cloned. Click **+Add** to add the target hosts, and provide the login credentials for them. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all**.
5. In the Middleware section, the values for Middleware Base and Java Home fields appear pre-populated. You can change the location details, if required.
 - For **Middleware Home**, enter the full path to the directory in which the Middleware Home is to be created.
 - For **Java Home**, enter the absolute path to the JDK directory to be used on the destination Host.

 Note:

You must have already installed JDK at the same path on all the hosts.

6. In the Domain section, the configuration for the source domain is displayed by default. You can change the following attributes to customize the domain properties:
 - **Domain Name:** The name of the domain. The generated components for the domain are stored under the specified Domain directory. For example, if you enter *mydomain*, your domain files are stored (by default) in *MW_HOME\user_projects\domains\mydomain*.
 - **Administrator Username:** The default Administrator account for the domain. This account is used to boot and connect to the domain's Administration Server. The username must not contain commas, tabs, or any of these characters: `< > # | & ? ( ) { }`.
 - **Administrator Password:** The password for the Administrator account. The password must be at least eight characters, and must contain at least one numeric character or at least one of the following characters: `! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [ \ ] ^ _ ` { | } ~`
 - **Unique Domain Identifier:** The unique domain identifier is a prefix to the domain name.
 - **Deploy JVMD Agent:** Deploys the JVMD agent on the servers or clusters of a Weblogic Domain. Java Virtual Machine Diagnostics (JVMD) enables administrators to diagnose performance problems in Java applications. After the JVMD Agent is deployed on the targeted JVM (the one running a production WebLogic Server), it collects real-time data and transmits it to the JVM Diagnostics Engine.
 - **Available JVMD Endpoints:** This is the list of available JVMD endpoints. For a successful JVMD deployment, the JVMD engine must be active, if the JVMD engine is not active the agent is not deployed.
 - **Targets:** This is the list of targets on clusters or managed servers. If no targets are provided, the JVMD agent will be deployed by default to all the clusters and non-clustered servers of the domain. If a cluster is selected as target, then all the servers that are part of that cluster automatically get the agent deployed.
 - **Domain Directory:** The location in which your domain directory will be stored. By default, the same path as the source is used to create this directory, but it can be changed. For example: If the Middleware Home is located at `/user/mwh`, the domain directory is created as `/user/domains`. The domain location can be anywhere on your local drive or network. On Windows, you must include the drive letter in this path.
 - **Applications Directory:** The directory in which the applications will be deployed on the destination host. By default, the same path as the source is used to create this directory, but it can be changed. For example: If the Middleware Home is located at `/user/mwh`, the application directory is created as `/user/mwh/applications`.
 - **Domain Mode:** The domain can operate in one of the following modes:

Production: The domain is used for production. In this mode, the security configuration is relatively stringent, requiring a username and password to deploy applications.

Development: The domain is used for development. In this mode, the security configuration is relatively relaxed, allowing you to auto-deploy applications.

- **Server Startup Mode:** you can start the server in one of the following modes depending on your requirement:

Start all Servers and Node Managers: This is the default option. Typically, you will select this option if you have no changes to make, and if the procedure has run as expected.

Start only Administration Server: This option starts only Administrator Server and Node Manager. Typically, you will select this option if you want to add a custom step to invoke the WLST online script, and then start the servers.

Do not start any Server or Node Manager: This option does not start any server or Node Manager. Typically, you will select this option if you have to customize the domain before starting any server.

7. In the Clusters section, you can modify the name of the cluster, enter the cluster address that identifies the Managed Servers in the cluster. You can select either Unicast or Multicast as the messaging mode. If you select **Multicast** mode, enter the address and port number that will be dedicated for multicast communications on the cluster. Click **+Add** to add one or more clusters to the configuration.

A SOA or OSB domain can be configured with a owsm cluster. Click **+Add OWSM Cluster** to configure SOA or OSB domain with an OWSM cluster,

8. In the Machines section, enter configuration information for machines in the domain. A Machine is a logical representation of the system that hosts one or more WebLogic Server instances. The Administration Server and Node Manager use the Machine definition to start remote servers. Click **+Add** to add one or more machine configurations. Enter the following details:

- **Node Manager Home:** The directory in which the Node Manager is installed. By default, the Node Manager is installed under the parent directory of the Middleware Home directory, but this can be modified.

\

Note:

The Node Manager home must always be installed inside the Administration Server domain home.

- **Node Manager Credentials:** You can use the same credentials as Administrative Server credentials or deselect the option to provide separate Node Manager Credentials.
- **Machine Name:** Enter a unique name for your machine.
- **Host:** Select the host name from the menu.
- **Node Manager Listen Address:** Enter the listen address used by Node Manager to listen for connection requests. By default, the IP addresses defined for the local system and localhost are shown in the drop-down list. The default value is the same as specified in the source domain.

 Note:

If multiple machines are running on the same host, the Node Manager Home location must be different for each host.

- **Node Manager Listen Port:** Enter a valid value for the listen port used by Node Manager to listen for connection requests. The valid Node Manager listen port range is 1 to 65535. The default value is 5556. The port number must be available on the destination machine.
9. In the Server section, enter the configuration information for the Administration Server and one or managed servers.
- Select the **Configure Coherence** check box
 - **Coherence Port:** This field is enabled only when you select the Configure Coherence option. You can retain the port number that is populated by default or change it.
 - **Listen Port:** By default, this option is selected. The values for the listen ports are pre-populated. You may enter any value from 1 to 65535. The port number you enter here must be available on the destination machine.

 Note:

If a domain was registered on the host with a port number whose status is down, you need to select a different port or manually de-register the domain before launching the deployment procedure.

- **SSL Listen Port:** If you enable SSL Listen Port, enter the port number of the SSL Listen Port for secure requests. You must ensure that the port numbers you specify for the Listen Port and SSL Listen Port are available. 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.
- **Administration Port:** This port enables the domain to run in Administration Mode. For this, all the servers needs to be SSL enabled.
- **Host:** Select the host on which the Administration Server or Managed Server is to be installed.
- **Listen Address:** Enter the listen address to be used to connect to the Administration Server or the Managed Server instance.
- **Machine:** Select the machine configuration that is to be associated with the Administration Server or the Managed Server
- **Cluster:** Select the cluster to be associated with the Managed Server.
- **Server Start:** Click to enter the Server Startup Parameters. Usually, the Node Manager can start a server without requiring you to specify startup options, however, since you have customized your environment you must specify the startup options in the Server Startup Parameter dialog box.
- **Transaction Log Directory:** Transaction log stores information about committed transactions that are coordinated by the server and that may not have been completed. Oracle WebLogic Server uses this transaction log for recovery.

ery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the Managed Servers within a cluster, store the transaction log in a location accessible to the Managed Server and its backup servers.

- **Domain Directory for Managed Server:** The directory in which the Managed Servers are installed. By default, they are installed under the parent directory of the Middleware Home directory, but this can be modified.
10. In the JMS Servers section, click **+Add** to add new JMS persistent stores and JMS servers. The storage type can be one of the following:
 - A JMS file store is a disk-based file in which persistent messages can be saved. You can modify the JMS file stores configured in your domain.
 - A JDBC Store:
 - Data Source
 - Targets
 11. In the Database section, depending on the profile being provisioned the following options are possible:
 - **From an Installation Media Profile:**

Search and select the database target. A default value for the Schema Prefix is populated; you may change this if required.

 - The **Create Schema** option allows you to create a new schema for the RCU profile selected by default. Provide credentials for the database target and the new schema.
 - Select **Same password for all** option if you would like to use the same username and password for all the data sources on the selected database target

If you deselect the default schema creation option, then an existing schema on the database target is used.
 - **From an Oracle Home Profile:**

If you deselect the default schema creation option, then an existing schema on the database target is used.
 12. *(optional)* In the Identity and Security section, you must enter the OID target name and the OID credentials. These are mandatory fields for creating an LDAP Authenticator and/or for reassociating the Domain Credential Store.

 Note:

If you are provisioning a WebCenter profile, this section is displayed only for the Production Topology which is the default option.

In addition to this, you must provide the following sets of inputs in the OID section:

- **Configure LDAP Authenticator Inputs.** This section describes how to create the LDAP authenticator:
 - Authenticator Name: Enter a name for the OIDAuthenticator provider. For example: MyOIDAuthenticator.

- User Base DN: Specify the DN under which your Users start. For example, `cn=users,dc=us,dc=mycompany,dc=com`
- Group Base DN: Specify the DN that points to your Groups node. For example: `cn=groups,dc=us,dc=mycompany,dc=com`

 Note:

As a prerequisite, you must have already provisioned the users and groups in the LDAP.

- **Configure Security Store Inputs:** In this section, provide the JPS Root Node information. The JPS root node is the target LDAP repository under which all data is migrated. The format is `cn=nodeName`.

 Note:

As a prerequisite, you must have already created the root node in the LDAP.

13. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (`http|https://host-name:port`)

For example: `http://host.example.com:80`

 Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (`http|https|tcp:// host-name:port`)

For example: `tcp://wcp-prov-ucm.example.com:4444`

 Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

14. (For a WebCenter Production Domain only) To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.

- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, mycrawladmin.
- The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is SEARCHSYS.
- The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, we-search.
- In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: `http://search_server_listenAddress:search_server_listenPort`. For example, `http://host.example.com:80`.

Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy `webcenter_portal_ses_admin.zip` to the host where SES Domain is present.

Unzip the `webcenter_portal_ses_admin.zip` to find the files: `facet.xml` and `searchAttrSortable.xml`.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -c http://host.example.com:80/search/api/admin/Admin-
Service createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://host.example.com:80/search/api/admin/AdminService -p
welcome1 updateAll searchAttr -a overwrite -i searchAttrSorta-
ble.xml
```

Take a snapshot of the WebCenter Content server manually.

15. In the Custom Scripts section, you can select the scripts stored as directives in Software Library to customize the deployment procedure. Following options are possible, you may or may not choose to pass the scripts with parameters:
 - a. You can pass a script with input parameters. For more information, see [Using Custom Scripts with Input Parameters](#).

 Note:

If you pass the input parameter, ensure that you allow the default option of **Input File** to be selected. For example, in the Pre Script field, you can choose **My Custom Script With Parameters** script that you earlier created. For more information see Storing Custom Scripts With Input Parameters. [How?](#)

Following are the contents of a sample input properties (input.properties) file:

```
ADMIN_SERVER_LISTEN_ADDRESS=host.example.com
ADMIN_SERVER_LISTEN_PORT=7001
ADMIN_PROTOCOL=t3
MIDDLEWARE_HOME=/scratch/usr1/soa/middleware
```

- b. You can alternatively choose to pass a script without any input parameters. For more information, see [Using Custom Scripts Without Inputs Parameters](#).

Note: If you do not want to pass an input parameter, you should deselect the **Input File** option. For example, in the Pre Script field, you can choose **My Custom Script Without Parameters** script that you earlier created. For more information, see Storing Custom Scripts Without Input Parameters. [How?](#)

You can pass the following scripts to customize your procedure:

16. Click **Next** to schedule the procedure. If you click **Submit**, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, you can login with designer role and create a template with lockdowns, and assign the desired privileges to other administrator/operators to run this template. Click **Cancel** to exit the procedure configuration.
17. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page. For more information about this, see [About Deployment Procedures](#).
18. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Provisioning a Fusion Middleware Domain from an Existing Oracle Home

To provision a fresh domain from an existing Oracle home, follow these steps:



Note:

To provide inputs and further customize the destination environment, click **Advance**. To understand the settings and configuration parameters that can be customized, see [Customizing the Destination Environment from an Existing Oracle Home](#).

1. On Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**
2. On the Middleware Provisioning home page, from Deployment Procedure section select **Provision Fusion Middleware**, then click **Launch**.
3. On the Middleware Provisioning page, in the General section, select **Oracle Home**. Click the search icon, the Select target dialog box is displayed. Search and select the WebLogic Server home that you want to reuse.
4. In the Hosts section, click **+Add** to search and select the destination hosts where the cloned Middleware Home will reside. For the cloning operation, you will need to provide the login credentials for the destination hosts. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all** option.
5. In the Middleware section, provide the domain Administrator credentials.
6. In the Database section, select a Database Target, choose an existing schema and provide the schema password.
7. *(optional)* In the Identity and Security section, you must enter the OID target name and the OID credentials. These are mandatory fields for creating an LDAP Authenticator and/or for reassociating the Domain Credential Store.



Note:

If you are provisioning a WebCenter profile, this section is displayed only for the Production Topology which is the default option.

In addition to this, you must provide the following sets of inputs in the OID section:

- **Configure LDAP Authenticator Inputs.** This section describes how to create the LDAP authenticator:
 - Authenticator Name: Enter a name for the OIDAuthenticator provider. For example: MyOIDAuthenticator.
 - User Base DN: Specify the DN under which your Users start. For example, cn=users, dc=us, dc=mycompany, dc=com
 - Group Base DN: Specify the DN that points to your Groups node. For example: cn=groups, dc=us, dc=mycompany, dc=com

 Note:

As a prerequisite, you must have already provisioned the users and groups in the LDAP.

- **Configure Security Store Inputs:** In this section, provide the JPS Root Node information. The JPS root node is the target LDAP repository under which all data is migrated. The format is cn=nodeName.

 Note:

As a prerequisite, you must have already created the root node in the LDAP.

8. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (http|https://host-name:port)

For example: http://wcp-prov.example.com:80

 Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (http|https|tcp:// host-name:port)

For example: tcp://wcp-prov-ucm.example.com:4444

 Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

9. *(For a WebCenter Production Domain only)* To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.
- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- a. Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, mycrawladmin.
- b. The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is SEARCHSYS.
- c. The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, we-search.
- d. In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: `http://search_server_listenAddress:search_server_listenPort`. For example, `http://slc01rsk.us.example.com:5720`.

 Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy `webcenter_portal_ses_admin.zip` to the host where SES Domain is present.

Unzip the `webcenter_portal_ses_admin.zip` to find the files: `facet.xml` and `searchAttrSortable.xml`.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -c http://slc01rsk.us.example.com:5720/search/api/
admin/AdminService createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://slc01rsk.us.example.com:5720/search/api/admin/AdminService -p
welcome1 updateAll searchAttr -a overwrite -i searchAttrSortable.xml
```

Take a snapshot of the WebCenter Content server manually.

10. Click **Next** to schedule the procedure. If you click **Submit**, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, if you want to create a template with lockdowns, and allow other users with operator privilege to run the template multiple times with minor modifications.

Click **Cancel** to exit the procedure configuration.

11. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page. For more information about this, see [About Deployment Procedures](#).
12. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Customizing the Destination Environment from an Existing Oracle Home

To customize the destination environment, click **Advance** available on the Fusion Middleware Provisioning wizard. Note that the **Advance** option is enabled only after you select the destination hosts.

Follow these steps:

Note:

Oracle recommends that you allow the default option **Typical** to remain selected, and provide all the details. Following which, you can click **Advance** to customize your destination environment. This way, the default values for most of the parameters appear pre-populated, and you will need to enter only the remaining (delta) details. Also, note that if you switch from **Advance** mode to **typical** mode, you will lose all the changes that you have made so far.

1. On Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**.
2. On the Middleware Provisioning home page, from Deployment Procedure section, select **Provision Fusion Middleware**, then click **Launch**.
3. On the Provision Fusion Middleware Configuration page, in the General section, the Oracle Home appears pre-selected.

By default, the **Provision Domain** option is selected. Basically, a fresh domain is provisioned with an existing Oracle home.

To clone only the Middleware Home, deselect the **Provision Domain**.

4. In the Hosts section, search and select the destination hosts on which the Middleware Home and WebLogic Domain need to be cloned. Click **+Add** to add the target hosts, and provide the login credentials for them. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all**.
5. In the Domain section, the configuration for the domain is displayed by default. You can change the following attributes to customize the domain properties:
 - **Domain Name:** The name of the domain. The generated components for the domain are stored under the specified Domain directory. For example, if you enter *mydomain*, your domain files are stored (by default) in *MW_HOME\user_projects\domains\mydomain*. Ensure that you provide a unique domain name.

- **Administrator Username:** The default Administrator account for the domain. This account is used to boot and connect to the domain's Administration Server. The username must not contain commas, tabs, or any of these characters: `< > # | & ? ( ) { }`.
 - **Administrator Password:** The password for the Administrator account. The password must be at least eight characters, and must contain at least one numeric character or at least one of the following characters: `! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [ \ ] ^ _ ` { | } ~`
 - **Unique Domain Identifier:** The unique domain identifier is a prefix to the domain name.
 - **Domain Directory:** The location in which your domain directory will be stored. By default, the same path as the source is used to create this directory, but it can be changed. For example: If the Middleware Home is located at `/user/mwh`, the domain directory is created as `/user/domains`. The domain location can be anywhere on your local drive or network. On Windows, you must include the drive letter in this path.
 - **Applications Directory:** The directory in which the applications will be deployed on the destination host. By default, this directory is created under the parent directory of the Middleware Home. For example: If the Middleware Home is located at `/user/mwh`, the application directory is created as `/user/mwh/applications`.
 - **Domain Mode:** The domain can operate in one of the following modes:
 - Production:** The domain is used for production. In this mode, the security configuration is relatively stringent, requiring a username and password to deploy applications.
 - Development:** The domain is used for development. In this mode, the security configuration is relatively relaxed, allowing you to auto-deploy applications.
 - **Server Startup Mode:** you can start the server in one of the following modes depending on your requirement:
 - Start all Servers and Node Managers:** This is the default option. Typically, you will select this option if you have no changes to make, and if the procedure has run as expected.
 - Start only Administration Server:** This option starts only Administrator Server and Node Manager. Typically, you will select this option if you want to add a custom step to invoke the WLST online script, and then start the servers.
 - Do not start any Server or Node Manager:** This option does not start any server or Node Manager. Typically, you will select this option if you have to customize the domain before starting any server.
6. In the Clusters section, you can modify the name of the cluster, enter the cluster address that identifies the Managed Servers in the cluster. You can select either Unicast or Multicast as the messaging mode. If you select **Multicast** mode, enter the address and port number that will be dedicated for multicast communications on the cluster. Click **+Add** to add one or more clusters to the configuration.
 7. In the Machines section, enter configuration information for machines in the domain. A Machine is a logical representation of the system that hosts one or more WebLogic Server instances. The Administration Server and Node Manager use the Machine definition to start remote servers. Click **+Add** to add one or more machine configurations. Enter the following details:

- **Node Manager Home:** The directory in which the Node Manager is installed. By default, the Node Manager is installed under the parent directory of the Middleware Home directory, but this can be modified.

Note that the Node Manager home must always be installed inside the Administration Server domain home.
 - **Node Manager Credentials:** You can use the same credentials as Administrative Server credentials or deselect the option to provide separate Node Manager Credentials.
 - **Machine Name:** Enter a unique name for your machine.
 - **Host:** Select the host name from the menu.
 - **Node Manager Listen Address:** Enter the listen address used by Node Manager to listen for connection requests. By default, the IP addresses defined for the local system and localhost are shown in the drop-down list. The default value is the same as specified in the source domain. Note that if multiple machines are running on the same host, the Node Manager Home location must be different for each host.
 - **Node Manager Listen Port:** Enter a valid value for the listen port used by Node Manager to listen for connection requests. The valid Node Manager listen port range is 1 to 65535. The default value is 5556. The port number must be available on the destination machine.
8. In the Server section, enter the configuration information for the Administration Server and one or managed servers.
- Select the **Configure Coherence** check box
 - **Coherence Port:** This field is enabled only when you select the Configure Coherence option. You can retain the port number that is populated by default or change it.
 - **Listen Port:** By default, this option is selected. The values for the listen ports are pre-populated. You may enter any value from 1 to 65535. The port number you enter here must be available on the destination machine.

Note If a domain was registered on the host with a port number whose status is down, you need to select a different port or manually de-register the domain before launching the deployment procedure.
 - **SSL Listen Port:** If you enable SSL Listen Port, enter the port number of the SSL Listen Port for secure requests. You must ensure that the port numbers you specify for the Listen Port and SSL Listen Port are available. 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.
 - **Administration Port:** This port enables the domain to run in Administration Mode. For this, all the servers needs to be SSL enabled.
 - **Host:** Select the host on which the Administration Server or Managed Server is to be installed.
 - **Listen Address:** Enter the listen address to be used to connect to the Administration Server or the Managed Server instance.
 - **Machine:** Select the machine configuration that is to be associated with the Administration Server or the Managed Server
 - **Cluster:** Select the cluster to be associated with the Managed Server.

- **Server Start:** Click to enter the Server Startup Parameters. Usually, the Node Manager can start a server without requiring you to specify startup options, however, since you have customized your environment you must specify the startup options in the Server Startup Parameter dialog box.
 - **Transaction Log Directory:** Transaction log stores information about committed transactions that are coordinated by the server and that may not have been completed. Oracle WebLogic Server uses this transaction log for recovery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the Managed Servers within a cluster, store the transaction log in a location accessible to the Managed Server and its backup servers.
 - **Domain Directory for Managed Server:** The directory in which the Managed Servers are installed. By default, they are installed under the parent directory of the Middleware Home directory, but this can be modified.
9. In the JMS Servers section, click **+Add** to add new JMS persistent stores and JMS servers. The storage type can be one of the following:
 - A JMS file store is a disk-based file in which persistent messages can be saved. You can modify the JMS file stores configured in your domain.
 - A JDBC Store:
 - Data Source
 - Targets
 10. In the Database section, choose an existing schema and provide the schema password. Select **Same password for all** option if you would like to use the same username and password for all the data sources on the selected database target.
 11. *(optional)* In the Identity and Security section, you must enter the OID target name and the OID credentials. These are mandatory fields for creating an LDAP Authenticator and/or for re-associating the Domain Credential Store.

 Note:

If you are provisioning a WebCenter profile, this section is displayed only for the Production Topology which is the default option.

In addition to this, you must provide the following sets of inputs in the OID section:

- **Configure LDAP Authenticator Inputs.** This section describes how to create the LDAP authenticator:
 - Authenticator Name: Enter a name for the OIDAuthenticator provider. For example: MyOIDAuthenticator.
 - User Base DN: Specify the DN under which your Users start. For example, cn=users, dc=us, dc=mycompany, dc=com
 - Group Base DN: Specify the DN that points to your Groups node. For example: cn=groups, dc=us, dc=mycompany, dc=com

 Note:

As a prerequisite, you must have already provisioned the users and groups in the LDAP.

- **Configure Security Store Inputs:** In this section, provide the JPS Root Node information. The JPS root node is the target LDAP repository under which all data is migrated. The format is `cn=nodeName`.

 Note:

As a prerequisite, you must have already created the root node in the LDAP.

12. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (`http|https://host-name:port`)

For example: `http://host.example.com:80`

 Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (`http|https|tcp:// host-name:port`)

For example: `http://host.example.com:80`

 Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

13. (For a WebCenter Production Domain only) To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.
- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- a. Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, mycrawladmin.
- b. The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is SEARCHSYS.
- c. The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, we-search.
- d. In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: `http://search_server_listenAddress:search_server_listenPort`.

For example, `http://host.example.com:80`.

 Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy `webcenter_portal_ses_admin.zip` to the host where SES Domain is present.

Unzip the `webcenter_portal_ses_admin.zip` to find the files: `facet.xml` and `searchAttrSortable.xml`.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -c http://host.example.com:80/search/api/admin/Admin-
Service createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://host.example.com:80/search/api/admin/AdminService -p
welcome1 updateAll searchAttr -a overwrite -i searchAttrSorta-
ble.xml
```

Take a snapshot of the WebCenter Content server manually.

14. In the Custom Scripts section, you can select the scripts stored as directives in Software Library to customize the deployment procedure. Following options are possible, you may or may not choose to pass the scripts with parameters:

- a. You can pass a script with input parameters. For more information, see [Using Custom Scripts with Input Parameters](#).

Note: If you pass the input parameter, ensure that you allow the default option of **Input File** to be selected. For example, in the Pre Script field, you can choose **My Custom Script With Parameters** script that you earlier created. For more information see [Storing Custom Scripts With Input Parameters](#).

Following are the contents of a sample input properties (input.properties) file:

```
ADMIN_SERVER_LISTEN_ADDRESS=host.example.com
ADMIN_SERVER_LISTEN_PORT=7001
ADMIN_PROTOCOL=t3
MIDDLEWARE_HOME=/scratch/usr1/soa/middleware
```

- b. You can alternatively choose to pass a script without any input parameters.

Note: If you do not want to pass an input parameter, you should deselect the **Input File** option. For example, in the Pre Script field, you can choose **My Custom Script Without Parameters** script that you earlier created. For more information, see [Using Custom Scripts Without Inputs Parameters](#).

You can pass the following scripts to customize your procedure:

- **Pre Script:** This script runs soon after the prerequisite checks and before the Oracle Home or WebLogic Domain is deployed.
 - **Post Administration Server Start Script:** This script runs after the Administration Server has been started.
 - **Post Script:** This script runs after all the Managed Servers have started.
15. Click **Next** to schedule the procedure. If you click **Submit**, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, you can login with designer role and create a template with lockdowns, and assign the desired privileges to other administrator/operators to run this template. Click **Cancel** to exit the procedure configuration.
16. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page. For more information about this, see [About Deployment Procedures](#).
17. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Cloning from an Existing WebLogic Domain Based-Profile

To clone an existing domain, follow these steps:


Note:

To provide inputs and further customize the destination environment, click **Advance**. To understand the settings and configuration parameters that can be customized, see [Customizing the Destination Environment from an Existing WeLogic Domain Based-Profile](#).



Note:

For more information on cloning a database in WebLogic domain, see [Use Case 6 - Cloning a database in WebLogic Domain](#).





Note:

Middleware Provisioning supports RAC only with GridLink Data Sources.

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**.
2. On the Middleware Provisioning home page, from the Profiles table select a WebLogic Domain profile, then click **Provision**.
3. On the Provision Fusion Middleware page, in the General section, the WebLogic Domain profile appears pre-selected.
4. In the Hosts section, search and select the destination hosts where the cloned WebLogic Domain will reside. Click **+Add** to add the target hosts, and provide the login credentials for them. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all**.

Note: The domain configuration for the destination host and the source host must be exactly the same for cloning operation to be successful. For example, if the source domain has servers on two different hosts, then you will need to select two different destinations hosts. You will be prompted to do so before you proceed.

5. In the Middleware section the value for Java Home appears pre-populated. Provide the domain Administrator credentials.



Note:

The vendor and version of the JAVA in use at source and destination should preferably be identical or at least be of the same major version.

For more information, see Movement Scripts and Move Plans of the *Fusion Middleware Administering Oracle Fusion Middleware Guide*.

6. In the Database section, select the cloned database target, and provide the schema password.



Note:

This section appears only if the source database had data sources.

7. *(optional)* The Identity and Security section, enter the OID target name, and the credential. As a prerequisite, you must have cloned the source OID to a destination environment.
8. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (http|https://host-name:port)

For example: `http://wcp-prov.example.com:80`



Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (`http|https|tcp:// host-name:port`)

For example: `http://host.example.com:80`



Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

9. *(For a WebCenter Production Domain only)* To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.
- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- a. Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, `mycrawladmin`.
- b. The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is `SEARCHSYS`.
- c. The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, `we-search`.
- d. In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: `http://search_server_listenAddress:search_server_listenPort`. For example, `http://slc01rsk.us.example.com:5720`.

 Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy webcenter_portal_ses_admin.zip to the host where SES Domain is present.

Unzip the webcenter_portal_ses_admin.zip to find the files: facet.xml and searchAttrSortable.xml.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -c http://slc01rsk.us.example.com:5720/search/api/
admin/AdminService createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://slc01rsk.us.example.com:5720/search/api/admin/AdminSer-
vice -p welcome1 updateAll searchAttr -a overwrite -i search-
AttrSortable.xml
```

Take a snapshot of the WebCenter Content server manually.

10. Click **Next** to schedule the procedure. If you click Submit, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, if you want to create a template with lockdowns, and allow other users with operator privilege to run the template multiple times with minor modifications.

Click **Cancel** to exit the procedure configuration.

11. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page. For more information about this, see [About Deployment Procedures](#).
12. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Customizing the Destination Environment from an Existing WeLogic Domain Based-Profile

To customize the destination environment, click Advance available on the Fusion Middleware Provisioning wizard. Note that the Advance option is enabled only after you select the destination hosts.

Follow these steps:

 Note:

Oracle recommends that you allow the default option **Typical** to remain selected, and provide all the details. Following which, you can click **Advance** to customize your destination environment. This way, the default values for most of the parameters appear pre-populated, and you will need to enter only the remaining (delta) details. Also, note that if you switch from Advance mode to typical mode, you will lose all the changes that you have made so far.

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, and then click **Middleware Provisioning**.
2. On the Middleware Provisioning home page, from Deployment Procedure section, select **Provision Fusion Middleware**, then click **Launch**.
3. On the Middleware Provisioning page, in the General section, select the WebLogic domain profile that contains the source domain details that is to be cloned.

By default, you will provision the WebLogic domain along with Oracle home. To provision only the domain (*Bitless profile*), you must deselect **Provision Middleware Home** and select **Provision Domain**. Alternatively, to provision only the middleware home retain Provision Middleware Home option, and deselect Provision Domain.

Use Shared Storage is particularly useful when you have multiple destination hosts. This option allows you to use a mounted location that is accessible by all the hosts.

4. In the Hosts section, search and select the destination hosts where the cloned WebLogic Domain will reside. Click **+Add** to add the target hosts, and provide the login credentials for them. If you have selected multiple hosts, and the login credentials for all of them are the same, then you can select **Same Credentials for all**.

Note: The domain configuration for the destination host and the source host must be exactly the same for cloning operation to be successful. For example, if the source domain has servers on 2 different hosts, then you will need to select two different destinations hosts. You will be prompted to do so before you proceed.

5. In the Middleware section, by default, the values for Middleware Home and Java Home are pre-populated. You can change the location details if required.
 - For **Middleware Home**, enter the full path to the directory in which the Middleware Home is to be created.
 - For **Java Home**, enter the absolute path to the JDK directory to be used on the destination Host. You need to specify this path if a similar configuration is detected on the source machine.

 Note:

The vendor and version of the JAVA in use at source and destination should preferably be identical or at least be of the same major version.

For more information, see Movement Scripts and Move Plans of the *Fusion Middleware Administering Oracle Fusion Middleware Guide*.

6. In the Domain section, the configuration for the source domain is displayed by default. You can change the following attributes to customize the domain properties:
 - a. **Domain Name:** The name of the domain. The generated components for the domain are stored under the specified Domain directory. For example, if you enter *mydomain*, your domain files are stored (by default) in *MW_HOME\user_projects\domains\mydomain*.
 - b. **Administrator Username:** The default Administrator account for the domain. This account is used to boot and connect to the domain's Administration Server. The username must not contain commas, tabs, or any of these characters: `< > # | & ? ( ) { }`.
 - c. **Administrator Password:** The password for the Administrator account. The password must be at least eight characters, and must contain at least one numeric character or at least one of the following characters: `! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [ \ ] ^ _ ` { | } ~`
 - d. **Unique Domain Identifier:** The unique domain identifier is a prefix to the domain name.
 - e. **Deploy JVMD Agent:** Deploys the JVMD agent on the servers or clusters of a Weblogic Domain. Java Virtual Machine Diagnostics (JVMD) enables administrators to diagnose performance problems in Java applications. After the JVMD Agent is deployed on the targeted JVM (the one running a production WebLogic Server), it collects real-time data and transmits it to the JVM Diagnostics Engine.
 - f. **Available JVMD Endpoints:** This is the list of available JVMD endpoints. For a successful JVMD deployment, the JVMD engine must be active, if the JVMD engine is not active the agent is not deployed.
 - g. **Targets:** This is the list of targets on clusters or managed servers. If no targets are provided, the JVMD agent will be deployed by default to all the clusters and non-clustered servers of the domain. If a cluster is selected as target, then all the servers that are part of that cluster automatically get the agent deployed.
 - h. **Domain Directory:** The location in which your domain directory will be stored. By default, the same path as the source is used to create this directory, but it can be changed. For example: If the Middleware Home is located at */user/mwh*, the domain directory is created as */user/mwh/domains*. The domain location can be anywhere on your local drive or network. On Windows, you must include the drive letter in this path.
 - i. **Applications Directory:** The directory in which the applications will be deployed on the destination host. By default, the same path as the source is used to create this directory, but it can be changed. For example: If the Mid-

dleware Home is located at /user/mwh, the application directory is created as /user/mwh/applications.

- j. **Server Startup Mode:** you can start the server in one of the following modes depending on your requirement:

Start all Servers and Node Managers: This is the default option. Typically, you will select this option if you have no changes to make, and if the procedure has run as expected.

Start only Administration Server: This option starts only Administrator Server and Node Manager. Typically, you will select this option if you want to add a custom step to invoke the WLST online script, and then start the servers.

Do not start any Server or Node Manager: This option does not start any server or Node Manager. Typically, you will select this option if you have to customize the domain before starting any server.

7. In the Clusters section, all the clusters available in the source domain are provisioned on the destination host.
8. In the Machines section, the configuration information in the source domain are pre-populated. A Machine is a logical representation of the system that hosts one or more WebLogic Server instances. The Administration Server and Node Manager use the Machine definition to start remote servers. You can not customize any of the values here.
9. In the Servers section, all the configuration information for the Administration Server and managed servers are picked up from the source domain. You can customize the following:
 - a. Select the Configure Coherence check box
 - b. Coherence Port: the port value is pre-populated that can not be customized.
 - c. **Custom Identity and Custom Trust:** Use this option to specify custom certificates when configuring a domain in SSL mode.
10. In the JMS Servers section, all JMS servers configured for the source domain are cloned on the destination hosts. You can not customize any value here.
11. In the Database section, you can change the schema username and password for your data sources.



Note:

This section appears only if the source database had data sources.

12. (optional) The Identity and Security section, enter the OID target name, and the credential. As a prerequisite, you must have cloned the source OID to a destination environment.
13. In the WebTier section, click **+Add** to search and select an Oracle HTTP Server (OHS) target. In the **Credential** field provide the credentials to access the target. If you have more than one Oracle HTTP Server, then you need to provide an **External Load Balancer URL**.

The format of the URL for an External Load Balancer: (http|https://host-name:port)

For example: http://wcp-prov.example.com:80

 Note:

If you are provisioning a WebCenter target, you will additionally need to provide an **Internal Load Balancer URL**.

The format of the URL for an Internal Load Balancer: (http|https|tcp:// host-name:port)

For example: tcp://wcp-prov-ucm.example.com:4444

 Note:

If you are provisioning a plain WebLogic Domain, this section will not be displayed.

14. (For a WebCenter Production Domain only) To configure an SES Domain with a WebCenter Production Domain, follow these steps:

Prerequisites:

- The SES Domain must be up and running.
- The SES Domain must be configured with the same OID as the WebCenter Domain.
- You must have created a Crawl Administrator Username in the OID.

To configure SES for your WebCenter Domain, enter the following details:

- a. Provide the credentials for Crawl admin user in OID. You must enter the same Crawl Administrator Username and Crawl Administrator Password that was created as a part of the prerequisite step. The Crawl Administration users in Spaces, and in the Identity Management System, are required to crawl certain Space objects, such as lists, pages, spaces, and people connections profiles. For example, mycrawladmin.
- b. The Search Administrator Username and Search Administrator Password are the credentials that were used for creating the SES Domain. For example, for the Oracle SES 11.2.2.2 release, the Search Administrator user is SEARCHSYS.
- c. The Search User Username and Search User Password are the credentials of the Oracle SES federation trusted entity. These get created while installing the SES. Each Oracle SES instance must have a trusted entity for allowing WebCenter Portal end users to be securely propagated at search time. A trusted entity allows the WebCenter Portal application to authenticate itself to Oracle SES and assert its users when making queries on Oracle SES. This trusted entity can be any user that either exists on the Identity Management Server behind Oracle SES or is created internally in Oracle SES. For example, we-search.
- d. In the SES Search URL field, enter the URL of Search Administration Tool. The format of this URL should be: http://search_server_listenAddress:search_server_listenPort. For example, http://slc01rsk.us.example.com:5720.

 Note:

After configuring the WebCenter Domain with SES, if you encounter an issue with searching the portal, you must perform in the following steps manually:

From the WebCenter Portal Oracle Home, copy webcenter_portal_ses_admin.zip to the host where SES Domain is present.

Unzip the webcenter_portal_ses_admin.zip to find the files: facet.xml and searchAttrSortable.xml.

Run these XML files as follows, for example:

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -p
welcome1 -c http://slc01rsk.us.example.com:5720/search/api/
admin/AdminService createAll facetTree -i facet.xml
```

```
/scratch/SES/oracle/middleware/Oracle_SES1/bin/searchadmin -c
http://slc01rsk.us.example.com:5720/search/api/admin/AdminSer-
vice -p welcome1 updateAll searchAttr -a overwrite -i search-
AttrSortable.xml
```

Take a snapshot of the WebCenter Content server manually.

15. In the Custom Scripts section, you can select the scripts stored as directives in Software Library to customize the deployment procedure. Following options are possible, you may or may not choose to pass the scripts with parameters:

- a. You can pass a script with input parameters. For more information, see [Using Custom Scripts with Input Parameters](#).

Note: If you pass the input parameter, ensure that you allow the default option of **Input File** to be selected. For example, in the Pre Script field, you can choose **My Custom Script With Parameters** script that you earlier created. For more information see [Storing Custom Scripts With Input Parameters. How?](#)

Following are the contents of a sample input properties (input.properties) file:

```
ADMIN_SERVER_LISTEN_ADDRESS=slc01.example.com
ADMIN_SERVER_LISTEN_PORT=7001
ADMIN_PROTOCOL=t3
MIDDLEWARE_HOME=/scratch/usr1/soa/middleware
```

- b. You can alternatively choose to pass a script without any input parameters. For more information, see [Using Custom Scripts Without Inputs Parameters](#).

Note: If you do not want to pass an input parameter, you should deselect the **Input File** option. For example, in the Pre Script field, you can choose **My Custom Script Without Parameters** script that you earlier created. For more information, see [Storing Custom Scripts Without Input Parameters. How?](#)

You can pass the following scripts to customize your procedure:

- **Pre Script:** This script runs soon after the prerequisite checks and before the Oracle Home or WebLogic Domain is deployed.
- **Post Administration Server Start Script:** This script runs after the Administration Server has been started.

- **Post Script:** This script runs after all the Managed Servers have started.
16. Click **Next** to schedule the procedure. If you click Submit, the procedure is submitted for execution right away. Click **Save** to save this as a template; this feature is particularly useful with lockdowns. For example, if you want to create a template with lockdowns, and allow other users with operator privilege to run the template multiple times with minor modifications.

Click **Cancel** to exit the procedure configuration.
 17. After submitting, you can track the progress of the provisioning operation from the Procedure Activity page. For more information about this, see [About Deployment Procedures](#).
 18. To view the newly provisioned target, from the **Targets** menu, select **Middleware**.

Provisioning the SOA Domain and Oracle Homes

This chapter describes how you can use the Middleware Provisioning solution offered in Enterprise Manager Cloud Control to provision a SOA Domain or/and an Oracle Home.

 Note:

Before provisioning a SOA Domain/ Oracle Home, sign in to My Oracle Support Web site at <https://support.oracle.com> and download Patch ID 20046898.

In particular, this chapter contains the following topics:

- [Getting Started with Provisioning SOA Domain and Oracle Home](#)
- [Source Environment and Destination Environment after SOA Provisioning](#)
- [Supported Versions of SOA for Provisioning](#)
- [Before you Begin Provisioning SOA Domain and Oracle Home](#)
- [Use Case 1: First Time Provisioning of a SOA Domain](#)
- [Use Case 2: Provisioning from a SOA Oracle Home Based Provisioning Profile](#)
- [Use Case 3: Cloning from a Provisioning Profile based on an Existing SOA Domain](#)
- [Use Case 4: Provisioning from an Existing SOA Home](#)
- [Use Case 5: Scaling Up an Existing SOA Domain](#)
- [Use Case 6 - Cloning a database in WebLogic Domain](#)

 Note:

Enterprise Manager does not support creating domain profiles for SOA, OSB domain targets of version 12.2.1.3.0 and above only. If you attempt to create such a profile by clicking **Create Button**, you would run into errors.

Getting Started with Provisioning SOA Domain and Oracle Home

This section helps you get started by providing an overview of the steps involved in provisioning WebLogic Domain and Middleware Home using the Fusion Middleware Deployment procedure.

Table 26-1 Getting Started with SOA Provisioning

Step	Description	Reference Links
Step 1	Selecting the Use Case. This chapter covers the use cases for provisioning Oracle SOA Domain and Oracle SOA Home. Select the use case that best matches your requirements.	- To learn about first time provisioning of a SOA Domain, see Use Case 1: First Time Provisioning of a SOA Domain. - To learn about provisioning from an SOA home based provisioning profile, see Use Case 2: Provisioning from a SOA Oracle Home Based Provisioning Profile. - To learn about cloning from an existing SOA domain, see Use Case 3: Cloning from a Provisioning Profile based on an Existing SOA Domain. - To learn about provisioning from an existing SOA home, see Use Case 4: Provisioning from an Existing SOA Home. - To learn about scaling out a SOA domain, see Use Case 5: Scaling Up an Existing SOA Domain. - To learn about cloning a database in WebLogic domain, see Use Case 6 - Cloning a database in WebLogic Domain
Step 2	Meeting Prerequisites to Provision a Middleware Profile Before you run the Fusion Middleware Deployment Procedure, there are a few prerequisites that you must meet.	To learn about the prerequisites for provisioning a SOA domain or home, see Before you Begin Provisioning SOA Domain and Oracle Home .
Step 3	Running the Fusion Middleware Deployment Procedure Run this deployment procedure to successfully provision a Weblogic Domain and/or an Oracle Home.	To learn about provisioning from an Installation Media Profile or an Oracle Home Profile, see Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile . To learn about provisioning from a WebLogic Domain Profile, see Provisioning a Fusion Middleware Domain from an Existing Oracle Home . To provision from an existing home, see Cloning from an Existing WebLogic Domain Based-Profile . To scale out from a SOA domain, see Scaling Up / Scaling Out Fusion Middleware Domains .

Source Environment and Destination Environment after SOA Provisioning

This section describes the middleware components the source environment contains before provisioning, and the components that get provisioned after you run the Provi-

sion Fusion Middleware Deployment Procedure. Primarily, two use cases are being described here using a typical SOA topology. First use case is of a fresh provisioning, where you start with an Installation Media based profile or an Oracle Home based profile, and provision the fresh domain and/or Oracle Home. The second use case describes how you can clone an existing SOA domain.

This section contains the following topics:

- [Source and Destination Environments for a Fresh SOA Provisioning Use Case](#)
- [Source and Destination Environments for SOA Cloning Use Case](#)


Source and Destination Environments for a Fresh SOA Provisioning Use Case

For a fresh SOA provisioning use case, before you begin, you must ensure that you have met the following topology requirements:

- Oracle HTTP Server has been installed and discovered.
- Virtual IP Address 1 and Virtual IP Address 2 have already been reserved.
- APPHOST1 and APPHOST2 must be discovered in Cloud Control.
- Database should have been discovered.

Note:

Ensure that Oracle HTTP Server, APPHOST1, APPHOST2, and the RAC database are being monitored as managed targets in Cloud Control.



After running the Fusion Middleware Deployment Procedure, all the products that are displayed inside the green box in the destination environment get provisioned.

Source and Destination Environments for SOA Cloning Use Case


For a SOA cloning use case, before you begin, ensure that you have met the following topology requirements:

- Oracle HTTP Server has been installed and discovered.
- Virtual IP Address 3 and Virtual IP Address 4 have already been reserved.
- APPHOST3 and APPHOST4 must be discovered in Cloud Control.
- Database must be cloned and discovered.
- If source environment is configured with Oracle ID, then OID must be cloned and discovered.



Note:

Ensure that Oracle HTTP Server, APPHOST3, APPHOST4, and the RAC database are being monitored as managed targets in Cloud Control.



For the cloning case WLS_SOA1 and WLS_SOA2 are part of APPHOST1 and APPHOST2 respectively are cloned into APPHOST3 and APPHOST4. The RAC DB is cloned separately.

Supported Versions of SOA for Provisioning

The versions listed in the following table support SOA provisioning.

Product	Version
Oracle Repository Creation Utility (RCU)	11g
Oracle SOA	12c, 11g

Before you Begin Provisioning SOA Domain and Oracle Home

You must keep the things to keep in mind before you actually start creating middle-ware profiles and provisioning from these profiles.

In particular, this section contains the following topics:

- [Create Middleware Roles and Assign Privileges to them](#)
- [Setting Named Credentials and Privileged Credentials for the Middleware Targets](#)
- [\(Applicable only for a Cloning WebLogic Domain Use Case\) Cloning a Database](#)

Create Middleware Roles and Assign Privileges to them

In a typical data center, the main users of Deployment Procedures are Designers (Lead Administrators) and Operators. Deployment Procedure privileges enable users to perform some design-time activities such as, creating the profiles, granting access-es on the profile, creating profile lock-downs, and run-time activities like running the Deployment Procedure to provision software applications.

For Middleware Provisioning, you will need to create Administrators with the following roles:

Role: EM Super Administrator

[Table 26-2](#) lists the roles based on which you can create administrators for Middleware Provisioning.

Table 26-2 Creating Administrators with the Required Roles

Enterprise Manager Role	Description
EM_PROVISIONING_DE-SIGNER	Role has privileges for provisioning designer
EM_PROVISIONING_OPER-ATOR	Role has privileges for provisioning operator

For instructions to create administrators with these roles, see [Creating Enterprise Manager User Accounts](#).

Setting Named Credentials and Privileged Credentials for the Middleware Targets

Oracle recommends that you set the Named Credentials for normal operating system user account (*Oracle*) and Named Credentials for privileged user accounts (*root*) to perform any of the provisioning tasks in Enterprise Manager Cloud Control.

For instructions to set the Named Credentials, see [Setting Up Credentials](#).

(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database

You must have cloned a database from the source domain at the same time that the domain was being cloned. To clone a database, you must have discovered the source database as a managed target in Enterprise Manager, following which you can create a profile out of the source database, and then provision the profile to complete the cloning process.

Use Case 1: First Time Provisioning of a SOA Domain

This use case describes how you can perform a first time install of Oracle Fusion Middleware software, and first time provisioning of a WebLogic Domain. This is particularly useful when you do not have an existing domain in your data center, or if you do not wish to clone from a provisioning profile based upon an existing domain. To do so, follow these steps:



Note:

If you use a Windows machine to provision the SOA Domain, after you have discovered the source SOA domain, you must bring the Node Manager down, and only then proceed with the SOA Domain Profile creation.

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create a Installation Media Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on an Installation Media](#).
2. (*optional*) You may choose to create some lock-downs and save the profile as a template after it passes the prerequisite checks. Doing so can be useful when you have to run the same profile multiple times for provisioning middleware products. The added benefit of saving the profile as a template is that you can grant access to Operators so they can run the profiles and carry out the Middleware Provisioning.
3. Before you provision a middleware profile based on an Installation Media profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
4. Log in with Operator Privileges (EM_PROVISIONING_OPERATOR). Select the template from the Deployment Procedure table, and click **Launch**.

If you have not created a template out of the profile, you can select your profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.

5. For provisioning a SOA Domain and Oracle Home from an Installation Media, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
6. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 2: Provisioning from a SOA Oracle Home Based Provisioning Profile

This use case describes how you can create a Gold Image from an existing Oracle Home (SOA) that has perhaps been patched, and provision this using the Oracle Home profile. This is particularly useful when you need to install Oracle Fusion Middleware software with or without a new WebLogic Domain.

1. Log in with EM_PROVISIONING_DESIGNER role to create an Oracle Home (SOA) Profile. To do so, follow the steps mentioned in [Creating Middleware Provisioning Profiles](#).
2. Before you provision from a middleware profile based on an Oracle Home, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
3. Select the profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For creating a clone of an existing domain's Oracle Home (with binaries and patches) but no domain configuration, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 3: Cloning from a Provisioning Profile based on an Existing SOA Domain

This use case describes how you can clone a WebLogic Domain from a provisioning profile based upon an existing WebLogic Domain. This is particularly useful when you need to install Oracle Fusion Middleware software and configure a new WebLogic Domain.

1. Log in with EM_PROVISIONING_DESIGNER role to create a WebLogic Domain Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on a WebLogic Domain](#).
2. Before you provision a middleware profile based on an WebLogic Domain profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the WebLogic Domain Profile](#).

3. Select the profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For provisioning a SOA Domain and Oracle Home from a profile, follow the steps mentioned in [Cloning from an Existing WebLogic Domain Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Existing WebLogic Domain Based-Profile](#).

Use Case 4: Provisioning from an Existing SOA Home

If you have an Oracle Home that you want to provision as it is (without having to create a profile), then you can do so by selecting the Oracle Home source target in the Provision Fusion Middleware procedure. For more information, see [Provisioning a Fusion Middleware Domain from an Existing Oracle Home](#).

Use Case 5: Scaling Up an Existing SOA Domain

To scale up a SOA Domain to include one or more managed servers, run the Scaleup/Scale Out Middleware procedure from the Deployment Procedures table on the Middleware Provisioning page. For more information, see [Scaling Up / Scaling Out Fusion Middleware Domains](#).

Use Case 6 - Cloning a database in WebLogic Domain

A database can be cloned from the source domain at the same time the domain is being cloned. To clone a database, the source database must be discovered as a managed target in Enterprise Manager. This constitutes of creating a profile of the source database and provisioning the profile.

For an overview of database cloning methods, see [Overview of Classic Cloning Methods](#).

The database cloning for WebLogic Server domain, SOA, Oracle Service Bus, WebCenter can be accomplished in the following ways:

1. The user can manually clone the database from the source domain to the destination domain before provisioning. The destination database inherits the schema from the original database.
2. The user can choose to export the database metadata as part of provisioning. This exports the product metadata necessary to clone the domain. The user can clone the application related schema later. Exporting the metadata is supported from SOA release 11g onwards and the schema creation is supported from SOA release 12c onwards.



Note:

While cloning the WebLogic domain, the entire database along with schema must be cloned in the destination domain before provisioning. Else, the cloning fails.

The following tables list the mapping of database cloning process supported with schema for the respective domains and releases.

Table 26-3 Database cloning process with schema supported for SOA domain

Case	SOA 12c	SOA 11g
User exports metadata during profile creation and provides a blank database with no schema during provisioning. The schema is created through Provisioning user interface in release 12c onwards.	Yes	No
User exports metadata during profile creation and provides a blank database with empty schema pre-created during provisioning.	Yes	Yes
User exports metadata during profile creation and provides a cloned database during provisioning.	No	No
Metadata not exported during profile creation, blank database with no schema provided during provisioning.	No	No
Metadata not exported during profile creation, blank database with schema provided during provisioning.	No	No
Metadata not exported during profile creation, cloned database provided during provisioning.	Yes	Yes

Table 26-4 Database cloning process with schema supported for Service Bus domain

Case	Service Bus 12c	Service Bus 11g
User exports metadata during profile creation and provides a blank database with no schema during provisioning. The schema is created through Provisioning user interface in release 12c onwards.	Yes	No
User exports metadata during profile creation and provides a blank database with empty schema pre-created during provisioning.	Yes	Yes

Table 26-4 (Cont.) Database cloning process with schema supported for Service Bus domain

Case	Service Bus 12c	Service Bus 11g
User exports metadata during profile creation and provides a cloned database during provisioning.	No	No
Metadata not exported during profile creation, blank database with no schema provided during provisioning.	No	No
Metadata not exported during profile creation, blank database with schema provided during provisioning.	No	No
Metadata not exported during profile creation, cloned database provided during provisioning.	Yes	Yes

Table 26-5 Database cloning process with schema supported for WebCenter domain

Case	WebCenter 12c	WebCenter 11g
User exports metadata during profile creation and provides a blank database with no schema during provisioning. The schema is created through Provisioning user interface in release 12c onwards.	Yes	No
User exports metadata during profile creation and provides a blank database with empty schema pre-created during provisioning.	Yes	Yes
User exports metadata during profile creation and provides a cloned database during provisioning.	No	No
Metadata not exported during profile creation, blank database with no schema provided during provisioning.	No	No
Metadata not exported during profile creation, blank database with schema provided during provisioning.	No	No

Table 26-5 (Cont.) Database cloning process with schema supported for WebCenter domain

Case	WebCenter 12c	WebCenter 11g
Metadata not exported during profile creation, cloned database provided during provisioning.	Yes	Yes

 Note:

In case of WebCenter, it is recommended to clone the database instead of exporting the metadata.

Table 26-6 Database cloning process with schema supported for JRF domain

Case	JRF
User exports metadata during profile creation and provides a blank database with no schema during provisioning. The schema is created through Provisioning user interface in release 12c onwards.	No
User exports metadata during profile creation and provides a blank database with empty schema pre-created during provisioning.	No
User exports metadata during profile creation and provides a cloned database during provisioning.	No
Metadata not exported during profile creation, blank database with no schema provided during provisioning.	No
Metadata not exported during profile creation, blank database with schema provided during provisioning.	No
Metadata not exported during profile creation, cloned database provided during provisioning.	Yes

 Note:

Metadata includes MDS (Metadata Services) and OPSS (Oracle Platform Security Services) data.

Provisioning the Service Bus Domain and Oracle Homes

This chapter describes how you can use the Middleware Provisioning solution offered in Enterprise Manager Cloud Control to provision an Service Bus Domain or/and an Oracle Home.

 Note:

Before provisioning an Service Bus Domain/Oracle Home, you must download and apply the patch 20046866.

In particular, this chapter contains the following topics:

- [Getting Started with Provisioning Service Bus Domain and Oracle Home](#)
- [Supported Versions of Service Bus for Provisioning](#)
- [Before you Begin Provisioning Service Bus Domain and Oracle Home](#)
- [Use Case 1: First Time Provisioning of a Service Bus Domain](#)
- [Use Case 2: Provisioning from a Service Bus Home Based Provisioning Profile](#)
- [Use Case 3: Cloning from a Provisioning Profile based on an Existing Service Bus Domain](#)
- [Use Case 4: Provisioning from an Existing Service Bus Home](#)
- [Use Case 5: Scaling Up an Existing Service Bus Domain](#)

Getting Started with Provisioning Service Bus Domain and Oracle Home

This section helps you get started by providing an overview of the steps involved in provisioning WebLogic Domain and Middleware Home using the Fusion Middleware Deployment procedure.

Table 27-1 Getting Started with Service Bus Provisioning

Step	Description	Reference Links
Step 1	Selecting the Use Case. This chapter covers the use cases for provisioning an Service Bus Domain and Service Bus Home. Select the use case that best matches your requirements.	• To learn about first time provisioning of an Service Bus Domain, see Use Case 1: First Time Provisioning of a Service Bus Domain. • To learn about provisioning from an Service Bus home based provisioning profile, see Use Case 2: Provisioning from a Service Bus Home Based Provisioning Profile. • To learn about cloning from an existing Service Bus domain, see Use Case 3: Cloning from a Provisioning Profile based on an Existing Service Bus Domain. • To learn about provisioning from an existing Service Bus home, see Use Case 4: Provisioning from an Existing Service Bus Home. • To learn about scaling out an Service Bus domain, see Use Case 5: Scaling Up an Existing Service Bus Domain.
Step 2	Meeting Prerequisites to Provision a Middleware Profile Before you run the Fusion Middleware Deployment Procedure, there are a few prerequisites that you must meet.	To learn about the prerequisites for provisioning an Service Bus domain or home, see Before you Begin Provisioning Service Bus Domain and Oracle Home .
Step 3	Running the Fusion Middleware Deployment Procedure Run this deployment procedure to successfully provision a Weblogic Domain and/or an Oracle Home.	To learn about provisioning from an Installation Media Profile or an Oracle Home Profile, see Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile. To learn about provisioning from a WebLogic Domain Profile, see Provisioning a Fusion Middleware Domain from an Existing Oracle Home. To provision from an existing home, see Cloning from an Existing WebLogic Domain Based-Profile. To scale out from an Service Bus domain, see Scaling Up / Scaling Out Fusion Middleware Domains.

Supported Versions of Service Bus for Provisioning

The versions listed in the following table support Service Bus provisioning.

Product	Version
Oracle Repository Creation Utility (RCU)	11g
Oracle Service Bus	11g

Before you Begin Provisioning Service Bus Domain and Oracle Home

You must keep the things to keep in mind before you actually start creating middle-ware profiles and provisioning from these profiles.

In particular, this section contains the following topics:

- [Create Middleware Roles and Assign Privileges](#)
- [Setting Named Credentials and Privileged Credentials for the Middleware Targets](#)
- [\(Applicable only for a Cloning WebLogic Domain Use Case\) Cloning a Database](#)

Create Middleware Roles and Assign Privileges

In a typical data center, the main users of Deployment Procedures are Designers (Lead Administrators) and Operators. Deployment Procedure privileges enable users to perform some design-time activities such as creating the profiles, granting accesses on the profile, creating profile lock-downs, and run-time activities like running the Deployment Procedure to provision software applications.

For Middleware Provisioning, you will need to create Administrators with the following roles:

Role: EM Super Administrator

[Table 27-2](#) lists the roles based on which you can create administrators for Middleware Provisioning.

Table 27-2 Creating Administrators with the Required Roles

Enterprise Manager Role	Description
EM_PROVISIONING_DE-SIGNER	Role has privileges for provisioning designer
EM_PROVISIONING_OPER-ATOR	Role has privileges for provisioning operator

For instructions to create administrators with these roles, see [Creating Enterprise Manager User Accounts](#).

Setting Named Credentials and Privileged Credentials for the Middle-ware Targets

Oracle recommends that you set the Named Credentials for normal operating system user account (*Oracle*) and Named Credentials for privileged user accounts (*root*) to perform any of the provisioning tasks in Enterprise Manager Cloud Control.

For instructions to set the Named Credentials, see [Setting Up Credentials](#).

(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database

You must have cloned a database from the source domain at the same time that the domain was being cloned. To clone a database, you must have discovered the source database as a managed target in Enterprise Manager, following which you can create a profile out of the source database, and then provision the profile to complete the cloning process.

Use Case 1: First Time Provisioning of a Service Bus Domain

This use case describes how you can perform a first time install of Oracle Fusion Middleware software, and first time provisioning of a WebLogic Domain. This is particularly useful when you do not have an existing domain in your data center, or if you do not wish to clone from a provisioning profile based upon an existing domain. To do so, follow these steps:

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create a Installation Media Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on an Installation Media](#).
2. *(optional)* You may choose to create some lock-downs and save the profile as a template after it passes the prerequisite checks. Doing so can be useful when you have to run the same profile multiple times for provisioning middleware products. The added benefit of saving the profile as a template is that you can grant access to Operators so they can run the profiles and carry out the Middleware Provisioning.
3. Before you provision a middleware profile based on an Installation Media profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
4. Log in with Operator Privileges (EM_PROVISIONING_OPERATOR). Select the template from the Deployment Procedure table, and click **Launch**.

If you have not created a template out of the profile, you can select your profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
5. For provisioning a Service Bus Domain and Oracle Home from an Installation Media, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
6. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 2: Provisioning from a Service Bus Home Based Provisioning Profile

This use case describes how you can create a Gold Image from an existing Oracle Home (Service Bus) that has perhaps been patched, and provision this using the Ora-

cle Home profile. This is particularly useful when you need to install Oracle Fusion Middleware software with or without a new WebLogic Domain.

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create an Oracle Home (Service Bus) Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on an Oracle Home](#).
2. Before you provision a middleware profile based on an Oracle Home, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
3. Select the profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For creating a clone of an existing domain's Oracle Home (with binaries and patches) but no domain configuration, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 3: Cloning from a Provisioning Profile based on an Existing Service Bus Domain

This use case describes how you can clone a WebLogic Domain from a provisioning profile based upon an existing WebLogic Domain. This is particularly useful when you need to install Oracle Fusion Middleware software and configure a new WebLogic Domain.

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create a WebLogic Domain Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on a WebLogic Domain](#).
2. Before you provision a middleware profile based on a WebLogic Domain profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the WebLogic Domain Profile](#).
3. Select the profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For provisioning a Service Bus Domain and Oracle Home from a profile, follow the steps mentioned in [Cloning from an Existing WebLogic Domain Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Existing WebLogic Domain Based-Profile](#).

Use Case 4: Provisioning from an Existing Service Bus Home

If you have an Oracle Home that you want to provision as it is (without having to create a profile), then you can do so by select the Oracle Home source target in the Provision Fusion Middleware procedure. For more information, see [Provisioning a Fusion Middleware Domain from an Existing Oracle Home](#).

Use Case 5: Scaling Up an Existing Service Bus Domain

To scale up a Service Bus Domain to include one or more managed servers, run the Scaleup/Scale Out Middleware procedure from the Deployment Procedures table on the Middleware Provisioning page. For more information, see [Scaling Up / Scaling Out Fusion Middleware Domains](#) .

Provisioning the Oracle WebCenter Domain and Oracle Homes

This chapter describes how you can use the Middleware Provisioning solution offered in Enterprise Manager Cloud Control to provision a WebCenter Domain or/and an Oracle Home.

In particular, this chapter contains the following topics:

- [Getting Started with Provisioning WebCenter Domain and Oracle Home](#)
- [About WebCenter Topologies Supported in Enterprise Manager](#)
- [Supported Versions of WebCenter for Provisioning](#)
- [Source Environment and Destination Environment after WebCenter Provisioning](#)
- [Before you Begin Provisioning WebCenter Domain and Oracle Home](#)
- [Use Case 1: First Time Provisioning of a WebCenter Portal with Lock-downs](#)
- [Use Case 2: Provisioning a WebCenter Home](#)
- [Use Case 3: Cloning an Existing WebCenter Portal Environment](#)
- [Use Case 4: Provisioning from an Existing WebCenter Home](#)
- [Use Case 5: Scaling Up an Existing WebCenter Domain](#)

Getting Started with Provisioning WebCenter Domain and Oracle Home

This section helps you get started by providing an overview of the steps involved in provisioning WebLogic Domain and Middleware Home using the Fusion Middleware Deployment procedure.

Table 28-1 Getting Started with WebCenter Provisioning

Step	Description	Reference Links
Step 1	Selecting the Use Case. This chapter covers the use cases for provisioning WebCenter Domain and WebCenter Home. Select the use case that best matches your requirements.	- To learn about first time provisioning of a WebCenter Domain, see Use Case 1: First Time Provisioning of a WebCenter Portal with Lock-downs. - To learn about provisioning from a WebCenter home based provisioning profile, see Use Case 2: Provisioning a WebCenter Home. - To learn about cloning from an existing WebCenter domain, see Use Case 3: Cloning an Existing WebCenter Portal Environment. - To learn about provisioning from an existing WebCenter home, see Use Case 4: Provisioning from an Existing WebCenter Home. - To learn about scaling out a WebCenter domain, see Use Case 5: Scaling Up an Existing WebCenter Domain.
Step 2	Meeting Prerequisites to Provision a Middleware Profile Before you run the Fusion Middleware Deployment Procedure, there are a few prerequisites that you must meet.	To learn about the prerequisites for provisioning a WebCenter domain or home, see Before you Begin Provisioning WebCenter Domain and Oracle Home .
Step 3	Running the Fusion Middleware Deployment Procedure Run this deployment procedure to successfully provision a WebLogic Domain and/or an Oracle Home.	To learn about provisioning from an Installation Media Profile or an Oracle Home Profile, see Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile . To learn about provisioning from a WebLogic Domain Profile, see Provisioning a Fusion Middleware Domain from an Existing Oracle Home . To provision from an existing home, see Cloning from an Existing WebLogic Domain Based-Profile . To scale out from a WebCenter domain, see Scaling Up / Scaling Out Fusion Middleware Domains .

About WebCenter Topologies Supported in Enterprise Manager

Oracle WebCenter Portal provisioning is now supported from Enterprise Manager Cloud Control.

Primarily, the following use cases are supported:

Table 28-2 WebCenter Provisioning Use cases

Task	Description
1. WebCenter Provisioning - Fresh Install	This use case is useful if you want to provision a fresh vanilla WebCenter (Oracle WebCenter Portal + Oracle WebCenter Content) environment.

Table 28-2 (Cont.) WebCenter Provisioning Use cases

Task	Description
2. WebCenter Provisioning - Cloning (like-to-like)	This use case is useful if you already have installed WebCenter environment (Oracle WebCenter Portal + Oracle WebCenter Content), and a later point want to copy/clone from that environment.

The following topologies are supported for provisioning WebCenter:


Table 28-3 Supported Topologies for WebCenter

Topologies	Components	External Components
Development Topology	WCP, WCC, OHS, DB, embedded LDAP, SAML/SSO	DB, OHS
Production Topology (HA)	WCP, WCC, DB, OHS, LBR, SES, LDAP w/SSO	DB, OHS, OTD, SES, LDAP

Development Topology Details

Development topology supports only a single node environment which is consistent with the following topology definition.


WebCenter Provisioning - Dev Topology



Production Topology Details

Production topology supports multi-node cluster environment, which means that you can select however many nodes per cluster. After provisioning, you can even reconfigure the node count for clusters.

WebCenter Provisioning – Prod Topology



Source Environment and Destination Environment after WebCenter Provisioning

This section describes what the source environment might contain, and what gets provisioned after you run the Fusion Middleware Deployment Procedure. You can see two topologies being displayed. First topology is a fresh provisioning use case where you start with an Installation Media based profile or an Oracle Home based profile, and provision that. Second is a cloning topology, where you are creating a copy of an existing domain.

This section contains the following topics:

- [Source and Destination Environments for a Fresh WebCenter Provisioning Use Case](#)
- [Source and Destination Environments for WebCenter Cloning Use Case](#)


Source and Destination Environments for a Fresh WebCenter Provisioning Use Case

Before you begin, ensure that you have met the following topology requirements:

- Oracle HTTP Servers must have been installed and discovered.
- APPHOST1 and APPHOST2 must be discovered in Cloud Control.
- Database should have been discovered.

Note:

Ensure that Oracle HTTP Server, APPHOST1, APPHOST2, and the RAC database are being monitored as managed targets in Cloud Control.



Source and Destination Environments for WebCenter Cloning Use Case


Before you begin, ensure that you have met the following topology requirements:

- Oracle HTTP Servers have been installed and discovered.
- APPHOST3 and APPHOST4 must be discovered in Cloud Control.
- Database must be cloned and discovered.

- If source environment is configured with Oracle ID, then OID must be cloned and discovered.

 Note:

Ensure that Oracle HTTP Server, APPHOST3, APPHOST4, and the RAC database are being monitored as managed targets in Cloud Control.



Supported Versions of WebCenter for Provisioning

The versions listed in the following table support WebCenter provisioning.

Product	Version
Oracle Repository Creation Utility (RCU)	11g
Oracle WebCenter Portal	11g
Oracle WebCenter Content	11g

Before you Begin Provisioning WebCenter Domain and Oracle Home

You must keep the things to keep in mind before you actually start creating middle-ware profiles and provisioning from these profiles.

 Note:

To provision a WebCenter Domain using LDAP in the Production mode, ensure that a Weblogic user is present with Administrator group privileges in the LDAP.

Ensure that you have identical topology of servers on all nodes in a single cluster. For example, in case of a two node cluster, the same set of servers must be available on node 1 and node 2 of the cluster.

In particular, this section contains the following topics:

- [Create Middleware Roles and Assign Privileges to them](#)
- [Setting Named Credentials and Privileged Credentials for the Middleware Targets](#)
- [\(Applicable only for a Cloning WebLogic Domain Use Case\) Cloning a Database](#)

Create Middleware Roles and Assign Privileges to them

In a typical data center, the main users of Deployment Procedures are Designers (Lead Administrators) and Operators. Deployment Procedure privileges enable users to perform some design-time activities like creating the profiles, granting accesses on the profile, creating profile lock-downs, and run-time activities like running the Deployment Procedure to provision software applications.

For Middleware Provisioning, you will need to create Administrators with the following roles:

Role: EM Super Administrator

[Table 28-4](#) lists the roles based on which you can create administrators for Middleware Provisioning.

Table 28-4 Creating Administrators with the Required Roles

Enterprise Manager Role	Description
EM_PROVISIONING_DE-SIGNER	Role has privileges for provisioning designer
EM_PROVISIONING_OPER-ATOR	Role has privileges for provisioning operator

For instructions to create administrators with these roles, see [Creating Enterprise Manager User Accounts](#).

Setting Named Credentials and Privileged Credentials for the Middle-ware Targets

Oracle recommends that you set the Named Credentials for normal operating system user account (*Oracle*) and Named Credentials for privileged user accounts (*root*) to perform any of the provisioning tasks in Enterprise Manager Cloud Control.

For instructions to set the Named Credentials, see [Setting Up Credentials](#).

(Applicable only for a Cloning WebLogic Domain Use Case) Cloning a Database

You must have cloned a database from the source domain at the same time that the domain was being cloned. To clone a database, you must have discovered the source database as a managed target in Enterprise Manager, following which you can create a profile out of the source database, and then provision the profile to complete the cloning process.

Use Case 1: First Time Provisioning of a WebCenter Portal with Lock-downs

This use case describes how you can perform a first time install of Oracle Fusion Middleware software, and first time provisioning of a WebLogic Domain. This is particularly useful when you do not have an existing domain in your data center, or if you do not wish to clone from a provisioning profile based upon an existing domain. To do so, follow these steps:

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create a Installation Media Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on an Installation Media](#).
2. Before you provision a middleware profile based on an Installation Media profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
3. *(optional)* While provisioning the profile, you can create a template from the inputs that has already been entered in this deployment procedure. For this, you need to use the lockdown feature offered in Enterprise Manager, which enables you to lock select inputs thereby restricting other users to the change these values in future. For example, you can lock the values entered in the Middleware, Database, Identity and Security, and WebTier sections. By saving the Deployment Procedure with these lockdowns, other administrators can leverage the defaulted inputted values and submit the procedure with very few inputs required from them.
4. Log in with Operator Privileges (EM_PROVISIONING_OPERATOR). Select the template from the Deployment Procedure table, and click **Launch**.

If you have not created a template out of the profile, you can select your profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
5. For provisioning a WebCenter Domain and Oracle Home from an Installation Media, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
6. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 2: Provisioning a WebCenter Home

This use case describes how you can create a Gold Image from an existing Oracle Home (WebCenter) that has perhaps been patched, and provision this using the Ora-

cle Home profile. This is particularly useful when you need to install Oracle Fusion Middleware software with or without a new WebLogic Domain.

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create an Oracle Home (WebCenter) Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on an Oracle Home](#).
2. Before you provision a middleware profile based on an Oracle Home, meet the prerequisites mentioned in [Prerequisites for Provisioning the Installation Media Profile or the Oracle Home Profile](#).
3. Log in with Designer/Operator privileges, select your profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For creating a clone of an existing domain's Oracle Home (with binaries and patches) but no domain configuration, follow the steps mentioned in [Provisioning of a new Fusion Middleware Domain from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Installation Media Based-Profile or an Oracle Home Based-Profile](#).

Use Case 3: Cloning an Existing WebCenter Portal Environment

This use case describes how you can clone a WebLogic Domain from a provisioning profile based upon an existing WebLogic Domain. This is particularly useful when you need to install Oracle Fusion Middleware software and configure a new WebLogic Domain.

1. Log in with Designer privileges (EM_PROVISIONING_DESIGNER role) to create a WebLogic Domain Profile. To do so, follow the steps mentioned in [Creating a Provisioning Profile Based on a WebLogic Domain](#).
2. Before you provision a middleware profile based on a WebLogic Domain profile, meet the prerequisites mentioned in [Prerequisites for Provisioning the WebLogic Domain Profile](#).
3. Log in with Designer/Operator Privileges, select your profile from the Profiles table on the Middleware Provisioning page, then click **Provision**.
4. For provisioning a WebCenter Domain and Oracle Home from a profile, follow the steps mentioned in [Cloning from an Existing WebLogic Domain Based-Profile](#).
5. If you want to customize the settings in the destination environment, follow the steps mentioned in [Customizing the Destination Environment from an Existing WebLogic Domain Based-Profile](#).

Use Case 4: Provisioning from an Existing WebCenter Home

If you have an Oracle Home that you want to provision as it is (without having to create a profile), then you can do so by select the Oracle Home source target in the Provision Fusion Middleware procedure. For more information, see [Provisioning a Fusion Middleware Domain from an Existing Oracle Home](#).

Use Case 5: Scaling Up an Existing WebCenter Domain

To scale up a WebCenter Domain to include one or more managed servers, run the Scaleup/Scale Out Middleware procedure from the Deployment Procedures table on the Middleware Provisioning page. For more information, see [Scaling Up / Scaling Out Fusion Middleware Domains](#) .

Middleware Provisioning using the EM CLI

This chapter describes how you can use the command line option Enterprise Manager Command Line Interface offered by Oracle to create, describe, list, delete, and customize the Middleware Profiles.

In particular, this chapter covers the following:

- [Creating Middleware Provisioning Profiles](#)
- [Submitting the Procedure using EM CLI](#)
- [Listing Middleware Provisioning Profiles](#)
- [Describing Provisioning Profiles](#)
- [Deleting Provisioning Profiles](#)

Creating Middleware Provisioning Profiles

Profiles are like templates that you can create and store in Software Library. Once a profile is created, it can be launched numerous times to provision WebLogic Domain and/or Oracle Home. The advantage of using a profile is that you can ensure that future WebLogic installations follow a standard, consistent configuration.

This chapter covers following sections:

- [Creating a WebLogic Domain Profile](#)
- [Creating an Oracle Home Profile](#)
- [Creating an Installation Media Profile](#)

Creating a WebLogic Domain Profile

To create a Fusion Middleware Provisioning Profile from a WebLogic Domain, use the EM CLI verb `create_fmws_domain_profile`.

```
emcli create_fmws_domain_profile
-name="Profile Name"
-ref_target="Reference Target Name"
[-description="Profile Description"]
[-oh_cred="Oracle Home Owner Credentials"]
[-includeOh]
[-schedule=
start_time:yyy/MM/dd HH:mm;
[tz:{java timezone ID}];
[grace_period:xxx];
]
```

[] indicates that the parameter is optional.

Description for the options:

-name

Name of the WebLogic Domain profile.

`-ref_target`

Name of the reference target used to create the WebLogic Domain profile.

`-description`

A short description for the WebLogic Domain profile you create.

`-oh_cred`

Named credential that will be used to access the reference host.

Format: CREDENTIAL_NAME:CREDENTIAL_OWNER.

All operations will be performed on the Administration Server host.

Credentials of the Oracle Home owner on the Administration Server host are required.

If no named credential is provided, then preferred host credentials for the Oracle Home target will be used.

`-wls_cred`

Named credential used to access the Administration Server. This is an optional parameter.

To pass the credential parameter, enter a name:value pair in the following format:

credential_name:credential_owner.

Where,

Credential_name is the name of the named credential.

Credential_owner is the credentials of the Administrator of the WebLogic Domain.

All operations are performed in online mode (using T2P) in case of a Fusion Middleware domain.

If no named credential is provided, the preferred administrator credentials for the domain target will be used.

`-includeOh`

Whether the Oracle Home binaries have to be included in the profile or not.

`-schedule`

The schedule for the Deployment Procedure.

If not specified, the procedure will be executed immediately.

start_time: when the procedure should start.

tz: the timezone ID.

grace_period: grace period in minutes.

Following examples describe how to create a WebLogic Domain profile:

- A WebLogic Domain called BitlessDomainProfile is created using the reference target /Farm01_base_domain/base_domain at the specified schedule. Since the Oracle home parameter is not passed, a plain WebLogic domain without the Oracle home binaries is created. Also, since the Oracle home credentials haven't been specified, the preferred host credentials for the target home is used.

```
emcli create_fmw_domain_profile
```

```
-name="BitlessDomainProfile"
```

```
-ref_target="/Farm01_base_domain/base_domain"
```

```
-description="A domain profile without software bits"
```

```
-schedule="start_time:2014/6/21 21:23;tz:America/New_York;grace_period:60"
```

- A WebLogic Domain profile along with Oracle home binaries called *DomainProfile-withBits* is created immediately. This profile is created from the reference target /Farm01_base_domain/base_domain using the specified named credentials.

```
emcli create_fmw_domain_profile
-name="DomainProfileWithBits"
-ref_target="/Farm01_base_domain/base_domain"
-oh_cred="MY_HOST_CRED:SYSMAN"
-includeMWH
```

- The is created in the Software Library in Fusion Middleware Provisioning/Profiles directory. Since the schedule is not mentioned, the job runs immediately.

```
emcli create_fmw_domain_profile
-name=D=SoaProfile
-ref_target="Farm01_SoaDomain/SoaDomain"
```

Output:

For example you will see following attributes when the profile has been submitted successfully:

- instance_name: 'CreateFmwProfile-SoaProfile_SYS-MAN_07_09_2014_11_36_AM'
- instance_guid: 'FDC7FC56E2CF2972E04373B1F00A1512'



Note:

To track the status of the profile being created, use the command:

```
emcli get_instance_status -instance=FDC7FC56E2CF2972E04373B1F00A1512 -xml -details -showJobOutput
```

Creating an Oracle Home Profile

To create a Fusion Middleware Provisioning Profile from an Oracle Home target, use the EM CLI verb `create_fmw_home_profile_verb`.

```
emcli create_fmw_home_profile
-name="Profile Name"
-ref_target="Reference Target Name"
[-description="Profile Description"]
[-oh_cred="Oracle Home Owner Credentials"]
[-schedule=
start_time:yyy/MM/dd HH:mm;
[tz:{java timezone ID}];
[grace_period:xxx];
]
```

[] indicates that the parameter is optional.

Description for the options:

-name

Name of the Oracle home profile.

-ref_target

Name of the Oracle Home target used as reference to create the profile.

-description

A short description for the Oracle home profile you create.

-oh_cred

Named credentials used to access the reference host. This is an optional parameter. To pass the credential parameter, enter a name:value pair in the following format: credential_name:credential_owner.

Where,

Credential_name is the name of the named credential.

Credential_owner is the credentials of the Oracle home owner on the Administration Server host.

If no named credential is provided, the preferred host credential for the Oracle home target will be used.

-schedule

The schedule for the Deployment Procedure.

If not specified, the procedure will be executed immediately.

To specify a value, enter:

start_time: when the procedure should start.

tz: the timezone ID.

grace_period: grace period in minutes.

Following examples describe how to create an Oracle Home profile:

- An Oracle Home profile is created in the Software Library in Fusion Middleware Provisioning/Profiles directory. Since the schedule is not mentioned, the job runs immediately.

```
emcli create_fmw_home_profile
-name=SoaProfile
-ref_target="WebLogicServer10_3_6_0_<agent host name>.example.com_6316"
```

Output:

You will see the following attributes when the profile has been submitted successfully:

- instance_name: 'CreateFmwProfile-SoaProfile_SYS-MAN_07_09_2014_11_36_AM'
- instance_guid: 'FDC7FC56E2CF2972E04373B1F00A1512'

 Note:

To track the status of the profile being created, use the command:

```
emcli get_instance_status -instance=FDC7FC56E2CF2972E04373B1F00A1512 -xml -details -showJobOutput
```

- An Oracle home profile called OhProfile1 is created using the reference host target at the specified schedule using preferred credentials.

```
emcli create_fmw_home_profile
-name="OhProfile1"
-ref_target="WebLogicServer_10.3.6.0_host.example.com_5033"
-description="An Oracle Home profile"
-schedule="start_time:2014/6/21 21:23;tz:America/New_York;grace_period:60"
```

- An Oracle home profile named OhProfile2 is created immediately from the reference host target using named credentials.

```
emcli create_fmware_profile
-name="OhProfile2" -ref_target="WebLogicServer_11.4.1.0_host.example.com_5023"
-oh_cred="MY_HOST_CRED:SYSMAN"
```

Creating an Installation Media Profile

To create a profile using the Oracle Fusion Middleware installation media, use the EM CLI verb `create_inst_media_profile_verb`

```
emcli create_inst_media_profile
-name="Profile Name"
-host="Reference Target Name"
-version="Reference Target Name"
-platform="Reference Target Name"
[-description="Profile Description"]
[-host_cred="Oracle Home Owner Credentials"]
-files=
WebLogic:WLSFile1;
SOA:SOAFile1,SOAFile2;
OSB:OSBFile;
RCU:RCUFile;
```

[] indicates that the parameter is optional.

Description for the options:

`-name`

Name of the installation media profile.

`-host`

Name of the host target where the installation media files are stored.

`-version`

Version of the installation media files.

`-platform`

Platform for which the installation media is applicable.

`-description`

A short description for the installation media profile.

`-host_cred`

Named credentials used to access the reference host. This is an optional parameter. To pass the credential parameter, enter a name:value pair in the following format: `credential_name:credential_owner`.

Where,

`Credential_name` is the name of the named credential.

`Credential_owner` is the credentials of the Oracle home owner on the Administration Server host.

If no named credential is provided, the preferred host credential for the host target will be used.

.

`-files`

List of files that have to be uploaded to Software Library. These files are passed in the format `product:file1,file2`. The Installation Media profile supports the following products: WebLogic, SOA, OSB, and RCU. Note that, to create any of these profiles, you must upload the WebLogic media file.

Following examples describe how to create an Oracle Home profile:

- Upload an installation media profile named WebLogic1036Installer containing WebLogic jar file called wls1036_generic.jar to the Software Library from the host myhost.example.com. Additionally, you must provide the platform and version details for the installation media file. Since you have not passed any credential as a parameter, the preferred host credentials will be used to access the WebLogic jar file.

```
emcli create_inst_media_profile
-name="WebLogic1036Installer"
-host="host.example.com"
-description="WebLogic Server 10.3.6.0 installer"
-version="10.3.6.0"
-platform="Generic"
-files="WebLogic:/u01/media/weblogic/wls1036_generic.jar"
```

- Upload the SOA and WLS installation media files to the Software Library from the host myhost.example.com. Additionally, you must provide the platform and version details for the installation media file. The profile called SOA+WLSInstaller is created using the named credentials specified.

```
emcli create_inst_media_profile
-name="SOA+WLSInstaller"
-host="myhost.example.com"
-description="SOA 11.1.1.7.0 and WebLogic Server 10.3.6.0 installer"
-version="11.1.1.7.0"
-platform="Generic"
-host_cred="MY_HOST_CRED:SYSMAN"
-files="WebLogic:/u01/media/weblogic/wls1036_generic.jar;SOA:/u01/media/soa/soa1.zip,/u01/media/soa/soa2.zip"
```

- An Oracle Home profile is created in the Software Library in Fusion Middleware Provisioning/Profiles directory. Since the schedule is not mentioned, the job runs immediately.

```
emcli create_inst_media_profile
-name="WebCenter Install Media"
-host="slc03qtn.example.com"
-description="WC 11.1.1.8.0 and WebLogic Server 10.3.6.0 installer and RCU"
-version="11.1.1.8.0"
-platform="2000"
-host_cred="BBANTHIA:SYSMAN"
-files="WebLogic:/net/adcnas438/export/emgpqa/provisioning/fmwprov/linux64/shiphomes/upinsmed_wls/wls1036_generic.jar;WCP:/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/shiphomes/wcp_11.1.1.8.0/wc.zip;WCC:/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/shiphomes/wcc_11.1.1.8.0/ecm_main1.zip,/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/shiphomes/wcc_11.1.1.8.0/ecm_main2.zip;RCU:/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/shiphomes/wrcu_11.1.1.8.0/rcuHome.zip"
```

Submitting the Procedure using EM CLI

To automate the process of provisioning using the command line, submit your procedure using the Enterprise Manager Command Line Interface (EMCLI) utility. The EMCLI enables you to access Enterprise Manager Cloud Control functionality from text-based consoles (shells and command windows) for a variety of operating systems. You can call Enterprise Manager functionality using custom scripts, such as SQL*Plus, OS shell, Perl, or Tcl, thus easily integrating Enterprise Manager functionality with your company's business process.

 Note:

Use the procedure below to generate a template input file based on configuration deployment procedure functionality:

1. From the **Middleware Provisioning** page launch the profile.
2. Enter the necessary values.
3. Provide a name to the provisioning deployment procedure and save.
4. Log in to EMCLI using the following commands:

```
emcli login -username=sysman -password=sysman
```

```
emcli sync
```
5. Run the following command to fetch the GUID of the configuration deployment procedure:

```
emcli get_procedures | grep <deployment procedure name>
```
6. Run the following command to fetch an example input file which contains all the values:

```
emcli describe_procedure_input -procedure=1965463666C90D9CE053E37BF00A8C65
```

The result of the above command becomes the content of the input file.
7. Make the required changes to the input file including the password.
8. Run the following commands to submit the deployment procedure:

```
emcli submit_procedure -input_file=data:/scratch/<input file name>.xml -name=ProvisionMiddleware -instance_name=<new deployment procedure name>
```

Listing Middleware Provisioning Profiles

This section describes how you can list all the Fusion Middleware Provisioning profiles that you have created in Enterprise Manager.

In particular, this section covers the following topics:

- [Listing All the Profiles](#)
- [Listing All the WebLogic Domain Profiles](#)
- [Listing All the Oracle Home Profiles](#)
- [Listing All the Installation Media Profiles](#)

Listing All the Profiles

To list all the Fusion Middleware Profiles, run the following command:

```
emcli list_fmw_profiles  
[-source_type="Profile Source"]
```


Description for the options:

-source_type
Valid values: weblogic_domain, oracle_home or install_media.
Profiles of only this type as source will be shown.

Following examples describe how to list an Oracle Home profile:

- To list all the Middleware Provisioning Profiles, run the following command:

```
emcli list_fmw_profiles
```

Output:

Location	Products	Platform	Version	Owner
Fusion Middleware Provisioning/ Profiles/SoaBitlessProfile	Oracle SOA	Linux x86-64	10.3.6.0	SYSMAN
Fusion Middleware Provisioning/ Profiles/SoaGoldImage	Oracle SOA	Linux x86-64	10.3.6.0	SYSMAN
Fusion Middleware Provisioning/ Profiles/WebCenter Install Media	Oracle WebCenter Content	Generic Platform	11.1.1.8.0	SYSMAN

Listing All the WebLogic Domain Profiles

To list all the WebLogic Domain Profiles, run the following command:

```
emcli list_fmw_profiles  
[-source_type="Profile Source"]
```

Description for the options:

-source_type
Valid values: weblogic_domain, oracle_home or install_media.
Profiles of only this type as source will be shown.

Following examples describe how to list a domain based profile:

- To list all the WebLogic Domain Provisioning Profiles, run the following command:

```
emcli list_fmw_profiles  
-source_type=weblogic_domain
```

Output:

Location	Products	Platform	Version	Owner
Fusion Middleware Provisioning/ Profiles/SoaBitlessProfile	Oracle SOA	Linux x86-64	10.3.6.0	SYSMAN

Listing All the Oracle Home Profiles

To list all the Oracle Home Profiles, run the following command:

```
emcli list_fmw_profiles  
[-source_type="Profile Source"]
```

Description for the options:

-source_type

Valid values: weblogic_domain, oracle_home or install_media.
Profiles of only this type as source will be shown.

Following examples describe how to list a Oracle Home based profile:

- To list all the WebLogic Domain Provisioning Profiles, run the following command:

```
emcli list_fmw_profiles  
-source_type=oracle_home
```

Output:

Location	Products	Platform	Version	Owner
Fusion Middleware Provisioning/ Profiles/SoaGoldImage	Oracle SOA	Linux x86-64	10.3.6.0	SYSMAN

Listing All the Installation Media Profiles

To list all the Installation Media profiles, run the following command:

```
emcli list_fmw_profiles  
[-source_type="Profile Source"]
```

Description for the options:

-source_type
Valid values: weblogic_domain, oracle_home or install_media.
Profiles of only this type as source will be shown.

Following examples describe how to list an Oracle Home profile:

- To list all the Install Media based Profiles, run the following command:

```
emcli list_fmw_profiles  
-source_type=install_media
```

Output:

Location	Products	Platform	Version	Owner
Fusion Middleware Provisioning/ Profiles/WebCenter Install Media	Oracle WebCenter Content	Generic Platform	11.1.1.8.0	SYSMAN

Describing Provisioning Profiles

To view the summary of the Fusion Middleware Provisioning Profile that you have created in Cloud Control, run the following command:

Usage:

```
emcli describe_fmw_profile  
-location="Profile Location"
```

Description for the options:

-location
Complete Software Library path to the Profile.
Use list_fmw_profiles to know the complete path.

In particular, this section covers the following topics:

- [Describing a WebLogic Domain Profile](#)
- [Describing an Oracle Home Profile](#)
- [Describing an Installation Media Profile](#)

Describing a WebLogic Domain Profile

To view a summary of the WebLogic Domain Profile, run the following command:

Usage:

```
emcli describe_fmw_profile  
-location="Fusion Middleware Provisioning/Profiles/SoaBitlessProfile"
```

Output

```
Location: Fusion Middleware Provisioning/Profiles/SoaBitlessProfile  
Created By: SYSMAN  
Created On: Sun, 6 Jul 2014  
URN: oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1
```

Attachments:

Name	Size	Content Type	Added By	Added On
nmMovePlan.xml	6.95 KB	text/plain	SYSMAN	Sun, 6 Jul 2014
mdm-url-resolver.xml	268 Bytes	text/plain	SYSMAN	Sun, 6 Jul 2014
domainMovePlan.xml	56.91 KB	text/plain	SYSMAN	Sun, 6 Jul 2014

```
Source: WebLogic Domain  
Oracle Home Included: No  
Node Manager Included: Yes  
Platform: Linux x86-64  
Version: 10.3.6.0
```

Oracle Homes:

```
Product: Oracle SOA  
Version: 11.1.1.7.0
```

WebLogic Domain:

```
Domain Name: SoaDomain  
Size: 611.45 MB  
Products: Oracle SOA
```

Name	Host	Listen Address	Port Machine	Cluster	Maximum Heap	Average CPU Usage
AdminServer	<agent host name>.example.com	<agent host name>.example.com	7001 LocalMachine		1365	0.361
Soa_server1	<agent host name>.example.com	<agent host name>.example.com	8001LocalMachine	cluster_soa	1365	0.717

Data Sources:

Name	Username	URL	Target
EDNDataSource	SOA_SOAINFRA	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
EDNLocalTxData-Source	SOA_SOAINFRA	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
OraSDPMDData-Source	SOA_SOAINFRA	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
OraSDPMDData-Source	SOA_SOAINFRA	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
SOALocalTxData-Source	SOA_SOAINFRA	dbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
mds-owsm	SOA_MDS	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa
mds-soa	SOA_MDS	jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com	cluster_soa

JMS Servers:

Name	Persistent Store	Directory	Target
BPMJMSServer_auto_1	BPMJMSFileStore_auto_1	UMSJMSFileStore_auto_1	soa_server1
PS6SOAJMSServer_auto_1	PS6SOAJMSFileStore_auto_1	UMSJMSFileStore_auto_1	soa_server1
SOAJMSServer_auto_1	SOAJMSFileStore_auto_1	SOAJMSFileStore_auto_1	soa_server1
UMSJMSServer_auto_1	UMSJMSFileStore_auto_1	UMSJMSFileStore_auto_1	soa_server1

Describing an Oracle Home Profile

To view a summary of the Oracle Home Profile, run the following command:

Usage:

```
emcli describe_fmw_profile
-location="Fusion Middleware Provisioning/Profiles/SoaGoldImage"
```

Output

```
Location: Fusion Middleware Provisioning/Profiles/SoaGoldImage
Created By: SYSMAN
Created On: Tue, 8 Jul 2014
URN: oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1
Source: Oracle Home
Platform: Linux x86-64
Version: 10.3.6.0
```

Oracle Homes:

```
Product: Oracle SOA
Version: 11.1.1.7.0
```

Describing an Installation Media Profile

To view a summary of the Installation Media Profile, run the following command:

Usage:

```
emcli describe_fmw_profile  
-location="Fusion Middleware Provisioning/Profiles/WebCenter Install Media"
```

Output:

```
Location: Fusion Middleware Provisioning/Profiles/WebCenter Install Media  
Description: WC 11.1.1.8.0 and WebLogic Server 10.3.6.0 installer and RCU  
Created By: SYSMAN  
Created On: Sun, 6 Jul 2014  
URN: oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMedia:FD8AA7EEEEA369A5E04373B1F00AE77C:0.1  
Source: Installation Media  
Product: Oracle WebCenter  
Platform: Generic Platform  
Version: 11.1.1.8.0
```

Files:

Product	File	Size
Oracle WebCenter Content	ecm_main1.zip	2.34 GB
	ecm_main2.zip	1.91 GB
Oracle RCU	PS6SOAJMSFileStore_auto_1	496.16 MB
Oracle WebCenter Portal	SOAJMSFileStore_auto_1	1.96 GB
Oracle WebLogic Server	UMSJMSFileStore_auto_1	1,019.01 MB

Describing an Exported Domain/Partiton Profile

To view a summary of the Exported Domain/Partiton Profile, run the following command:

Usage:

```
emcli describe_fmw_profile  
-location="WebLogic Domain Partition Provisioning/Profiles/exportDomain"
```

Output

```
Location: WebLogic Domain Partition Provisioning/Profiles/exportDomain  
Created By: SYSMAN  
Created On: Mon, 16 Oct 2017  
URN: oracle:oracle:defaultService:em:provisioning:1:cmp:COMP_Component:PartitionProfile:5BB150A40646CFC6E0535B31C40A1737:0.1
```

Attachments:

Name	Size	Content Type	Added By	Added On
base_domain-attributes.json	1.64 KB	text/plain	SYSMAN	Mon, 16 Oct 2017

Type: WebLogic Domain Partition Profile
Oracle Home Included: No
Platform: Linux x86-64
Version: 12.2.1.3.0
Source: WebLogic Domain
Domain Name: base_domain

Virtual Targets

Name	URI	Prefix	Host Name	Target
\${PARTITION_NAME}	-AdminServer	-virtualTarget	/\${PARTITION_NAME}	AdminServer
\${PARTITION_NAME}	-Cluster_1	-virtualTarget	/\${PARTITION_NAME}	Cluster_1

Example Response:

```
{
  "location": "WebLogic Domain Partition Provisioning/Profiles/exportedDomain",
  "createdBy": "SYSMAN",
  "createdOn": "Mon, 16 Oct 2017",
  "urn": "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:PartitionProfile:5BAF9D80FF705B3CE0535B31C40ABEFC:0.1",
  "attachments": [
    {
      "fileName": "base_domain-attributes.json",
      "size": "1.64 KB",
      "contentType": "text/plain",
      "addedBy": "SYSMAN",
      "addedOn": "Mon, 16 Oct 2017"
    }
  ],
  "source": "Oracle Home",
  "platform": null,
  "wlsVersion": null
}
```

Deleting Provisioning Profiles

To delete the Fusion Middleware Provisioning Profile from Software Library, run the following:

Usage:

```
emcli delete_fmw_profile
    -location="Profile Location"
```

Description for the options:

-location
Complete Software Library path to the Profile.
Use list_fmw_profiles to know the complete path.

Examples:

To delete a profile by the name MyProfile, run the following command:

```
emcli delete_fmw_profile
    -location="Fusion Middleware Provisioning/Profiles/MyProfile"
```

To delete a DomainProfile, run the following command:

```
emcli delete_fmware_profile
    -location="Fusion Middleware Provisioning/Profiles/DomainProfile"
```

Output: The profile Fusion Middleware Provisioning/Profiles/DomainProfile has been deleted successfully.

You can delete only one Fusion Middleware Provisioning Profile at a time.

Describing scaledown_domain usage

To remove specified managed servers from a WebLogic Domain, run the following:

Usage:

```
emcli scaledown_domain
    -domain_target="Domain Target Name"
    -server_list="Server List"
    [-host_cred="Host Credentials"]
    [-wls_cred="WebLogic Administrator Credentials"]
    [-schedule=
        start_time:yyyy/MM/dd HH:mm;
        [tz:{java timezone ID}];
        [grace_period:xxx];
    ]
```

[] indicates that the parameter is optional.

Description for the options:

- -domain_target
Name of the WebLogic domain target.
- -server_list
Comma separated list of the managed servers names to be removed.
- -host_cred
Named credential list for OS user that has write permission on the managed servers domain homes. To pass the credential parameter, enter a name:value pair in the following format: credential_name:credential_owner.
where,
 - Credential_name is the name of the named credential.
 - Credential_owner is the credentials of the Oracle home owner on the Administration Server host.

Note:

All the operations will be performed with the same credential if only one is provided, if each managed server requires a different credential, then a list of comma separated credentials will be matched against the list of managed servers. If no named credential is provided, the preferred host credentials will be used.

- -wls_cred

Named credential used to access the Administration Server. This is an optional parameter. To pass the credential parameter, enter a name:value pair in the following format: `credential_name:credential_owner`.

where,

- `Credential_name` is the name of the named credential.
- `Credential_owner` is the credentials of the Administrator of the WebLogic Domain.



Note:

If no named credential is provided, the preferred administrator credentials for the domain target will be used.

- `-schedule`

Specify when to run the deployment procedure. If no value is entered, by default, the procedure runs immediately. To schedule a procedure, provide:

- `start_time`: when the procedure should start
- `tz`: the timezone ID
- `grace_period`: grace period in minutes

Example 1

A single Managed Server named `mServer` needs to be removed from the WebLogic Domain target `/Farm01_base_domain/base_domain` at the specified schedule. Since the Managed Server host credentials haven't been specified, the preferred credentials for the target host and domain target are used.

```
emcli scaledown_domain
  -domain_target="/Farm01_base_domain/base_domain"
  -server_list="mServer"
  -schedule="start_time:2014/6/21
21:23;tz:America/New_York;grace_period:60"
```

Describing deploy_jvmd usage

To deploy JVMD agent in a WebLogic Domain, run the following:

Usage:

```
emcli deploy_jvmd
  -domain_name="Domain Target Name"
  [-engine="JVMD Engine URL"]
  [-clusters="Clusters Name"]
  [-servers="Servers Name"]
  [-host_cred="Oracle Home Owner Credentials"]
  [-wls_cred="WebLogic Administrator Credentials"]
  [-deploy_mode="Deploy Mode"]
  [-java_args="Client java arguments"]
  [-schedule=
    start_time:yyyy/MM/dd HH:mm;
    [tz:{java timezone ID}];
    [grace_period:xxx];
  ]
```



```
[ -workDir="Working Directory Location"]  
[ -libDir="Libraries Directory Location"]
```

[] indicates that the parameter is optional.

Description for the options:

- **-domain_name**
Name of the WebLogic Domain Target where the agent will be deployed.
- **-engine**
JVMD Engine URL the agent will communicate to.
- **-clusters**
Name of the cluster(s) the JVMD agent will be targeted.
- **-servers**
Name of the server(s) the JVMD agent will be targeted.
- **-host_cred**
Named credential used to access the reference host. This is an optional parameter. To pass the credential parameter, enter a name:value pair in the following format: `credential_name:credential_owner`.

where,

- `Credential_name` is the name of the named credential.
- `Credential_owner` is the credentials of the Oracle home owner on the Administration Server host.

Note:

All the operations will be performed on the Administration Server host. If no named credential is provided, the preferred host credentials for the Oracle home target will be used.

- **-wls_cred**
Named credential used to access the Administration Server. This is an optional parameter. To pass the credential parameter, enter a name:value pair in the following format: `credential_name:credential_owner`.

where,

- `Credential_name` is the name of the named credential.
- `Credential_owner` is the credentials of the Oracle home owner on the Administration Server host.

Note:

If no named credential is provided, the preferred administrator credentials for the domain target will be used.

- **-deploy_mode**

Deployment mode of the JVMD agent. This is an optional parameter. The valid modes are deploy, redeploy and undeploy. If no mode is provided, the 'deploy' mode will be used by default.

- -java_args

Client java arguments that will be used to connect to the Administration Server. This is an optional parameter.

- -schedule

Specify when to run the deployment procedure. If no value is entered, by default, the procedure runs immediately. To schedule a procedure, provide:

- start_time: when the procedure should start
- tz: the timezone ID
- grace_period: grace period in minutes

- -workDir

Specify the working directory to be used. This is an optional parameter.

- -libDir

Default directory for extracting native libraries. This is an optional parameter.

Example 1

A JVMD agent is deployed in the domain base_domain at the specified schedule and targeted to two clusters. Also, since the credentials haven't been specified, the preferred credentials for the target are used.

```
emcli deploy_jvmd
  -domain_name="/Farm01_base_domain/base_domain"
  -engine="protocol://myhost.mycompany.com:port"
  -clusters="Cluster1,Cluster2"
  -schedule="start_time:2016/6/21
21:23;tz:America/New_York;grace_period:60"
  -workDir="/u01/mytemp"
```

Middleware Profiles Using REST APIs

This chapter describes how you can use REST APIs to create, describe, list, and delete Middleware Profiles. The operations like creating/listing/describing/deleting middleware provisioning profiles, which were only possible from the Cloud Control console is now being additionally supported using a REST request/response interactions. Profiles are like templates that you can create and store in Software Library. Once a profile is created, it can be launched numerous times to provision WebLogic Domain and/or Oracle Home. The advantage of using a profile is that you can ensure that future WebLogic installations follow a standard, consistent configuration.

In particular, this chapter covers the following:

- [Creating Middleware Provisioning Profiles](#)
- [Listing Middleware Provisioning Profiles](#)
- [Describing Provisioning Profiles](#)
- [Deleting Provisioning Profiles](#)

Creating Middleware Provisioning Profiles

This section describes how you can use the REST APIs to create the three different types of provisioning profiles that are supported in Enterprise Manager Cloud Control:

- [Creating a WebLogic Domain Profile](#)
- [Creating an Oracle Home Profile](#)
- [Creating an Installation Media Profile](#)

Creating a WebLogic Domain Profile

To create a profile from a WebLogic Domain, follow these steps:

1. Perform the GET operation on the URL to view the inputs that you need to provide to create a profile.

Table 30-1 GET Request Configuration for Creating a WebLogic Domain Profile

Feature	Description
URL	<code>:/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/weblogic_domain</code>
Body	None
Parameters/Constraints	None
Request method	GET

Table 30-1 (Cont.) GET Request Configuration for Creating a WebLogic Domain Profile

Feature	Description
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  name: "Name of the profile to be created."
  targetName: "Name of the WebLogic Domain target that will be used as reference
to create the profile."
  description: "[Optional] Description of the profile that will be created."
  credential: "[Optional] Named credential that will be used to access the refer-
ence host. Format: CREDENTIAL_NAME:CREDENTIAL_OWNER. All operations will be per-
formed on the Administration Server host. Credentials of the Oracle Home owner
on the Administration Server host are required. If no named credential is provid-
ed, then preferred host credentials for the Oracle Home target will be used."
  wlsCredential: "[Optional] Named credential that will be used to access the Ad-
ministration Server. Format: CREDENTIAL_NAME:CREDENTIAL_OWNER. If no named cre-
dential is provided, then preferred administrator credentials for the domain tar-
get will be used."
  includeOh: "[Optional] Whether the Oracle Home binaries have to be included in
the profile or not. Value: true/false"
  schedule: "[Optional] The schedule for the Deployment Procedure. If not speci-
fied, the procedure will be executed immediately. Format: start_time:yyyy/MM/dd
HH:mm; [tz:{java timezone ID}]; [grace_period:xxx];
}
```

2. Update the values, and then perform a POST operation.

Table 30-2 POST Request Configuration for Creating a WebLogic Domain Profile

Feature	Description
URL	em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/weblogic_domain
Request header	
Body	{ "name": "DomainProfileViaRestAPI", "targetName": "/Farm01_SoaDomain/SoaDomain", "credential": "HOSTCRED:SYSMAN", "description": "Domain Profile", "includeOh": "false" }
Parameters	NA
Request method	POST
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  instanceName: "CreateFmwProfile-DomainProfileViaRest_SYSMAN_07_09_2014_09_47_AM"
  instanceGuid: "FDBDDDB0C6767690CE04373B1F00A09E8"
  executionGuid: "FDBDDDB0C676A690CE04373B1F00A09E8"
  executionUrl: "http://<agent host name>:7802/em/faces/core-jobs-procedureExecutionTracking?instanceGUID=FDBDDDB0C6767690CE04373B1F00A09E8&showProcAct-Link=yes&executionGUID=FDBDDDB0C676A690CE04373B1F00A09E8"
  name: "DomainProfileViaRestAPI"
  targetName: "/Farm01_SoaDomain/SoaDomain"
  description: "Domain Profile"
  credential: "HOSTCRED:SYSMAN"
  include0h: "false"
}
```

Creating an Oracle Home Profile

To create a provisioning profile from an Oracle Home, follow these steps:

1. Perform the GET operation on the URL to view the inputs that you need to provide to create a profile.

Table 30-3 GET Request Configuration for Creating an Oracle Home Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/oracle_home
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  name: "Name of the profile to be created."
  targetName: "Name of the Oracle Home target that will be used as reference to create the profile."
  description: "[Optional] Description of the profile that will be created."
  credential: "[Optional] Named credential that will be used to access the reference host. Format: CREDENTIAL_NAME:CREDENTIAL_OWNER. If no named credential is provided, then preferred host credentials for the Oracle Home target will be used."
  schedule: "[Optional] The schedule for the Deployment Procedure. If not specified, the procedure will be executed immediately. Format: start_time:yyyy/MM/dd HH:mm; [tz:{java timezone ID}]; [grace_period:xxx];"
}
```

2. Update the values, and then perform a POST operation.

Table 30-4 POST Request Configuration for Creating an Oracle Home Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/oracle_home
Body	<pre>{ "name": "OhProfileViaRestAPI", "targetName": "WebLogicServer10_3_6_0_slc00tkv.example.com_6316", "credential": "HOSTCRED:SYSMAN", "description": "OH Profile" }</pre>
Parameters	NA
Request method	POST
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  instanceName: "CreateFmwProfile-OhProfileViaRestAPI_SYSMAN_07_09_2014_09_54_AM"
  instanceGuid: "FDC6776FB9B31CCDE04373B1F00A2B1C"
  executionGuid: "FDC6776FB9B61CCDE04373B1F00A2B1C"
  executionUrl: "http://<agent host name>:7802/em/faces/core-jobs-procedureExecutionTracking?instanceGUID=FDC6776FB9B31CCDE04373B1F00A2B1C&showProcAct-Link=yes&executionGUID=FDC6776FB9B61CCDE04373B1F00A2B1C"
  name: "OhProfileViaRestAPI"
  targetName: "WebLogicServer10_3_6_0_slc00tkv.example.com_6316"
  description: "OH Profile"
  credential: "HOSTCRED:SYSMAN"
}
```

Creating an Installation Media Profile

To create a provisioning profile from an Installation Media, follow these steps:

1. Perform the GET operation on the URL to view the inputs that you need to provide to create a profile.

Table 30-5 GET Request Configuration for Creating an Installation Media Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/install_media
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  name: "Name of the profile to be created."
  targetName: "Name of the Host target that where all the installation files are
  stored."
  description: "[Optional] Description of the profile that will be created."
  credential: "[Optional] Named credential that will be used to access the files.
  Format: CREDENTIAL_NAME:CREDENTIAL_OWNER. If no named credential is provided,
  then normal preferred credentials for the Host target will be used."
  platform: "Platform for which the installation media is applicable."
  version: "Version of the installation media."
  -files: [2]
  -0:
  {
    product: "Product1"
    -files: [1]
    0: "file1"
  }
  -1:
  {
    product: "Product2"
    -files: [2]
    0: "file2"
    1: "file3"
  }
}
```

2. Update the values, and then perform a POST operation.

Table 30-6 POST Request Configuration for Creating an Installation Media Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/create/install_media

Table 30-6 (Cont.) POST Request Configuration for Creating an Installation Media Profile

Feature	Description
Body	<pre>{ "name": "ImProfileViaRestAPI", "targetName": "<agent host name>.example.com", "platform": "2000", "version": "11.1.1.8.0", "credential": "HOSTCRED:SYSMAN", "description": "IM Profile", "files": [{ "product": "WebLogic", "files": ["/net/adcnas438/export/emgpqa/provisioning/fmwprov/ linux64/shiphomes/upinsmed_wls/wls1036_generic.jar"] }, { "product": "WCC", "files": ["/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/ fmwprov/linux64/shiphomes/wcc_11.1.1.8.0/ecm_main1.zip", "/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/ fmwprov/linux64/shiphomes/wcc_11.1.1.8.0/ecm_main2.zip"] }, { "product": "WCP", "files": ["/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/ fmwprov/linux64/shiphomes/wcp_11.1.1.8.0/wc.zip"] }] }</pre>
Parameters	NA
Request method	POST
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  instanceName: "CreateFmwProfile-ImProfileViaRestAPI_SYSMAN_07_09_2014_10_07_AM"
  instanceGuid: "FDC6BDD54DA70346E04373B1F00A2DBE"
  executionGuid: "FDC6BDD54DAA0346E04373B1F00A2DBE"
  executionUrl: "http://slc03qtn:7802/em/faces/core-jobs-procedureExecutionTrack-
ing?instanceGUID=FDC6BDD54DA70346E04373B1F00A2DBE&showProcActLink=yes&execution-
GUID=FDC6BDD54DAA0346E04373B1F00A2DBE"
  name: "ImProfileViaRestAPI"
  targetName: "slc03qtn.example.com"
  description: "IM Profile"
```



```
credential: "HOSTCRED:SYSMAN"
platform: "2000"
version: "11.1.1.8.0"
-files: [3]
-0:
{
product: "WebLogic"
-files: [1]
0: "/net/adcnas438/export/emgpqa/provisioning/fmwprov/linux64/shiphomes/
upinsmed_wls/wls1036_generic.jar"
}
-1:
{
product: "WCC"
-files: [2]
0: "/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/
shiphomes/wcc_11.1.1.8.0/ecm_main1.zip"
1: "/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/
shiphomes/wcc_11.1.1.8.0/ecm_main2.zip"
}
-2:
{
product: "WCP"
-files: [1]
0: "/net/slcnas478/export/farm_em_repos/emgpqa/provisioning/fmwprov/linux64/
shiphomes/wcp_11.1.1.8.0/wc.zip"
}
}
```

Listing Middleware Provisioning Profiles

This section describes how you can list all the profiles that you have created in Enterprise Manager. The summary includes a basic link for each profile. This link can be used to view the summary of the profile with the *Describe* API. For more information on this, see [Describing Provisioning Profiles](#).

In particular, this section includes the following topics:

- [Listing All the Profiles](#)
- [Listing WebLogic Domain Profiles](#)
- [Listing Oracle Home Profile](#)
- [Listing Installation Media Profiles](#)

Listing All the Profiles

Perform the GET operation on the URL to get the canonical link, which can be later used to describe the profile.

Table 30-7 GET Request Configuration for Listing All the Profiles

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/list
Body	NA

Table 30-7 (Cont.) GET Request Configuration for Listing All the Profiles

Feature	Description
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  -profiles: [3]
  -0:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
    nent:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1"
    location: "Fusion Middleware Provisioning/Profiles/SoaBitlessProfile"
    products: "Oracle SOA"
    source: "WebLogic Domain"
    urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBun-
    dle:FD8AEF47A15C0874E04373B1F00AEC31:0.1"
  }
  -1:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
    nent:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1"
    location: "Fusion Middleware Provisioning/Profiles/SoaGoldImage"
    products: "Oracle SOA"
    source: "Oracle Home"
    urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBun-
    dle:FDB228F176CA1B47E04373B1F00AF708:0.1"
  }
  -2:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
    nent:SUB_InstallationMedia:FD8AA7EEEE369A5E04373B1F00AE77C:0.1"
    description: "WC 11.1.1.8.0 and WLS 10.3.6.0 installer and RCU"
    location: "Fusion Middleware Provisioning/Profiles/WebCenter Install Media"
    -products: [2]
    0: "Oracle WebCenter Content"
    1: "Oracle WebCenter Portal"
    source: "Installation Media"
    urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMe-
    dia:FD8AA7EEEE369A5E04373B1F00AE77C:0.1"
  }
}
```

Listing WebLogic Domain Profiles

Perform the GET operation on the URL to get the canonical link, which can be later used to describe the profile.

Table 30-8 GET Request Configuration for Listing the WebLogic Domain Profiles

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/list/weblogic_domain
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  -profiles:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
    nent:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1"
    location: "Fusion Middleware Provisioning/Profiles/SoaBitlessProfile"
    products: "Oracle SOA"
    source: "WebLogic Domain"
    urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBun-
    dle:FD8AEF47A15C0874E04373B1F00AEC31:0.1"
  }
}
```

Listing Oracle Home Profile

Perform the GET operation on the URL to get the canonical link, which can be later used to describe the profile.

Table 30-9 GET Request Configuration for Listing the Oracle Home Profiles

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/list/oracle_home
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  -profiles:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
```

```

nent:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1"
location: "Fusion Middleware Provisioning/Profiles/SoaGoldImage"
products: "Oracle SOA"
source: "Oracle Home"
urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBun-
dle:FDB228F176CA1B47E04373B1F00AF708:0.1"
}
}

```

Listing Installation Media Profiles

Perform the GET operation on the URL to get the canonical link, which can be later used to describe the profile.

Table 30-10 GET Request Configuration for Listing the Installation Media Profiles

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/list/install_media
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```

{
  -profiles:
  {
    canonicalLink: "http://slc03qtn:7802/em/websvcs/restful/extws/cloudservices/fmw/
    provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
    nent:SUB_InstallationMedia:FD8AA7EEEE369A5E04373B1F00AE77C:0.1"
    description: "WC 11.1.1.8.0 and WLS 10.3.6.0 installer and RCU"
    location: "Fusion Middleware Provisioning/Profiles/WebCenter Install Media"
    -products: [2]
    0: "Oracle WebCenter Content"
    1: "Oracle WebCenter Portal"
    source: "Installation Media"
    urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMe-
    dia:FD8AA7EEEE369A5E04373B1F00AE77C:0.1"
  }
}

```

Describing Provisioning Profiles

This section describes how you can view a summary of all the profiles that you have created in Enterprise Manager. For this, you'll need to run the List API and procure the canonical link. For more information on this, see [Listing Middleware Provisioning Profiles](#).

In particular, this section covers the following topics:

- [Describing a WebLogic Domain Profile](#)
- [Describing an Oracle Home Profile](#)
- [Describing an Installation Media Profile](#)

Describing a WebLogic Domain Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-11 GET Request Configuration for Describing the WebLogic Domain Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```
{
  location: "Fusion Middleware Provisioning/Profiles/SoaBitlessProfile"
  createdBy: "SYSMAN"
  createdOn: "Sun, 6 Jul 2014"
  urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1"
  -attachments: [3]
  -0:
  {
    fileName: "nmMovePlan.xml"
    size: "6.95 KB"
    contentType: "text/plain"
    addedBy: "SYSMAN"
    addedOn: "Sun, 6 Jul 2014"
  }
  -1:
  {
    fileName: "mdm-url-resolver.xml"
    size: "268 Bytes"
    contentType: "text/plain"
    addedBy: "SYSMAN"
    addedOn: "Sun, 6 Jul 2014"
  }
  -2:
  {
    fileName: "domainMovePlan.xml"
    size: "56.91 KB"
```

```
contentType: "text/plain"
addedBy: "SYSMAN"
addedOn: "Sun, 6 Jul 2014"
}
source: "WebLogic Domain"
oracleHomeIncluded: "No"
nodeManagerIncluded: "Yes"
platform: "Linux x86-64"
wlsVersion: "10.3.6.0"
-oracleHome:
{
-products: [1]
-0:
{
name: "Oracle SOA"
version: "11.1.1.7.0"
}
}
-domain:
{
name: "SoaDomain"
size: "611.45 MB"
-products: [1]
0: "Oracle SOA"
}
-servers: [2]
-0:
{
name: "AdminServer"
host: "<agent host name>.example.com"
listenAddress: "<agent host name>.example.com"
nonSslPortEnabled: "Yes"
nonSslPort: "7001"
sslPortEnabled: "No"
machine: "LocalMachine"
maxHeap: "1365"
averageCpuUsage(%): "0.361"
}
-1:
{
name: "soa_server1"
host: "slc00tkv.example.com"
listenAddress: "<agent host name>.example.com"
nonSslPortEnabled: "Yes"
nonSslPort: "8001"
sslPortEnabled: "No"
machine: "LocalMachine"
cluster: "cluster_soa"
maxHeap: "1365"
averageCpuUsage(%): "0.717"
}
-dataSources: [7]
-0:
{
name: "EDNDataSource"
schema: "SOA_SOAINFRA"
url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
-targets: [1]
0: "cluster_soa"
}
-1:
```

```

{
  name: "EDNLocalTxDataSource"
  schema: "SOA_SOAINFRA"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [1]
  0: "cluster_soa"
}
-2:
{
  name: "OraSDPMDDataSource"
  schema: "SOA_ORASDPM"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [1]
  0: "cluster_soa"
}
-3:
{
  name: "SOADDataSource"
  schema: "SOA_SOAINFRA"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [1]
  0: "cluster_soa"
}
-4:
{
  name: "SOALocalTxDataSource"
  schema: "SOA_SOAINFRA"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [1]
  0: "cluster_soa"
}
-5:
{
  name: "mds-owsm"
  schema: "SOA_MDS"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [2]
  0: "cluster_soa"
  1: "AdminServer"
}
-6:
{
  name: "mds-soa"
  schema: "SOA_MDS"
  url: "jdbc:oracle:thin:@<agent host name>.example.com:1521/dbv.example.com"
  targets: [2]
  0: "cluster_soa"
  1: "AdminServer"
}
-jmsServers: [4]
-0:
{
  name: "BPMJMSServer_auto_1"
  persistentStore: "BPMJMSFileStore_auto_1"
  storeType: "fileStore"
  directory: "BPMJMSFileStore_auto_1"
  target: "soa_server1"
}
-1:
{
  name: "PS6SOAJMSServer_auto_1"

```

```

persistentStore: "PS6SOAJMSFileStore_auto_1"
storeType: "fileStore"
directory: "PS6SOAJMSFileStore_auto_1"
target: "soa_server1"
}
-2:
{
name: "SOAJMSServer_auto_1"
persistentStore: "SOAJMSFileStore_auto_1"
storeType: "fileStore"
directory: "SOAJMSFileStore_auto_1"
target: "soa_server1"
}
-3:
{
name: "UMSJMSServer_auto_1"
persistentStore: "UMSJMSFileStore_auto_1"
storeType: "fileStore"
directory: "UMSJMSFileStore_auto_1"
target: "soa_server1"
}
}
}

```

Describing an Oracle Home Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-12 GET Request Configuration for Describing the Oracle Home Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```

{
location: "Fusion Middleware Provisioning/Profiles/SoaGoldImage"
createdBy: "SYSMAN"
createdOn: "Tue, 8 Jul 2014"
urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1"
source: "Oracle Home"
platform: "Linux x86-64"
wlsVersion: "10.3.6.0"
-oracleHome:

```



```

{
  size: "3.51 GB"
  -products: [1]
  -0:
  {
    name: "Oracle SOA"
    version: "11.1.1.7.0"
  }
}

```

Describing an Installation Media Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-13 GET Request Configuration for Describing the Installation Media Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMedia:FD8AA7EEEEA369A5E04373B1F00AE77C:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	GET
Supported Since Release	FMW 12.1.0.7

Example Response:

```

{
  location: "Fusion Middleware Provisioning/Profiles/WebCenter Install Media"
  description: "WC 11.1.1.8.0 and WebLogic Server 10.3.6.0 installer and RCU"
  createdBy: "SYSMAN"
  createdOn: "Sun, 6 Jul 2014"
  urn: "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMedia:FD8AA7EEEEA369A5E04373B1F00AE77C:0.1"
  source: "Installation Media"
  product: "Oracle WebCenter"
  platform: "Generic Platform"
  version: "11.1.1.8.0"
  -files: [4]
  -0:
  {
    product: "Oracle WebCenter Content"
    size: "2.34 GB"
    -files: [2]
    -0:
    {
      fileName: "ecm_main1.zip"
      size: "1.91 GB"
    }
  }
}

```

```
-1:
{
  fileName: "ecm_main2.zip"
  size: "437.6 MB"
}
}
-1:
{
  product: "Oracle RCU"
  size: "496.16 MB"
  -files: [1]
  -0:
  {
    fileName: "rcuHome.zip"
    size: "496.16 MB"
  }
}
-2:
{
  product: "Oracle WebCenter Portal"
  size: "1.96 GB"
  -files: [1]
  -0:
  {
    fileName: "wc.zip"
    size: "1.96 GB"
  }
}
-3:
{
  product: "Oracle WebLogic Server"
  size: "1,019.01 MB"
  -files: [1]
  -0:
  {
    fileName: "wls1036_generic.jar"
    size: "1,019.01 MB"
  }
}
}
```

Deleting Provisioning Profiles

This section describes how you can delete the profiles you have created in Enterprise Manager.

In particular, this section covers the following topics:

- [Deleting a WebLogic Domain Profile](#)
- [Deleting an Oracle Home Profile](#)
- [Deleting an Installation Media Profile](#)

Deleting a WebLogic Domain Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-14 DELETE Request Configuration for Deleting the WebLogic Domain Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FD8AEF47A15C0874E04373B1F00AEC31:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	DELETE
Supported Since Release	FMW 12.1.0.7

Example Response:

After the deletion is successful, you might see something on the lines of the following message:

Profile Fusion Middleware Provisioning/Profiles/MyProfile deleted successfully.

Deleting an Oracle Home Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-15 DELETE Request Configuration for Deleting the Oracle Home Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_FMWBundle:FDB228F176CA1B47E04373B1F00AF708:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	DELETE
Supported Since Release	FMW 12.1.0.7

Example Response:

After the deletion is successful, you might see something on the lines of the following message:

Profile Fusion Middleware Provisioning/Profiles/MyProfile deleted successfully.

Deleting an Installation Media Profile

Use the canonical link obtained from the LIST API in the GET operation to view the summary of the selected profile.

Table 30-16 DELETE Request Configuration for Deleting the Installation Media Profile

Feature	Description
URL	/em/websvcs/restful/extws/cloudservices/fmw/provisioning/profile/describe/oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_InstallationMedia:FD8AA7EEEEA369A5E04373B1F00AE77C:0.1
Request header	Is authorization required??
Body	NA
Parameters/Constraints	NA
Request method	DELETE
Supported Since Release	FMW 12.1.0.7

Example Response:

After the deletion is successful, you might see something on the lines of the following message:

```
Profile Fusion Middleware Provisioning/Profiles/MyProfile deleted successfully.
```

Scaling Up / Scaling Out Fusion Middleware Domains

This chapter explains how you can scale up and scale out a SOA Domain, an Service Bus Domain, a WebLogic Domain, and a WebCenter Domain using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started](#)
- [Prerequisites](#)
- [Running the Scale Up / Scale Out Middleware Deployment Procedure](#)
- [Middleware Provisioning and Scale Up / Scale Out Best Practices](#)

 Note:

To scaleup multiple products in one session, you can clone the servers from each products's cluster simultaneously.

Getting Started

A WebLogic Domain consists of a set of managed servers running independently or in a cluster, sharing the distributed resources. A WebLogic Server cluster consists of multiple WebLogic managed servers running simultaneously and working together to provide increased scalability and reliability. The server instances that constitute a cluster can run on the same machine, or be located on different machines. You can increase a cluster's capacity by adding additional server instances to the cluster on an existing machine, or by adding machines to the cluster to host the new server instances. You can use the Domain Scale Up / Scale Out deployment procedure to automate the scaling up or scaling out of a domain. You can:

- Scale up a domain by adding or cloning a managed server to a host that already exists in the domain or cluster.
- Scale out a domain by adding or cloning a managed server to a host that is not present in the domain or cluster.

Prerequisites

Before running the Scale Up / Scale Out Middleware deployment procedure, you must meet the prerequisites listed in this section.

**Note:**

For information about how to setup your infrastructure for Middleware Provisioning, see [Prerequisites for Provisioning from the Middleware Provisioning Profiles](#).

Meet the following prerequisites before you start extending the WebLogic Domain:

- The WebLogic Domain that is scaled up / scaled out must be an existing domain that has been discovered with Cloud Control.
- If you are scaling out a domain, ensure that the destination machine contains sufficient space. If the size of the 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 Middleware Home. The working directory is cleaned up after deployment procedure has been successfully completed.
- The Middleware Home directory you specify on the destination machine must be a new directory or must be empty.
- The Management Agent must be installed on the source (where the Administration Server is running) and the destination machines. The Administration Server for the domain must be up and running.
- The Administration Server and Managed Server (being cloned) must be up and running before you run the deployment procedure.
- The Managed Server and Node Manager ports must be free.
- Ensure that the user has the necessary permission to the stage directory inside the work directory. The scale up may fail with an error in case sufficient permissions are not granted to the user. In such an event, assign the necessary **700** permission to the stage directory and retry.
- For scaling out a domain, the user must have the following permissions:
 - Read permissions on:
 - * Administration Server Host Middleware Directory
 - * Administration Server Host Domain Directory
 - Write permissions on:
 - * Administration Server Host Working Directory
 - * Working Directory of all the destination Managed Server hosts
 - * Middleware Directory of all the destination Managed Server hosts
 - * Domain Directory of all the destination Managed Server hosts
- For scaling up a domain, the user must have the following permissions:
 - Read permissions on:
 - * Administration Server Host Working Directory
 - * Domain Directory of all the destination Managed Server hosts
- The domain being scaled up / out should not be in Edit mode. Ensure that there is a running WebLogic Console for this domain.

- If you choose to associate a new managed server with an existing Node Manager or a machine, ensure that the Node Manager is up and running. If not, the deployment procedure will fail.
- Ensure that you have discovered an existing OHS target in Enterprise Manager, if you want to front end the Oracle HTTP Server.
- Ensure that the target machine which is being scaled up and Enterprise Manager host where the OMS is running are in the same timezone, before you begin the scaleup process.
- Ensure that you do not delete the Aggregation Server when you are scaling down a domain. If the Aggregation Server is deleted, then a lot of dependent applications will stop running.

Running the Scale Up / Scale Out Middleware Deployment Procedure

A WebLogic Domain consists of a set of managed servers running independently or in a cluster, sharing the distributed resources. A WebLogic Server cluster consists of multiple WebLogic managed servers running simultaneously and working together to provide increased scalability and reliability. The server instances that constitute a cluster can run on the same machine, or be located on different machines. You can increase a cluster's capacity by adding additional server instances to the cluster on existing machines, or by adding machines to the cluster to host the new server instances.

The Scale Up / Scale Out Middleware wizard allows you to increase a cluster's capacity by adding additional server instances to the cluster on an existing machine, or by adding machines to the cluster to host the new server instances.

 Note:

Java Object Cache is always configured on a scaled out or scaled up middleware managed server whether it was available before scale up or not.

To scale up or scale out a domain, follow these steps:

1. Ensure that the prerequisites are met. See [Prerequisites](#).
2. Specify the source domain. See [WebLogic Domain Scaling Up: Scale up/Scale out Fusion Middleware Domain Page](#).
3. Specify when the scale up or scale out operation should be performed. [Weblogic Domain Scaling Up / Scaling Out : Schedule Page](#).



Note:

For scaling out a WebCenter Domain, you must perform the following steps manually:

1. Oracle HTTP Server must be installed, discovered, monitored in Cloud Control. Additionally, you must ensure that for the scale out instance the configuration file should be manually configured.
2. If the Spaces Server is scaled out, then you must run the following commands to attach the WebService Policy:

```
attachWebServicePolicy(application='WC_Spaces2/webcenter',
moduleName='webcenter', moduleType='web',
serviceName='SpacesWebService',
subjectName='SpacesWebServiceSoapHttpPort',
policyURI='oracle/wss11_saml_token_with_message_protection_serv-
ice_policy')
```

3. If the Discussion Server is scaled out, then you must run the following commands to attach the WebService Policy:

```
attachWebServicePolicy(application='WC_Collaboration2/owc_discussions',
moduleName='owc_discussions', moduleType='web',
serviceName='OWCDiscussionsServiceAuthenticated',
subjectName='OWCDiscussionsServiceAuthenticated',
policyURI='oracle/wss10_saml_token_service_policy')
```

WebLogic Domain Scaling Up: Scale up/Scale out Fusion Middleware Domain Page

You can automate the scaling up or scaling out of a domain or cluster using the Domain Scale Up / Out Deployment Procedure. You can add capacity to an existing WebLogic Domain and /or cluster by:

- Adding attributes of a new managed server to an existing cluster.
- Adding and copying attributes of a new managed server to an existing cluster.
- Cloning a managed server. If the source server is clustered, when you clone an existing Managed Server, another Managed Server will be created in the same cluster.

A wizard guides you through the process.

1. From the Targets menu, select **Middleware**.
2. A list of Middleware targets is displayed. Find the WebLogic Domain that you want to use as the source for the scaling up or out operation. Right click on that WebLogic Domain to access the context sensitive menu. From the menu, select **Provisioning** and then select **Scale Up / Scale Out WebLogic Domain**.

Alternatively, you can click WebLogic Domain link. On the domain home page, from the WebLogic Domain menu, select **Provisioning**, and click **Scale Up/Scale Out WebLogic Domain**.

The Scale up/Scale out Fusion Middleware Domain page is displayed.

You can also launch this as a procedure from Middleware Provisioning page. To do so, from **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware provisioning**. On the Middleware Provisioning page, from the deployment procedures table, select **Scale up/Scale out Middleware**, and click **Launch**. In this case you will need to provide the source information by selecting the WebLogic domain that needs to be extended.

The screenshot shows the 'Scale up/Scale out Fusion Middleware Domain: Configuration' page in Oracle Enterprise Manager Cloud Control 13c. The page is divided into several sections:

- WebLogic Domain:** Contains fields for 'Domain Target Name' (set to 'myfam_base_domainbase_domain'), 'Oracle Home Credential' (set to 'Select a credential'), 'WebLogic Administrator Credential' (set to 'Select a credential'), and 'Working Directory' (set to 'Agent's Working Directory <emidstageDir>OP Instance</emidstageDir>').
- Hosts:** A section to add destination hosts where new servers will be provisioned. It includes a table with columns 'Host' and 'Credential'.
- Machines:** A section to add destination hosts and then create new machines if required. It includes a table with columns 'Name', 'Host', 'Node Manager Listen Address', and 'Node Manager Listen Port'.
- Servers:** A section to add the destination hosts before adding new servers. It includes a table with columns 'Name', 'Host', 'Listen Address', 'Listen Port', 'Machine', 'Cluster', and 'Server Start'.

3. In the **Weblogic Domain** section, enter the Working Directory field details specifying the directory on the Administration Server machine where the domain scale up related files are temporarily stored. A minimum of one GB of directory space is required to store the temporary files. If this directory is not present, it will be created. When the scale up operation has been completed, the directory and its contents will be deleted.

Host Credentials: In this section, you must provide all the credentials for all the hosts on which the Managed Servers are running. To do so, click the add icon for each host, and in the Add New Credentials dialog box, enter a valid username and password. If you want the job to use these credentials for each target when the job runs, select the **Set As Preferred Credentials** check box.

Deploy JVMD agent: Deploys the JVMD agent on the servers or clusters of a Weblogic Domain. Java Virtual Machine Diagnostics (JVMD) enables administrators to diagnose performance problems in Java applications. After the JVMD Agent is deployed on the targeted JVM (the one running a production WebLogic Server), it collects real-time data and transmits it to the JVM Diagnostics Engine.

WebLogic Administrator Credentials: In this section, you must set the credentials to access WebLogic administration server. To do so, click the **Add** for each of the WebLogic Domains, and in the Add new administrator credentials dialog box, enter a valid administrator name and password. If you want the job to use these credentials for each target when the job runs, then select the **Set As Preferred Credentials** check box.

Note:

The Working Directory must not be created under the Middleware Home or the WebLogic Domain Home directory.

4. In the **Hosts** section, click **Add** to select the host you want to add managed servers to.
5. In the **Machines** section, you can **Add** the destination hosts and then create new machines if required.

 Note:

To perform administrative operations such as start and stop from the Enterprise Manager Cloud Control, the Node Manager must be configured and running on the machine.

6. In the **Servers** section, you can add a new managed server to a WebLogic Domain.

To add a new managed server to a WebLogic Domain, click **Add**. Enter configuration details for a new Managed Server like a unique name for the new Managed Server, listen address, and the SSL port.

 Note:

The **Add** button is only available for Clusterless domains. If the domain has clusters, then cloning is the preferred option.

If the domain has the administration port enabled, **Administration Port** column will appear in the servers table so that the user can enter the administration port for the new server. If the domain does not have the administration port enabled, this column will not be visible in the servers table. When a new server is added/cloned, automatically a new port is set. You must ensure that this new port is free, if not the deployment procedure will fail.

If you select **Custom Identity and Custom Trust**, you must provide valid inputs for all the KeyStore fields.

Cloning an Existing Managed Server: To do so, select the host you want to add the new server to and click Clone. This option adds a new managed server to the domain, and copies a pre-determined set of the attributes from the existing server.

Deleting a Managed Server: To do so, select the managed server and click Delete Server.

7. In the **Custom Scripts and Variables** section, you can add any custom properties that you want to pass to your scripts.
8. Click **Next**.

 Note:

If you are scaling a domain with products like SOA or OSB, then a **JMS Servers** section is displayed if the source domain has it configured.

Weblogic Domain Scaling Up / Scaling Out : Schedule Page

On the Schedule page, specify a Deployment Instance name. If you want to run the procedure immediately, retain the default selection, that is, **Immediately**.

If you want to run the procedure later, select **Later** and provide time zone, start date, and start time details.

You can set the notification preferences according to deployment procedure status.

If you want to run only prerequisites, select **Pause the procedure to allow me to analyze results after performing prerequisite checks**. This pauses the procedure execution after all prerequisite checks are performed.

Click **Next**.

Middleware Provisioning and Scale Up / Scale Out Best Practices

This section lists some of the best practices to be followed while using the Middleware Provisioning deployment procedure.

- **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, the working directory on the source and target machines is populated by default. If these values are changed, you must ensure that the working directory on the source and the destination machines are different. For example, if the working directory is /tmp/source for the source machine, it could be /tmp/dest on the destination directory. You must also ensure that the listen port number and SSL port numbers (if enabled) for the Administration Server and Managed Server are different on the source and destination servers.
- **JDBC Configuration:** While configuring the JDBC data sources, the database user and schema owner must enter appropriate passwords.
- **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 manual 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.
- **Port Conflicts:** While adding a new server on a host, you must enter ports that are unique to that server and are not used by any other existing managed server or process on that host.

Export and Import of Domain Partitions

This chapter provides an overview of exporting and importing Oracle WebLogic Domain Partitions. Oracle WebLogic Server supports WebLogic domain partition functionality. Partition functionality is introduced and supported through WebLogic Server Multitenancy feature. It consolidates applications into fewer domains with partition within domains and maintains secured isolation. It provides increased efficiency and reduced cost. For more information on WebLogic Server Multitenancy, see *WebLogic Server Multitenant* in the *Oracle® Fusion Middleware Using WebLogic Server Multitenant* guide. The WebLogic domain partition can be exported and imported between WebLogic Server domains. Export or import is also supported between domains residing on different Enterprise Manager deployments. In particular, this chapter covers the following:

- [About Exporting and Importing WebLogic Domain Partitions](#)
- [Exporting WebLogic Domain Partition](#)
- [Importing WebLogic Domain Partition](#)

About Exporting and Importing WebLogic Domain Partitions

A domain partition (partition for short) is an administrative and runtime slice of a WebLogic domain. Multi-Tenancy provides shared infrastructure for use by multiple organizations. These organizations are a conceptual grouping of tenants. Multiple tenants are supported in a single domain. It improves density and achieves more efficient use of resources and eliminates the following hurdles involved in sharing multiple applications:

1. Runtime Cross-Application impact
2. Security differences
3. Data Co-Mingling
4. Administrative challenges

A domain partition is an administrative portion of a domain that can be managed independently and can share runtime capacity in a domain, that is, the managed servers and clusters. Each domain partition has its own runtime copy of the applications and resources. It provides application isolation and efficiency. Multiple instances of an application can run in different domain partitions without modifying the application.

Domain partitions use fewer server and domain resources. This enables you to simplify the management of Software As a Service (SaaS) and Platform As a Service (PaaS) applications.

Oracle WebLogic Server supports WebLogic domain partition functionality. A single WebLogic domain can contain several partitions, each capable of having its own set of services like JDBC, JMS, SAF, applications, and security realm. Oracle WebLogic Multitenant Server provides resource isolation within a domain partition which is an administrative and runtime slice of a WebLogic domain. It runs application instances and related resources for a tenant. Domain partitions achieve greater density by allowing

application instances and related resources to share the domain, WebLogic Server, Java virtual machine, and operating system While isolating Tenant-Specific application data, configuration, and runtime traffic.

A partition can contain one or more resource groups. These resource groups can be converted to templates. A domain partition can be exported and imported into a new or same domain. This export or import functionality moves the services and applications that are part of the partition. This functionality is primarily used in Enterprise Manager to provision partitions on demand. The key capabilities and limitations of the WebLogic domain partition functionality are:

- Partition export or import is supported.
- An empty partition cannot be created.
- One partition can be exported at a time from a domain.
- One partition can be imported into one domain at a time.
- The applications and services running in the domain are cloned along with the partition. They can neither be deleted nor can new ones be added during provisioning.

Related Topics

- [Exporting WebLogic Domain Partition](#)
- [Importing WebLogic Domain Partition](#)

Exporting WebLogic Domain Partition



Note:

Exporting WebLogic domain partition functionality is supported from WLS version 12.2.1.0.0 onwards.

To export a domain partition, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. Expand the **Target Type** and check the option against **WebLogic** to view a list of WebLogic targets.
3. Right click on the WebLogic Domain target, and from the **Provisioning** menu, click **Create Provisioning Profile** and select **By Exporting WebLogic Domain Partition** option and click **Next**.
4. Enter a profile name in the **Name** field under the **Profile** tab.



Note:

As part of the export process, the Enterprise Manager creates a profile of the partition and stores in the software library. The profile contains exported partition archive and some metadata that helps the Enterprise Manager in the importing process.

5. Enter the description of the profile created in the **Description** field.
6. In the **WebLogic Domain Partition** field select the partition.
7. Verify the values in the fields **Domain Home** and **Oracle Home**.
8. Select the values for the remaining fields **Oracle Home Credential**, **WebLogic Administrator Credential**, and **Working Directory**.

**Note:**

The temporary staging directory of the agent is used by default, to copy the script and generate archives. In case this directory does not have enough permissions or space, the user can specify another directory.

9. In the Software Library Upload Location, select the Software Library storage details. Ensure to provide a valid storage type and upload location to update the profile.
10. Click **Create**.

**Note:**

Exporting WebLogic Domain Partition operation can be performed using EMCLI verb `create_partition_profile`. See `create_partition_profile` in the *Oracle® Enterprise Manager Command Line Interface*.

You can view the profile created in the **Middleware Provisioning** page. The new profile will be shown in the **Profiles** table, select the name of the profile to view more details.

Related Topics

- [About Exporting and Importing WebLogic Domain Partitions](#)
- [Importing WebLogic Domain Partition](#)

Importing WebLogic Domain Partition

**Note:**

Importing WebLogic domain partition functionality is supported from WLS version 12.2.1.0.0 onwards.

To import a domain partition into an existing WebLogic domain, follow these steps:

1. From the **Targets** menu, select **Middleware**.
2. Expand the **Target Type** and check the option against **WebLogic** to view a list of WebLogic targets.

3. To import the WebLogic Domain Partition, right click on the domain where you want to import the partition.
4. From the **Provisioning** menu, select **Import WebLogic Domain Partition**.
You can also import WebLogic Domain Partition from the **Middleware Provisioning** page. From the **Enterprise** menu, select **Provisioning and Patching** and click **Middleware Provisioning**. Select a profile under the **Profiles** table and click **Import Partition** button. The profile must have been created by either exporting a domain or by exporting a partition
5. Under the **General** tab, select a profile under the **WebLogic Domain Partition Profile** field.
6. Click on the **View** icon to view the profile details.
7. Upon selecting a profile, the values in the following fields are prefilled:
 - **WebLogic Domain Partition**
 - **Virtual Targets**
 - **Database**

 Note:

The Database field is filled only if the source partition has data sources.

8. The details of the profile can be viewed by clicking on the view icon available next to the **WebLogic Domain Partition Profile** field. The details include information on Data Sources, JMS Servers, Applications, and Virtual Targets.

If you like to make an advanced configuration, click **Use Custom JSON file** and enter the location of custom well-formed JSON file in the **JSON File Location** field.

 Note:

You can download the JSON file from **Profile Viewer**. Click the View icon next to the profile search icon and download it from the attachments table. You can then modify it and copy it to the Administration Server host of the domain and provide that location in the **JSON File Location** field.

9. Enter the credentials in fields **Oracle Home Credential** and **WebLogic Administrator Credential** by using the search icon against the fields.
10. Enter a unique name in the **Partition Name** field.
11. Under the **Virtual Targets** tab, in the **Targets** field choose the virtual targets from the drop down menu. The host name specified must have DNS entries that resolve to the correct server or load balancer. If the host name is not specified then it is equivalent of wildcarding the host name to match all incoming requests.

**Note:**

Two virtual targets cannot be mapped to the same target in a partition.

12. Choose a database under the **Database** tab against the **Connection String** field.

**Note:**

This is optional and has to be executed only if the source partition has data source. Ensure that the schema is already created, the database services are up and running and the database service has the same name mentioned in the data source.

While selecting the target in the Select target window the following database types are available:

- Cluster Database
- Database Instance
- Pluggable Database

13. Enter a password in the **Schema Password** field.

14. Click **Next** and confirm the changes.

Importing WebLogic Domain Partition operation can be performed using EMCLI verb too. Use the following steps.

1. The generic submit_procedure EMCLI verb is used to import WebLogic Domain Partition. Launch the import partition user interface from the desired domain target.
2. Enter all the required details.
3. Navigate to the next page.
4. Click **Save**.
5. Enter the unique SAVED_CONFIG_NAME.
6. Execute the following command:

```
emcli describe_procedure_input -name=SAVED_CONFIG_NAME > input.properties
```

7. Un-comment the secret variables like *passwords* and provide values for them.
8. Submit the procedure by executing the following command:

```
emcli submit_procedure -name=ProvisionPartition -input_file=data:input.properties [-instance_name=UNIQUE_SUBMISSION_NAME]
```

**Note:**

For more details on the verbs, see Deployment Procedure Verbs in the *Enterprise Manager Command Line Interface*.

Related Topics

- [About Exporting and Importing WebLogic Domain Partitions](#)
- [Exporting WebLogic Domain Partition](#)

Migrating a Domain to a Partition

This chapter provides an overview of exporting an existing domain and importing it as a partition into another domain with the help of Domain to Partition Conversion Tool (D-PCT). Enterprise Manager Cloud Control provides an user friendly user interface that helps you to convert domains into partitions.

The export tool component in D-PCT is used on the WebLogic Server installation that hosts the source domain. You can use this tool to create and configure partition, resource groups, and resource group templates. By default, this tool moves all applications, libraries, and resources to the new partition. Optionally, it also provides a mechanism for selecting individual applications, libraries, and resources.

In particular, this chapter covers the following:

- [Pre-Requisites and Limitations](#)
- [Exporting WebLogic Domain](#)
- [Importing WebLogic Domain Partition](#)

Pre-Requisites and Limitations

Prerequisites of Migrating a WebLogic Server Domain to a Domain Partition:

Before you configure D-PCT, you must fulfill the following prerequisites:

- The target domain must be configured on a WebLogic Server 12.2.1.1 or above installation.
- You must download and install JDK 8 on the host machine of the source domain. This is required for running the export tool on WebLogic Server 10.3.6.
- The target WebLogic Server 12.2.1.1 or above domain must be up and running.
- All servers in the domain must be up and running.
- For a straightforward conversion from a domain to a partition and to ensure correct operation of the import tool to create a virtual target based on the JavaScript Object Notation (JSON) file, it is important that the target and source domains have identical cluster and server names.
- Before using the import tool to create a new partition, you must ensure that the new domain is configured the same as the source domain with regard to servers, clusters, virtual targets, and security realms.

Limitations of Migrating a WebLogic Server Domain to a Domain Partition:

The following scenarios are not supported by Domain to Partition Conversion Tool (D-PCT) tool:

- The domain upgrade of libraries and resources to a new release is not supported at import. Administrators and users are responsible for possible necessary domain upgrades from the source domain before exporting the archive file.

- Neither the application runtime state nor the application specific runtime configuration is exported to the archive file. For example, the JMS messages in queues or users in an embedded LDAP realm are not exported.
- The remote clients compiled with WebLogic Server versions prior to 12.2.1 are not able to look up JNDI resources deployed on a partition. They need to be recompiled with the WebLogic Server releases 12.2.1 or later.
- JDBC does not support the use of the Logging Last Resource (LLR) feature in WebLogic Server Multitenant. Data sources with this option need to be converted to use a substitute setting that may or may not be an adequate replacement.

For more information on D-PCT limitations and considerations, see *Oracle Fusion Middleware Using Oracle WebLogic Server Multitenant* guide.

Exporting WebLogic Domain

To export the domain, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching** and click **Middleware Provisioning**.
2. Click **Create Profile** button.
3. Select **By Exporting a WebLogic Domain** option and click **Next**.
4. In the **Profile** tab, enter a profile name in the **Name** field.
5. Enter the description of the profile created in the **Description** field.
6. In the **Source Target** tab, enter the **Oracle Home Credential** by using the search icon against the field.
7. Enter the directory on the domain host in the **Java Home** field.

 Note:

JDK8 is required to perform an export operation.

8. In the **Domain to Partition Conversion Tool** section, enter the WebLogic Home in the **12.2.1.1.0+ Oracle Home** field by using the search icon against the field. The **Domain to Partition Conversion Tool** is present only in version 12.2.1.1.0 and above. Thus, the Oracle WebLogic Home version should be of version 12.2.1.1.0 or above.
9. Enter the operating system credential for Oracle Home in the **12.2.1.1.0+ Oracle Home Credential** field by using the search icon against the field.
10. Click **Create**.

Importing WebLogic Domain Partition

To import a domain partition into an existing WebLogic domain, see [Importing WebLogic Domain Partition](#).

Deploying / Redeploying / Undeploying Java EE Applications

This chapter explains how you can deploy, undeploy, and redeploy Java EE Applications using Oracle Enterprise Manager Cloud Control (Cloud Control).

In particular, this chapter covers the following:

- [Getting Started with the Deploy/Undeploy Java EE Application Deployment Procedure](#)
- [Deploying, Undeploying, or Redeploying Java EE Applications](#)
- [Supported Releases for Deploy/Undeploy Java EE Applications Deployment Procedure](#)
- [Prerequisites for Deploying/Undeploying Java EE Applications](#)
- [Creating a Java EE Application Component](#)
- [Java EE Applications Deployment Procedure](#)
- [Deploying / Redeploying / Undeploying Java EE Applications Using EMCLI](#)

Getting Started with the Deploy/Undeploy Java EE Application Deployment Procedure

This section provides an overview of the steps involved in deploying, redeploying, and undeploying Java EE Applications. The Deploy / Undeploy Java EE Applications Deployment Procedure allows you to perform the following operations:

- Deploy
- Undeploy
- Redeploy

Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision a Java EE Application. Click the reference links provided against the steps to reach the relevant sections that provide more information.

Table 34-1 Getting Started with Deploying, Undeploying, or Redeploying a Java EE Application

Step	Description	Reference Links
Step 1	Understanding the Deployment Procedure Understand the Deployment Procedures offered by Cloud Control to deploy, undeploy, or redeploy a Java EE Application. Know how the Deployment Procedure function, what use cases it covers, and so on.	To learn about the Deployment Procedure, see Deploying, Undeploying, or Redeploying Java EE Applications .
Step 2	Selecting the Use Case This chapter covers the use cases for deploying, undeploying, and redeploying Java EE Application. Select the use case that best matches your requirements.	- To deploy a Java EE Application, see Deploying a Java EE Application. - To redeploy a Java EE Application, see Redeploying a Java EE Application. - To undeploy a Java EE Application, see Undeploying a Java EE Application.
Step 3	Meeting the Prerequisites Before you run the Deployment Procedure, you must meet the prerequisites, such as configuring the Software Library and creating components to be provisioned as part of the Deploy / Undeploy Java EE Application deployment procedure.	To learn about the prerequisites for deploying, undeploying or redeploying Java EE Application, see Prerequisites for Deploying/Undeploying Java EE Applications .
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully deploy, redeploy, or undeploy one or more Java EE applications.	To run the Deploy / Undeploy Java EE Applications Deployment Procedure, follow the steps explained in Java EE Applications Deployment Procedure .

Deploying, Undeploying, or Redeploying Java EE Applications

This deployment procedure supports deployment of Java EE Applications packaged into .ear, .war, .jar, .rar or .gar files as per Java EE specifications. Administrators can use Cloud Control to deploy, redeploy, and undeploy one or more Java EE applications and need not drill down into the WebLogic Server Administration Console or the Oracle Enterprise Manager Fusion Middleware Control console to perform these tasks. The Java EE applications need to be pre-configured before you add them to the Cloud Control Software Library. You can deploy a pre-configured Java EE application to one or more WebLogic domains in Cloud Control.

The Java EE Application Deployment Procedure offers GUI-rich interactive screens that allow you to deploy / redeploy to, or undeploy a pre-configured Java EE application from one or more WebLogic Domains.

The Enterprise Manager Cloud Control 13c Release 2 Plug-ins Update 1 supports deploying/redeploying/undeploying to resource group templates, global resource groups, or partition resource groups.

Supported Releases for Deploy/Undeploy Java EE Applications Deployment Procedure

This deployment procedure can be used to deploy, undeploy, or redeploy Java EE Applications to and from WebLogic versions 10.3, 10.3.1 and later, including 12.1.1, 12.1.2, 12.1.3, 12.2.1, 12.2.1.x, and 12.1.x.

WebLogic versions 8.x and 9.x are not supported.

Prerequisites for Deploying/Undeploying Java EE Applications

Before running the Deploy / Undeploy Java EE Applications Deployment Procedure, ensure that the following prerequisites are met:

- Ensure that the Software Library is configured. See [Setting Up Oracle Software Library](#) for details.
- The Java EE Application component must have been created in the Software Library. See [Creating a Java EE Application Component](#) for details.
- The Management Agent must be installed on the Administration Server host machine for the WebLogic Domain to which you are deploying, undeploying, or redeploying the Java EE application. Deployment commands are executed from the Administration Server host machine.
- If a target execution script is to be run as part of the deployment procedure, then the Management Agent must be installed on the Managed Server's host machine where the target execution script is to be run.
- A target execution script can be used to set up the required environment or replace tokens in additional files such as property files. The script will be executed on selected targets.

Creating a Java EE Application Component

You can create a Java EE Application component which contains the archive, deployment plan, predeploy, postdeploy, target execution scripts, and other files required for deploying the Java EE application.

Note:

If you create a Java EE Application component on the host where the Software Library is configured, then you must ensure that this host name and the management agent host name should match, if not, you will see the following error while creating the component:

An error was encountered during saving entity samples.gar.
Please see the log for details.
oracle.sysman.emSDK.app.exception.EMSystemException

To create a Java EE Application component, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Create a folder or select a folder from the Software Library, select **Create Entity**, then select **Component**.
3. From the Create Entity: Component dialog box, select Java EE Application and click **Continue**.
4. In the Create Java EE Application: Describe page, enter the Name, Description, and click **Next**.
5. In the Create Java EE Application: Select Files page, select one or more files to be associated with the Java EE Application. You can upload files from a storage location in the Software Library. For Software Library to become usable, at least one upload file location must be configured. In the Specify Destination section, click the **Browse** button in the Upload Location field. Select either of the following:

- **OMS Shared File System:** An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Server (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running or in a shared location, so that it is accessible to all the OMS hosts. For multiple OMS environments, Oracle recommends that you configure the Software Library in a shared location so that the storage is accessible through NFS mount points to all Oracle Management Servers in the environment.

- **OMS Agent File System:** An OMS Agent File System location is a location that is accessible to one of the OMS host's Agent. This option is ideal for OMS installed on Windows hosts. By selecting this option for uploading files, you can avoid sharing a location between all participating OMS hosts.

Credentials must be set before using an OMS Shared File System or OMS Agent File System. For an OMS Shared File System, normal host credentials must be set before configuring a storage location. However, for OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

6. In the Specify Source section, you can add the standard Java EE archive files such as .ear, .war, .jar, .rar, .gar, and other optional files such pre and post-deploy scripts, target execution script, execution plan and additional files. You can either upload each file separately (Individual Files) or upload a zip file (Zip File) that contains the JavaEEAppComp.manifest file. You can upload the files from:
 - **Local Filesystem:** Click **Browse** and upload the files from your local system.
 - **Agent Filesystem:** You can upload the files from a remote filesystem monitored by the Management Agent. Click **Browse** and select a host machine from the list and click **Select**. Click **Add**. The Remote File Browser window is displayed. Click the **Login As** button and enter the credentials for the host machine. Specify the location in which the files are present, select one or more archive related files and click **Add**. The selected files are listed in the Current Selection section. Click **OK** to return to the Create Entity: Select Files page.
7. The files are listed in the table. Specify the type of the file by selecting the options in the Type field. Click **Next**.

8. Review and verify the information entered so far. Click **Save and Upload** to upload the files and create the Java EE Application component.

Java EE Applications Deployment Procedure

This section describes the Java EE Application deployment procedure. It covers the following:


- Deploying a Java EE Application
- Undeploying a Java EE Application
- Redeploying a Java EE Application

Deploying a Java EE Application

Follow these steps to deploy a Java EE Application:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. Select the Java EE Application procedure from the list and click **Launch**. You can also use the following method to launch the deployment procedure:
 - Click **Middleware** from the Targets menu.
 - Right click on a WebLogic Domain from the list and from the context sensitive menu, select **Provisioning**, then select Deploy / Undeploy Java EE Applications.
 - In the Deployment Procedure Manager page, select the Java EE Application Provisioning procedure and click **Launch**.
3. In the Deploy/Undeploy Java EE Applications: Select Targets page, choose the **Deploy** operation.

Figure 34-1 Deploy / Undeploy Java EE Applications: Select Target



Each domain has entities like clusters, admin server, managed server, global resource group, resource group templates, and partitions.

While deploying Multi-Tenancy partitions, each partition has one or more Resource Groups. The deployment is possible for only one of the selected Resource Groups of a partition.

Click the **radio button** to select the desired Resource Group for a partition.

While deploying applications to Resource Groups and Resource Group Templates, you must ensure that application deployment meets the uniqueness criteria of application names set by WebLogic. Deployment procedure wizard validates and ensures the uniqueness in simpler cases. In case of extended resource groups from resource group templates, the uniqueness check is left to WebLogic deployer as Enterprise Manager does not have the hierarchy details. Hence, you must ensure the uniqueness before hand to avoid a deployment procedure failure.

4. Select WebLogic Domains and select the targets on which the Java EE application is to be deployed. Click **Add WebLogic Domains**. Choose one or more WebLogic domains from the list and click **Select**.

 Note:

You can lock an operation, a target, or an application.. Click the **Lock** icon to lock an operation or the fields you are configuring in any of the pages in the wizard. Once the fields have been locked, the Operator needs to provide minimal input while running the deployment procedure

5. The selected WebLogic domains are listed in the Targets table. Select the targets (clusters, managed servers, admin servers, Global Resource Groups, Resource Group Templates, Partitions, or Resource groups) for each domain and click **Next**.
6. In the Deploy / Undeploy Java EE Applications: Select Applications page, add the archives and other related files that are to be deployed from the Software Library. Click **Add** to select one or more archives and other application related files or components from the Software Library. The Add Application popup is displayed. In the Component Name field, enter a file name or a wild card pattern to search for and retrieve components from the Software Library. Select the **Show Java EE Application Components Only** checkbox to list only the Java EE Application components in the Components in Software Library column. Select the archives and click the right arrow to move them to the Components Selected for Deployment section.

Figure 34-2 Deploy / Undeploy Java EE Applications: Add Applications

Add Application

Search Files
Select a Java EE Application component from the Software Library. Except for the archive, you can override any of the files present in the selected component by selecting a Generic component of the same type from the Software Library. If you do not have a Java EE Application component, uncheck the "Show Java EE Application Components Only" check box and select separate Generic components for Archive, Plan, Pre-Deploy Script, Post-Deploy Script, Target Execution Script, and any additional files. Specify the location on the WebLogic Server Host on which the application is to be deployed. This location can be an absolute path or a path relative to the \$WLS_HOME. See the Online Help for more details.

Component Name: riddles.ear ☒ Show Java EE Application Components Only

Components in Software Library		
Component Name	Path	Revision
Riddles.ear	testing/Riddles.ear	0.1

Components Selected for Deployment				
Component Name	Path	Revision	Type	Location On Target
No Components added yet				

Attributes

Application Name:

Application Version:

Deploy this archive as a library: ☐

Staging Mode: Default

Start Mode: ☒ Start in full mode (servicing all requests)
☐ Start application in admin mode (servicing only administration requests)
☐ Do not start

OK Cancel

**Note:**

In the image shown above, Application Version and Plan Version will not be visible if you are redeploying non-versioned application.

7. In the Type field, the type of each component is displayed. The Type can be:
 - **Archive:** This is the archive file which can be a .ear, .war, .jar, .rar, or .gar file.

**Note:**

If you have selected a .gar archive file, the WebLogic Domain to which the application is being deployed must be 12.1.2.0.0 or higher.

- **Plan:** This is an .xml file containing the deployment options for this application.
- **Pre Deploy Script:** This is a script containing WLST commands. The Management Agent runs this script on the Administration Server of each WebLogic domain before the application is deployed. You can use this script to create data sources, JMS end points, and any other resources that might be needed by the application that is being deployed.
- **Post Deploy Script:** This is a WLST script that is executed by the Management Agent on the Administration Server after the application is deployed. You can use this script to perform any post deployment configuration. For example, if you need to roll back and undo the changes made by the pre deploy script, you can select this option.

 Note:

The archive, plan, predeploy, and postdeploy scripts can be moved only to the Administration Server.

- **Additional File:** You can add one or more files that will be required by the application that are not part of the application archive. These files can be of any type and can be moved only to the selected targets (managed servers and clusters).
 - **Target Execution Script:** These scripts can be used to set up the required environment or replace tokens in the additional files like property files. These scripts will be executed on selected targets.
8. In the Location On Target field, for each component, specify the location on the WebLogic Server Host on which the application is to be deployed. This can be an absolute path or relative to the \$WLS_HOME for the selected targets.
 9. After selecting the required files for deployment, enter a unique name for the application and specify the Staging Mode which can be:
 - **Default:** Each server in the WebLogic Domain maintains two attributes which are Staging Mode and StagingDirectoryName. The Staging Mode is the default staging mode for the server and StagingDirectoryName is the location on which the staged files are stored. Select this option to use the default staging mode for all the targets.
 - **Stage:** Select this option if the archive files should be moved to the destination machine.
 - **No Stage:** Select this option if the archive files should not be moved to the destination machine.
 10. Select the **Deploy this archive as library** option if the application needs to be deployed as a shared library. You can select this option if one or more applications need the same set of files.
 11. Select the Start Mode for deployment which can be:
 - **Start in full mode (servicing all requests):** Select this option to make the deployed application available to all users.
 - **Start application in admin mode (servicing only administration requests):** If you select this option, the deployed application is available only to the Administrator.
 - **Do not start:** The application is deployed but not started. You can select this option if any manual post-deployment configuration is required.
 12. Click **OK** to add the archive and return to the Select Applications page. You can add more archives or click **Next** to proceed. If you have added more than one archive, select the **Skip on Failure** checkbox to skip any failed deployments and continue deploying the remaining applications.
 13. Click the **Lock** icon to lock the fields you have configured.

**Note:**

The Designer can lock the fields after configuring them. This ensures that the Operator can run the deployment procedure with minimal input.

14. Click **Next**. Specify the credentials for each domain you have selected, the host on which the Administration Server is running, and the hosts to which the additional files or execution scripts are to be moved. You can choose:

- **Preferred Credentials:** This option is selected by default and the preferred credentials stored in the Management Repository are used. This option is available only if it has already been defined in Cloud Control.
- **Named Credentials:** A named credential specifies the authentication information for a user and can be a combination of username / password, public and private key pair, and can be used to perform provisioning, patching, run jobs, and other system management tasks.

After selecting the credentials, click **Apply**.

For more information on setting up the credentials, see Enterprise Manager Security Guide.

15. Click the **Lock** icon to lock the fields you have configured. These fields cannot be edited once they are locked.
16. In the Schedule Deployment page, you can schedule the date on which the Java EE Application deployment procedure should be executed.
17. Click **Next**. 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 you submit the deployment procedure, you will return to the Procedure Activity page where you can view the status of the Deployment Procedure. After the Java EE Application has been deployed, you can search for the target and navigate to the Target Home page.

Redeploying a Java EE Application

Follow these steps to redeploy a Java EE application:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. Select the Java EE Application procedure from the list and click **Launch**. You can also use the following method to launch the deployment procedure:
 - Click **Middleware** from the Targets menu.
 - Right click on a WebLogic Domain from the list and from the context sensitive menu, select **Provisioning**, then select Deploy / Undeploy Java EE Applications.
 - In the Deployment Procedure Manager page, select the Java EE Application Provisioning procedure and click **Launch**.
3. In the Select Targets page, choose the **Redeploy** operation.

 Note:

Click the **Lock** icon to lock an operation or the fields you are configuring in any of the pages in the wizard. Once the fields have been locked, the Operator needs to provide minimal input while running the deployment procedure.

4. Click **Add WebLogic Domains** to add one or more WebLogic domains. In the list of targets displayed, choose a target and click Select.
5. The deployment targets are listed in the Targets table. Select the applications that need to be redeployed and click Next.
6. In the Select Applications page, a list of applications that can be redeployed are displayed. Select an application and click **Edit** to modify the archive details and other application related files. In the Application Details window, enter a file name or a wild card pattern to search for and retrieve files from the Software Library. Select the archives and click the right arrow to move them to the Components Selected for Deployment section.
7. In the Type field, the type of each component is displayed. The Type can be:
 - **Archive:** This is the archive file which can be a .ear, .war, .jar, or .rar file.
 - **Plan:** This is an .xml file containing the deployment options for this application.
 - **Pre Deploy Script:** This is a script containing WLST commands. The Management Agent runs this script on the Administration Server of each WebLogic domain before the application is deployed. You can use this script to create data sources, JMS end points, and any other resources that might be needed by the application that is being deployed.
 - **Post Deploy Script:** This is a WLST script that is executed by Management Agent on the Administration Server after the application is deployed. You can use this script to perform any post deployment configuration. For example, if you need to roll back and undo the changes made by the pre deploy script, you can select this option.

 Note:

The archive, plan, predeploy, and postdeploy scripts can be moved only to the Administration Server.

- **Additional File:** You can add one or more files that will be required by the application that are not part of the application archive. These files can be of any type and can be moved only to the selected targets (managed servers and clusters).
 - **Target Execution Script:** These scripts can be used to set up the required environment or replace tokens in the additional files like property files. These scripts will be executed on selected targets.
8. Review the default location on the target machine on which the component will reside. This can be an absolute path or relative to the \$WLS_HOME for the selected targets.

9. After selecting the required files for deployment, enter a unique name for the application and specify the Staging Mode which can be:
 - **Default:** Each server in the WebLogic Domain maintains two attributes which are Staging Mode and StagingDirectoryName. The Staging Mode is the default staging mode for the server and StagingDirectoryName is the location on which the staged files are stored. Select this option to use the default staging mode for all the targets.
 - **Stage:** Select this option if the archive files should be moved to the destination machine.
 - **No Stage:** Select this option if the archive files should not be moved to the destination machine.
 10. Select the Start Mode for deployment which can be:
 - **Start in full mode (servicing all requests):** Select this option to make the deployed application available to all users.
 - **Start application in admin mode (servicing only administration requests):** If you select this option, the deployed application is available only to the Administrator.
 - **Do not start:** The application is deployed but not started. You can select this option if any post-deployment configuration is required.
 11. Specify the Retirement Policy for the application. You can select:
 - **Allow the application to finish its current sessions and then retire:** Select this option if all the current sessions should be completed before retirement.
 - **Retire the previous version after retire timeout:** Specify a timeout period after which the application will be automatically retired.
-  **Note:**

The Retirement Policy field is applicable only when you are redeploying versioned application.
12. Click **OK** to add the archive and return to the Select Applications page. You can add more archives or click **Next** to proceed. If you have added more than one archive, select the **Skip on Failure** checkbox to skip any failed deployments and continue deploying the remaining applications.
 13. Click the **Lock** icon to lock the fields you have configured. These fields cannot be edited once they are locked.
 14. Click **Next**. Specify the credentials for each domain you have selected, the host on which the Administration Server is running, and the hosts to which the additional files or execution scripts are to be moved. You can choose:
 - **Preferred Credentials:** This option is selected by default and the preferred credentials stored in the Management Repository are used. This option is available only if it has already been defined in Cloud Control.
 - **Named Credentials:** A named credential specifies the authentication information for a user and can be a combination of username / password, public and private key pair, and can be used to perform provisioning, patching, run jobs, and other system management tasks.

For more information on setting up the credentials, see Enterprise Manager Security Guide.

15. Click the **Lock** icon to lock the fields you have configured. These fields cannot be edited once they are locked.
16. In the Schedule Deployment page, you can schedule the date on which the Java EE Application deployment procedure should be executed.
17. Click **Next**. 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 you submit the deployment procedure, you will return to the Procedure Activity page where you can view the status of the Deployment Procedure. After the Java EE Application has been redeployed, you can search for the target and navigate to the Target Home page.

Undeploying a Java EE Application

Follow these steps to undeploy a Java EE Application:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. Select the Java EE Application procedure from the list and click **Launch**. You can also use the following method to launch the deployment procedure:
 - Click **Middleware** from the Targets menu.
 - Right click on a WebLogic Domain from the list and from the context sensitive menu, select **Provisioning**, then select Deploy / Undeploy Java EE Applications.
 - In the Deployment Procedure Manager page, select the Java EE Application Provisioning procedure and click **Launch**.
3. In the Select Targets page, choose the **Undeploy** operation.
4. Click **Add WLS Domains** to add one or more WebLogic domains. In the list of targets displayed, choose a target and click **Select**.
5. The deployment targets are listed in the Targets table. Select the application that needs to be undeployed and click **Next**.

ORACLE Enterprise Manager Cloud Control 13c

SYSMAN

Select Targets | Select Applications | Set Credentials | Schedule | Review

Deploy/Undeploy Java EE Applications: Select Targets

Save As... | Back | Step 1 of 5 | Next | Cancel

Use this deployment procedure to deploy, redeploy to, or undeploy one or more Java EE applications from one or more WebLogic Domains. Before you run the deployment procedure, ensure that the application archives, deployment plan, scripts, or other application related files have been added to the Software Library.

Operation

Deploy | Redeploy | **Undeploy**

Targets

Select WebLogic Domains | Remove WebLogic Domain

Current Applications	Targets	Administration Server Host	Delete EM Target
EMGC_EMGC_DOMAIN\EMGC_DOMAIN		<Host_Name>	
emgc	EMGC_OMS1		☒
emps	EMGC_OMS1		☒
jmsengine(13.2.2.0.0)	EMGC_OMS1		☐

You can select the **Delete EM Target** checkbox for each selected application to automatically erase target history after the undeployment procedure.

6. Specify the credentials for each domain you have selected, the host on which the Administration Server is running, and the hosts to which the additional files or execution scripts are to be moved. You can choose:
 - **Preferred Credentials:** This option is selected by default and the preferred credentials stored in the Management Repository are used. This option is available only if it has already been defined in Cloud Control.
 - **Named Credentials:** A named credential specifies the authentication information for a user and can be a combination of username / password, public and private key pair, and can be used to perform provisioning, patching, run jobs, and other system management tasks.

For more information on setting up the credentials, see Enterprise Manager Security Guide.

7. Specify the deployment schedule and click **Next**.
8. Review the details and click **Undeploy**. You will return to the Procedure Activity page where you can check the status.

Deploying / Redeploying / Undeploying Java EE Applications Using EMCLI

This section lists the procedures to deploy, redeploy, or undeploy a Java EE application using EMCLI:

- [Deploying a Java EE Application Using EMCLI](#)
- [Undeploying a Java EE Application Using EMCLI](#)
- [Redeploying a Java EE Application Using EMCLI](#)

Deploying a Java EE Application Using EMCLI

To deploy a Java EE application using EMCLI, run the following command:

```
emcli submit_procedure -name="JAVAEE_APPLICATION_DEPLOYMENT" -input_file=data="/
input-deploy.properties"
[instance_name="javaee_deploy"]
```

Sample Input File

This section shows the format of a sample `input-deploy.properties` file.

```
stopOnError=true
operation=Deploy

#Host Creds for all the hosts part of the domain
domains.0.domainTargetName=/Farm01_wls1213_njrf_1c1ms/wls1213_njrf_1c1ms
domains.0.wlsCredType=New
domains.0.wlsUserName=weblogic1
domains.0.wlsPassword=welcome1$
#domains.0.credName=wls_creds:SYSMAN
```



```

#Host Creds for all the hosts part of the domain
domains.0.hosts.0.hostName=hostname.example.com
domains.0.hosts.0.hostCredType=New
domains.0.hosts.0.hostUserName=<USERID>
domains.0.hosts.0.hostPassword=<PASSWORD>
#domains.0.hosts.0.hostCredName=user1_creds:<USERID>

#For each application
domains.0.javaeeApps.0.appName=MyCalender
domains.0.javaeeApps.0.isSharedLib=false
domains.0.javaeeApps.0.stageMode=DEFAULT
domains.0.javaeeApps.0.startMode=full
domains.0.javaeeApps.0.retirementPolicy=true
domains.0.javaeeApps.0.retirementTimeout=0
domains.0.javaeeApps.0.targetets=Cluster_1
#domains.0.javaeeApps.0.deleteTarget=false

domains.0.javaeeApps.0.archiveSwLibPath=demo_folder/calendar_app
domains.0.javaeeApps.0.archiveFileName=Calendar.war
domains.0.javaeeApps.0.archiveStagingLocation=/my_arch/mystage

# If there is no plan, then this can be skipped
#domains.0.javaeeApps.0.planSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.planFileName=xyz.jar
#domains.0.javaeeApps.0.planStagingLocation=/my_plan/mystage

# If there is no predeploy script, this can be skipped
#domains.0.javaeeApps.0.preDeployScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.preDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.preDeployScriptStagingLocation=/my_prescript/
mystage

# If there is no postdeploy script, this can be skipped
#domains.0.javaeeApps.0.postDeployScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.postDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.postDeployScriptStagingLocation=/my_postscript/
mystage

# If there are no additional scripts, this can be skipped
#domains.0.javaeeApps.0.additionalFiles.0.swLibPath=my_folder/calpp
#domains.0.javaeeApps.0.additionalFiles.0.fileName=additional.zip
#domains.0.javaeeApps.0.additionalFiles.0.stagingLocation=/my_add/mystage

# If there is no target execution script, this can be skipped
#domains.0.javaeeApps.0.targetExecScriptSwLibPath=my_folder/my_script
#domains.0.javaeeApps.0.targetExecScriptFileName=script.pl
#domains.0.javaeeApps.0.targetExecScriptStagingLocation=/my_exec/mystage

```

Description of the Input File

- stopOnError: Set this flag to **true** if the deployment procedure should stop when an error is encountered while a Java EE application is being deployed.
- operation: The type of operation (**Deploy**) being performed.

- Domain Details: Specify the details of the WebLogic domains on which these applications are being deployed.
 - domains.0.domainTargetNameWhere:
 - * 0 is the index of the domain
 - * domainTargetName is the name of the target in Enterprise Manager.
 - domains.0.wlsCredType: The credentials for the domain. Possible values are Preferred, New, or Named Credentials.
 - domains.0.wlsUserName (optional): If New Credentials are selected, specify the user name here.
 - domains.0.wlsPassword (optional): If New Credentials are selected, specify the password here.
 - domains.0.credName (optional): If Named Credentials are selected, specify the name of the credentials followed by the name of the credential owner in the format <cred_name>:<owner_name>. For example, user1_creds:SYSMAN.
- Host Details:

Specify the Administration Server host details here.

 - hosts.0.hostName: where:
 - * 0 is the index for the Administration Server host. If additional files, and target execution scripts are present, indexes 1, 2, 3, and so on must be used for all the other hosts in the domain.
 - * hostName is the name of the Administration Server host.
 - domains.0.hosts.0.hostCredType: The credentials for the Administration server. Possible values are Preferred, New, or Named Credentials.
 - domains.0.wlsUserName(optional): If New Credentials are selected, specify the user name here.
 - domains.0.wlsPassword(optional): If New Credentials are selected, specify the password here.
 - domains.0.credName(optional): If Named Credentials are selected, specify the name of the credentials followed by the name of the credential owner in the format <cred_name>:<owner_name>. For example, user1_creds:SYSMAN.
- Application Details:

Specify the details of the Java EE applications that are being deployed.

 - javaeeApps.0.appName: where:
 - * javaeeApps.0.appName: The name of the Java EE application being deployed. Since multiple applications can be deployed, the index (0) identifies each application being deployed.
 - * isSharedLib: Set this flag to **true** if the application is to be deployed in a shared library.
 - * stageMode The stage mode can be **Default**, **Stage**, and **No Stage**.
 - * startMode: This can be **Full**, **Admin**, or **Do not start**.
 - * retirementPolicy: Set this flag to **true** if a retirement policy is available.
 - * retirementtimeout: Specify the retirement timeout period here.

- * targets: The target(s) on which the applications are to be deployed.
- * archiveSwLibPath: The location in which the archive is stored in the Software Library.
- * archiveFileName: The name of the archive file.
- * archiveStagingLocation: The location in which the archive should be staged on the Administration Server host.

**Note:**

In the remaining optional sections, you can specify the Software Library path, the file name, and the staging location as necessary.

Undeploying a Java EE Application Using EMCLI

To undeploy a Java EE application using EMCLI, run the following command:

```
emcli submit_procedure -name="JAVAEE_APPLICATION_UNDEPLOY" -input_file=data="/input-undeploy.properties"
[instance_name="javaee_undeploy"]
```

Sample Input File

This section shows the format of a sample `input-undeploy.properties` file.

```
stopOnError=true
operation=Undeploy

#Host Creds for all the hosts part of the domain
domains.0.domainTargetName=/Farm01_wls1213_njrf_1c1ms/wls1213_njrf_1c1ms
domains.0.wlsCredType=New
domains.0.wlsUserName=weblogic1
domains.0.wlsPassword=welcome1$
#domains.0.credName=wls_creds:SYSMAN

#Host Creds for all the hosts part of the domain
domains.0.hosts.0.hostName=hostname.example.com
domains.0.hosts.0.hostCredType=New
domains.0.hosts.0.hostUserName=user1
domains.0.hosts.0.hostPassword=KDsa0105
#domains.0.hosts.0.hostCredName=user1_creds:SYSMAN

#For each application
domains.0.javaeeApps.0.appName=MyCalendar
#domains.0.javaeeApps.0.isSharedLib=false
#domains.0.javaeeApps.0.stageMode=DEFAULT
#domains.0.javaeeApps.0.startMode=full
#domains.0.javaeeApps.0.retirementPolicy=true
#domains.0.javaeeApps.0.retirementTimeout=0
domains.0.javaeeApps.0.targets=Cluster_1
domains.0.javaeeApps.0.deleteTarget=true

#domains.0.javaeeApps.0.archiveSwLibPath=my_folder/calpp
```

```
#domains.0.javaeeApps.0.archiveFileName=file.jar
#domains.0.javaeeApps.0.archiveStagingLocation=/my_arch/mystage

# If there is no plan, then this can be skipped
#domains.0.javaeeApps.0.planSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.planFileName=xyz.jar
#domains.0.javaeeApps.0.planStagingLocation=/my_plan/mystage

# If there is no predeploy script, this can be skipped
#domains.0.javaeeApps.0.preDeployScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.preDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.preDeployScriptStagingLocation=/my_prescript/
mystage

# If there is no postdeploy script, this can be skipped
#domains.0.javaeeApps.0.postDeployScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.postDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.postDeployScriptStagingLocation=/my_postscript/
mystage

# If there are no additional scripts, this can be skipped
#domains.0.javaeeApps.0.additionalFiles.0.swLibPath=my_folder/calpp
#domains.0.javaeeApps.0.additionalFiles.0.fileName=additional.zip
#domains.0.javaeeApps.0.additionalFiles.0.stagingLocation=/my_add/mystage

# If there is no target execution script, this can be skipped
#domains.0.javaeeApps.0.targetExecScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.targetExecScriptFileName=tscrip.pl
#domains.0.javaeeApps.0.targetExecScriptStagingLocation=/my_exec/mystage
```

Description of the Input File

- stopOnError: Set this flag to **true** if the deployment procedure should stop when an error is encountered while a Java EE application is being undeployed.
- operation: The type of operation (**Undeploy**) being performed.
- Domain Details: Specify the details of the WebLogic domains on which these applications are being undeployed.
 - domains.0.domainTargetNamewhere:
 - * 0 is the index of the domain
 - * domainTargetName is the name of the target in Enterprise Manager.
 - domains.0.wlsCredType: The credentials for the domain. Possible values are Preferred , New. or Named Credentials.
 - domains.0.wlsUserName(optional): If New Credentials are selected, specify the user name here.
 - domains.0.wlsPassword(optional): If New Credentials are selected, specify the password here.
 - domains.0.credName (optional): If Named Credentials are selected, specify the name of the credentials followed by the name of the credential owner in the format <cred_name>:<owner_name>. For example, user1_creds:SYSMAN.
- Host Details:

Specify the Administration Server host details here.

- `hosts.0.hostName`: where:
 - * `0` is the index for the Administration Server host. If additional files, and target execution scripts are present, indexes 1, 2, 3, and so on must be used for all the other hosts in the domain.
 - * `hostName` is the name of the Administration Server host.
- `domains.0.hosts.0.hostCredType`: The credentials for the Administration Server. Possible values are Preferred, New, or Named Credentials.
- `domains.0.wlsUserName(optional)`: If New Credentials are selected, specify the user name here.
- `domains.0.wlsPassword(optional)`: If New Credentials are selected, specify the password here.
- `domains.0.credName(optional)`: If Named Credentials are selected, specify the name of the credentials followed by the name of the credential owner in the format `<cred_name>:<owner_name>`. For example, `user1_creds:SYSMAN`.

- Application Details:

Specify the details of the Java EE applications that are being undeployed.

- `javaeeApps.0.appName`: where:
 - * `javaeeApps.0.appName`: The name of the Java EE application being undeployed. Since multiple applications can be undeployed, the index (0) identifies each application being deployed.
 - * `isSharedLib`: Set this flag to **true** if the application being undeployed is present in a shared library.
 - * `stageMode`: The stage mode can be **Default**, **Stage**, and **No Stage**.
 - * `startMode`: This can be **Full**, **Admin**, or **Do not start**.
 - * `retirementPolicy`: Set this flag to **true** if a retirement policy is available.
 - * `retirementtimeout`: Specify the retirement timeout period here.
 - * `targets`: The target(s) from which the applications are to be undeployed.
 - * `deleteTarget`:

Set this flag to **true** if the target is to be deleted after the application has been undeployed.
 - * `archiveSwLibPath`: The location in which the archive is stored in the Software Library.
 - * `archiveFileName`: The name of the archive file.
 - * `archiveStagingLocation`: The location in which the archive should be staged on the Administration Server host.



Note:

In the remaining optional sections, you can specify the Software Library path, the file name, and the staging location as necessary.

Redeploying a Java EE Application Using EMCLI

To redeploy a Java EE application using EMCLI, run the following command:

```
emcli submit_procedure -name="JAVAEE_APPLICATION_REDEPLOYMENT" -input_file=data="/
input-redeploy.properties"
[instance_name="javaee_redeploy"
```

Sample Input File

```
stopOnError=true
operation=Redeploy

#Host Creds for all the hosts part of the domain
domains.0.domainTargetName=/Farm01_wls1213_njrf_1c1ms/wls1213_njrf_1c1ms
domains.0.wlsCredType=New
domains.0.wlsUserName=weblogic1
domains.0.wlsPassword=welcome1$
#domains.0.credName=wls_creds:SYSMAN

#Host Creds for all the hosts part of the domain
domains.0.hosts.0.hostName=hostname.example.com
domains.0.hosts.0.hostCredType=New
domains.0.hosts.0.hostUserName=user1
domains.0.hosts.0.hostPassword=KDsa0105
#domains.0.hosts.0.hostCredName=user1_creds:SYSMAN

#For each application
domains.0.javaeeApps.0.appName=MyCalender
domains.0.javaeeApps.0.isSharedLib=false
domains.0.javaeeApps.0.stageMode=DEFAULT
domains.0.javaeeApps.0.startMode=full
domains.0.javaeeApps.0.retirementPolicy=true
domains.0.javaeeApps.0.retirementTimeout=0
domains.0.javaeeApps.0.targets=Cluster_1
#domains.0.javaeeApps.0.deleteTarget=false

domains.0.javaeeApps.0.archiveSwLibPath=demo_folder/calendar_app
domains.0.javaeeApps.0.archiveFileName=Calendar.war
domains.0.javaeeApps.0.archiveStagingLocation=/my_arch/mystage

# If there is no plan, then this can be skipped
#domains.0.javaeeApps.0.planSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.planFileName=xyz.jar
#domains.0.javaeeApps.0.planStagingLocation=/my_plan/mystage

# If there is no predeploy script, this can be skipped
#domains.0.javaeeApps.0.preDeployScriptSwLibPath=my_folder/calpp
#domains.0.javaeeApps.0.preDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.preDeployScriptStagingLocation=/my_prescript/
mystage

# If there is no postdeploy script, this can be skipped
#domains.0.javaeeApps.0.postDeployScriptSwLibPath=my_folder/calpp
```

```
#domains.0.javaeeApps.0.postDeployScriptFileName=xyz.py
#domains.0.javaeeApps.0.postDeployScriptStagingLocation=/my_postscript/
mystage

# If there are no additional scripts, this can be skipped
#domains.0.javaeeApps.0.additionalFiles.0.swLibPath=my_folder/calpp
#domains.0.javaeeApps.0.additionalFiles.0.fileName=additional.zip
#domains.0.javaeeApps.0.additionalFiles.0.stagingLocation=/my_add/mystage

# If there is no target execution script, this can be skipped
#domains.0.javaeeApps.0.targetExecScriptSwLibPath=my_folder/my_script
#domains.0.javaeeApps.0.targetExecScriptFileName=script.pl
#domains.0.javaeeApps.0.targetExecScriptStagingLocation=/my_exec/mystage
```



Note:

See [Deploying a Java EE Application Using EMCLI](#) for details on the input details.

Provisioning Coherence Nodes and Clusters

This chapter explains how you can provision Coherence nodes or clusters across multiple targets in a farm using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started](#)
- [Supported Releases](#)
- [Deploying Coherence Nodes and Clusters](#)
- [Troubleshooting](#)

For more information about Oracle Coherence, see *Managing Oracle Coherence in the Enterprise Manager Cloud Control Oracle Fusion Middleware Management Guide*.

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 Cloud Control monitored target.
- Add a management node to an existing cluster and add this cluster as an Cloud Control monitored target.
- Add one or more nodes to a cluster that is already being monitored by Cloud Control.
- Update existing nodes by copying modified software components or configuration files and restart the nodes.
- Create a new Coherence cluster.

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 35-1 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 Cloud 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 Deploying Coherence Nodes and Clusters .
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 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 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 Deployment Procedure .

Supported Releases

This section lists the releases of Oracle Coherence supported by this Deployment Procedure:

- Oracle Coherence 3.5, 3.6, and 3.7.
- Oracle Coherence 12.1.2 Standalone Version.

Deploying Coherence Nodes and Clusters

This section describes how you can add one or more nodes to an existing cluster or create a new cluster using the Coherence Node Provisioning procedure.

This section covers the following:

- [Prerequisites](#)
- [Creating a Coherence Component](#)
- [Deployment Procedure](#)

Prerequisites

Before running the Deployment Procedure, meet the following prerequisites:

- The host on which a Coherence node is being added or updated must be a monitored target in Cloud Control.

- 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 Oracle Software Library. If the size of the zip file is more than 25 MB, it must be uploaded from a host monitored by the Management Agent. You can add specific configuration files as components to the Software Library which will override the default configuration files. These configuration files can be different depending on the type of node (storage, management, etc.). While adding a software component, it is recommended that you specify the Product Name as **Coherence**.
- If you are provisioning a new node on a host on which the Coherence binaries are not present, you must upload the coherence.zip file and the default-start-script.pl to the Oracle Software Library. Each file required for provisioning must be uploaded as an individual software component.

Creating a Coherence Component

You can create one or more Coherence components and save it to the Software Library. This components are required while provisioning Coherence nodes and clusters. To create a Coherence component, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Create a folder in which the components are to be stored. After the folder has been created, right click on the folder and select **Create Entity** and **Component** from the **Actions** menu.
3. A Create Component popup window appears. From the Select Subtype drop down list, select the **Generic Component** and click **Continue**.

Figure 35-1 Create Generic Component: Describe Page

Software Library

Describe Configure Select Files Set Directives Review

Create Generic Component : Describe

Back Step 1 of 5 Next Save Cancel

Parent Directory Components
Subtype Generic Component

Enter name, description and other attributes that describe the entity. These attributes are shared by all revisions of this entity. Additionally, attach any documents and keep notes.

* Name test

Description

Other Attributes

Name	Value	Description
Product Version		Product Version of the software component
Product		Product the software component represents
Vendor		Vendor of the software component

Attachments

+ Add - Remove

File Name	Size(KB)	Mime Type
No attachment has been added yet.		

Notes

New Note + Add

Note	Added By	Date
No note has been added yet.		

TIP Notes once added cannot be deleted or edited.

4. In the Create Generic Component: Describe page, enter the Product as Coherence. This is helpful when you are searching for files during Coherence Provisioning. Click **Next**.

5. In the Create Generic Component: Select Files page, check the **Upload Files** option.

Figure 35-2 Create Generic Component: Select Files Page

6. In the Specify Source section, select Agent Machine in the File Source drop down box and upload the following files:
 - \$AGENT_ROOT/plugins/oracle.sysman.emas.agent.plugin_12.1.0.1.0/archives/coherence/bulkoperationsmbean_11.1.1.jar
 - \$AGENT_ROOT/plugins/oracle.sysman.emas.agent.plugin_12.1.0.1.0/archives/coherence/coherenceEMIntg.jar
 - \$AGENT_ROOT/plugins/oracle.sysman.emas.agent.plugin_12.1.0.1.0/scripts/coherence/default-start-script.pl

If you are uploading a large zip file such as Coherence.zip, you must save it on the Agent machine. Specify the path on the Agent to upload the file. This zip file must be downloaded from <http://www.oracle.com/technetwork/middleware/coherence/downloads/index.htm>

7. Click **Save and Upload** to submit a file transfer job to upload the remote files to the specified upload location.

Deployment Procedure

To deploy a Coherence node or a cluster, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. Select the Coherence Node Provisioning deployment procedure and click **Launch**.

Note:

You can also use the following methods to launch the deployment procedure:

- From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**. Select the Coherence deployment procedure from the list and click **Launch**.
- Select the **Coherence Node Provisioning** option from the Coherence Home page menu.

3. The Source Selection page which is the first page of the Coherence Node Provisioning wizard is displayed.

Figure 35-3 Source Selection Page

Source

Add all software components needed for Coherence node provisioning or update. For each component, you need to provide destination directory information on the target machine. This can be an absolute path or a relative path from `${INSTALL_DIR}` or `${COHERENCE_HOME}`. The value of `INSTALL_DIR` variable needs to be defined in Target Selection step. `COHERENCE_HOME` will be set to `INSTALL_DIR/coherence`. You can skip adding software components if these are available on the target machine.

Component Name	Revision	Destination Directory	Location	File Name	Remove
Common Provisioning Module	0.1	<code>\${INSTALL_DIR}</code>	Common Provisioning Utilities/11.2.0.1.0/all_platforms/Common Provisioning Module	ProvCommon.pm	

4. You can add all the software components needed to add or update a Coherence cluster. If the Coherence Home has already been created, you can click **Next** to go to the next page.

If the Coherence Home does not exist, click **Add** to add the Coherence binaries and the Start script from the Software Library. The Select Source popup is displayed. All software components with the product name Coherence that are present in the Software Library are displayed. Select required components and click **Select**.

5. For each component you have selected, specify the destination directory on the target machine. This can be an absolute path or a relative path from `${INSTALL_DIR}` or `${COHERENCE_HOME}`. The contents of the coherence.zip file will be extracted to this directory.

Note:

- The value of the `$INSTALL_DIR` is defined in the Coherence Node Provisioning: Target Selection page and it is set as a level above the `COHERENCE_HOME` directory.
- If the software components are available in the target machine, this step can be skipped.

6. Click **Next**. The Target Selection page is displayed. On this page, you can:
 - **Add Nodes:** You can add a new node or make a copy of an existing node. Click the **Search** icon in the Target Name field and select a Coherence cluster from the list. The following details are displayed:
 - **Cluster Name:** The name of the cluster.
 - **Cluster Communication:** This can be Multicast or Well Known Address (WKA).
 - **License Mode:** The mode in which the cluster has been deployed.

Click **Add** in the New Nodes section to add new nodes to an existing Coherence cluster monitored by Cloud Control. The Add Coherence Node page is displayed. Specify the details of the node and click **Continue** to add the node

and return to the Coherence Node Provisioning: Select Target page. See [Adding a Coherence Node](#) for details.

- **Creating a Copy of an Existing Node:** You can make a copy of an existing node. When you select a cluster to which a node is to be added, a list of nodes present in the cluster are displayed in the Existing Nodes section. Select a node from this list and click **Create Like**. The Add Coherence Node page is displayed. Specify the details of the node and click **Continue** to return to the Coherence Node Provisioning: Select Target page. A copy of the selected node is now listed in the New Nodes section and can be deployed.
- **Create Cluster:** Click **Create Cluster** to create a new Coherence cluster. Enter Cluster Name along with the following details:
 - **Cluster Name:** Enter a unique name for the cluster.
 - **Cluster Communication:** Select **Multicast** or **Well Known Address (WKA)**. If you select Multicast, you are prompted for the Cluster Port and Cluster Address. If you select WKA, enter one or more sets of <host-name>:<port> entries separated by a comma and make sure that the WKA details have been specified in the Coherence Node Provisioning: Add Node page.

 Note:

- * If you select WKA, you need to create an override file (tangosol-coherence-override-em.xml) for the WKA entries and specify the Dtangosol.coherence.override parameter in this file. This file is created from the default-start-script.pl which can be modified if required.
- * If you select WKA for Cluster Communication while creating a copy of an existing node (**Create Like** option), you are prompted for the IP address instead of host name in the Well Known Address field.

- **License Mode:** Specify the mode in which the cluster is to be deployed. This can be Development, Evaluation, or Production.

Click **Add** in the New Nodes section to add nodes to the new cluster being created. The Add Coherence Node page is displayed. Specify the details of the node and click **Continue** to add the node and return to the Coherence Node Provisioning: Select Target page.

To create a cluster, you must have the following components, coherenceE-MIntg.jar, bulkoperationsmbean_11.1.1.jar, default-start-script.pl, and Coherence.zip.

7. Click **Next** to go to the next step in the wizard. In the Coherence Node Provisioning: Set Credentials page, you can set credentials for each host. You can apply the same credentials for multiple hosts by selecting multiple hosts from the list.
8. Select the host and specify the credentials which can be:
 - **Preferred Credentials:** This option is selected by default and the preferred credentials stored in the Management Repository are used. This option is available only if it has already been defined in Cloud Control.

- **Named Credentials:** You can override the preferred credentials and select a common set of credentials that will be used for all the hosts and WebLogic domains.
- **New Credentials:** You can override the preferred credentials and specify a separate set of credentials for each host.

Select the credentials and click **Apply** to apply the credentials to the selected hosts. For more information on setting up credentials, see the Enterprise Manager Security chapter in the Enterprise Manager Security Guide.

9. Click **Next**. 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.

**Note:**

If you set the Grace Period as **Indefinite**, Cloud Control 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.

10. Click **Next**. The Review page is displayed. You can review the details you have provided for deploying the node. If you are updating a node, you can view the node processes that will be stopped on this page. Click **Finish** to deploy or update the node.

After the new Management Node has been created, you must wait for the first collection before you add nodes to the cluster.

Adding a Coherence Node

You can add a node to an existing cluster or create a new cluster by adding one or more new nodes. To add a node, follow these steps:

1. Click **Add** in the Coherence Node Provisioning: Target Selection page. The Add Coherence Node page is displayed. Enter the following details:

Figure 35-4 Add Coherence Node Page

Add Coherence Node Cancel Continue

Host Details
Select an EM Agent managed host where this Coherence node would be added. Enter required number of nodes on this host depending on the available resources. All nodes created will use same software components from `INSTALL_DIR`. Port values will be incremented by one from the entered value.

* Host Name: Number of Nodes:

Node Details
Provide node identification details.

* Node Name: Site Name:
 Rack Name: Role Name:
 Unicast Address: Unicast Port:
 Well Known Address(WKA):

☐ Do not copy software components.

JVM Diagnostics Details
Enter JVM Manager Host and JVM Manager Port to monitor this node using JVM

JVM Manager Host: JVM Manager Port:

Management Node Details
Select the Management node with MBeanServer checkbox and provide additional information. Only the primary management node is used for monitoring, but it is recommended to have at least one another management node for fail-over. Select Primary management node for monitoring checkbox if this node is used for monitoring. Provide JMX user name and password if JMX authentication is enabled.

☒ Management node with MBeanServer

* JMX Remote Port: JMX User Name:
 JMX Password: ☒ Primary management node used for monitoring

☒ Use Bulk Operations MBean
Recommended for efficient monitoring.

Environment Details
Provide absolute paths to following home variables. Start script can be absolute path or relative to `$(INSTALL_DIR)`. COHERENCE_HOME is the full path to the `INSTALL_DIR/coherence` directory.

* Install Directory: * Start Script:

Cancel Continue

Table 35-2 Add Coherence Node Page

Field Name	Description
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 <code>tangosol_coherence_machine</code> environment variable to set the <code>tangosol.coherence.machine</code> parameter and <code>oracle_coherence_machine</code> to set the <code>oracle.coherence.machine</code> 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 depending on the machine and node configuration. If the value is more than 1, then all the nodes will have the following properties: - Each node will use the same <code>COHERENCE_HOME</code> and start script. - The Node Name value will be added as the prefix and a number will be appended to each node. For example <code><node_name>_1</code>, <code><node_name>_2</code> 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.
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.

Table 35-2 (Cont.) Add Coherence Node Page

Field Name	Description
Do not Copy Software Components	If you are adding a node on a machine on which a Coherence Home is already present, you must check the Do not copy Software Components checkbox. If you are copying the files onto a new host, unchecked the Do not copy Software Components checkbox to ensure that the binaries selected in the Coherence Node Provisioning: Source Selection page.
Well Known Address (WKA)	If Cluster Communication has been set to WKA in the Coherence Node Provisioning: Target Selection page, enter the host and port number in the format host1:port1, host2:port2 and so on.
JVM Diagnostics Details	
JVM Manager Host and Port	If this node is to be monitored by JVM Diagnostics Manager, specify the address and port number of the JVM Diagnostics Console.
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 <code>com.sun.management.jmxremote.authenticate=true</code> . The JMX User name and JMX Password need to be set in the <code>\$JDK_HOME/jre/lib/management/jmxremote.password</code> and <code>\$JDK_HOME/jre/lib/management/jmxremote.access</code> 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. This node is used to discover the Coherence cluster and any nodes added later will be added to the newly discovered cluster. 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.
Use Bulk Operations MBean	This checkbox is selected by default. When this option is selected, a new management node with BulkOperationsMBean will be started.
Install Directory	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 summarizes how the values specified during deployment will be used by the environment variables specified in the start script. The deployment procedure also sets the `JAVA_HOME` and `AGENT_HOME` variables by using the Agent installation details. You may override these by specifying appropriate values in your start script.

Table 35-3 Environment Variables

UI Parameter	Environment Variable Set	Coherence System Parameter
CLUSTER_NAME	tangosol_coherence_cluster	tangosol.coherence.cluster
CLUSTER_ADDRESS	tangosol_coherence_clusteraddress	tangosol.coherence.clusteraddress
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.port
LICENSE_MODE	license_mode	tangosol.coherence.mode
COHERENCE_HOME	coherence_home	oracle.coherence.home
START_SCRIPT	start_script	oracle.coherence.startscript
JVM_CONSOLE_HOST	jvm_console_host	jamconhost
JVM_CONSOLE_PORT	jvm_console_port	jamconport
WKA_PORT	wka_port	tangosol.coherence.override=em-coherence-override.xml

2. After adding the node, click **Continue** to return to the Coherence Node Provisioning: Target Selection page.
3. Click **Next** to go to the next step in the wizard.

Sample Scripts

The default-start-script.pl and generate-wka-override.pl scripts are present in the \$EMAS_PLUGIN_ROOT/scripts/coherence/directory.

default-start-script.pl

This script is the default start script used to start a Coherence node. A sample script is shown below:

```
#!/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, the 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.

#
# - 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 $agent_home=$ENV{'AGENT_HOME'};
my $wka_port=$ENV{'WKA_PORT'};
my $license_mode=$ENV{'LICENSE_MODE'};
my $jamhost=$ENV{'JAM_CONSOLE_HOST'};
my $jampport=$ENV{'JAM_CONSOLE_PORT'};

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 $bulkmban=$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="\\";
}

print
"\n\n*****\n"

print "Output from default-start-script\n";
print "Starting Node : $member\n";
print "Coherence Home : $coherence_home \n";

```

```
print "Start Script : $start_script \n";
print "Java Home : $java_home \n";
print "Agent Home : $agent_home \n";
print "WKA Port: $wka_port \n";
print "License Mode: $license_mode \n";

print "Site Name : $site \n";
print "Rack Name : $rack \n";
print "Machine Name : $machine \n";
print "Oracle Coherence Machine Name : $oracle_coherence_machine \n";
print "Role Name : $role \n";

print "Cluster Name : $cluster \n";
print "Cluster Addr : $clusteraddr \n";
print "Cluster Port : $clusterport \n";
print "JMX Port : $jmxport \n";
print "Bulk MBean : $bulkmbean \n";
print "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 limitaion
# 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";
}

if("$license_mode" ne "") {
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.mode=$license_mode";
}
#This is used to generate WKA override file. If you choose to use an existing
#override file, you can comment this out.
#Make sure you set "-Dtangosol.coherence.override" to the appropriate file name.
if("$wka_port" ne "") {
$wka_port = "\""$wka_port."\"";
$wka_script = $agent_
home.$dsep."sysman".$dsep."admin".$dsep."scripts".$dsep."coherence".$dsep."genera
te-wka-override.pl";
print "executing $wka_script $wka_port\n";
if ( !&IsWindows() ) {
system("chmod 0700 $wka_script");
}
if(fork() == 0) {
exec("$wka_script $wka_port") or die "Could not execute
generate-wka-override.xml\n";
}
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.override=em-coherence-override.xml";
}
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.$dsep."lib".$dsep."coherence.jar".$psep.$coherence
_home.$dsep."lib".$dsep."reporter.jar";
print "CLASSPATH: $CLASSPATH\n";
if($jamhost ne "" && $jamport ne "") {
$CLASSPATH=$CLASSPATH.$psep.$agent_
_home.$dsep."archives".$dsep."jlib".$dsep."jamagent.war";
my $jamjvmid="$cluster/$member";
print "Using Oracle JVMID - $jamjvmid\n";
$SYS_OPT="$SYS_OPT -Doracle.coherence.jamjvmid=$jamjvmid";
$SYS_OPT="jamconshost=$jamhost $SYS_OPT";
$SYS_OPT="jamconspport=$jamport $SYS_OPT";
$SYS_OPT="oracle.ad4j.groupidprop=$jamjvmid $SYS_OPT";
}
if ($bulkmbean ne "" && $jmxport ne "") {

```

```

# Management node with Bulk Operation MBean.
# add Oracle supplied jars for Bulk Operation MBean
$CLASSPATH=$CLASSPATH.$psep.$agent
_ home.$dsep."..".$dsep."..".$dsep."lib".$dsep."coherenceEMIntg.jar".$psep.$agent
_ home.$dsep."..".$dsep."..".$dsep."dependencies".$dsep."bulkoperationsmbean
_11.1.1.jar";
# Start MBeanServer
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=all";
print "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 $startup
_class";
} elsif ($jmxport ne "") {
# Management Node with out Bulk Operation MBean
# Start MBeanServer
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=all";
print "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";
} else {
# A simple managed node. Do not start MBeanServer.
$SYS_OPT="$SYS_OPT -Dtangosol.coherence.management=none";
print "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";
}
if ( !&IsWindows() ) {
if (fork() == 0) {
print "Executing start script from child process... $cmd \n";
exec("$cmd") or die "Could not execute $cmd\n";
}
} else {
print "Command used to start node = $cmd\n";
exec($cmd);
}
print "exiting default start script\n";
exit 0;
sub IsWindows {
$osname = $^O;
if ( $osname eq "Windows_NT"
|| $osname eq "MSWin32"
|| $osname eq "MSWin64" )
{
return 1;
}
else {
return 0;
}
}
}

```

generate-wka-override.pl

If Cluster Communication has been set to WKA in the Coherence Node Provisioning: Target Selection page, this script is launched by the default-start-script.pl. The generate-wka-override.pl is used to generate the override file. If you have your own override file, you can comment out the part that uses the generate-wka-override.pl script in the default-start-script.pl.

```

#!/usr/local/bin/perl
#
# $Header: emas/sysman/admin/scripts/coherence/generate-wka-override.pl /main/1
# 2011/02/01 16:51:33 $
#
# generate-wka-override.pl
#
# Copyright (c) 2011, 2011, Oracle and/or its affiliates. All rights reserved.
#
# NAME
# generate-wka-override.pl - <one-line expansion of the name>
#
# DESCRIPTION
# <short description of component this file declares/defines>
# expects input args as:
# host1:port1,host2:port2,host3:port3
# writes the wka information to em-coherence-override.xml
# Sample xml file:
#<coherence xml-override="/tangosol-coherence-override-{mode}.xml">
# <cluster-config>
# <unicast-listener>
# <well-known-addresses>
# <socket-address id="1">
# <address>10.232.129.69</address>
# <port>8088</port>
# </socket-address>
# <socket-address id="2">
# <address>10.232.129.69</address>
# <port>8089</port>
# </socket-address>
# </well-known-addresses>
# <port>8088</port>
# </unicast-listener>
# </cluster-config>
#</coherence>
#
use Cwd;
use IPC::Open3;
my $host_port = $ARGV[0];
@host_port_array = split(',', $host_port);
$size = @host_port_array;

my $xmlfile="em-coherence-override.xml";
print "$xmlfile\n";
open(XMLFL,"> $xmlfile");
print XMLFL "<coherence
xml-override=\"/tangosol-coherence-override-{mode}.xml\">\n";
print XMLFL "<cluster-config>\n";
print XMLFL "<unicast-listener>\n";
print XMLFL "<well-known-addresses>\n";

my $id = 1;
for($i = 0; $i < $size; $i++) {
    $single_host_port = $host_port_array[$i];
    $single_host_port =~ s/^\s+|\s+$//g;
    @single_host_port_array = split(':', $single_host_port);
    $wka_host = $single_host_port_array[0];
    $wka_port = $single_host_port_array[1];
    $id = $id + $i;
    print XMLFL "<socket-address id=\"$id\">\n";
    print XMLFL "<address>$wka_host</address>\n";

```

```
print XMLFL "<port>$wka_port</port>\n";
print XMLFL "</socket-address>\n";
}
print XMLFL "</well-known-addresses>\n";
print XMLFL "</unicast-listener>\n";
print XMLFL "</cluster-config>\n";
print XMLFL "</coherence>";
close(XMLFL);
exit 0;
```

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 Management 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.

Provisioning SOA Artifacts and Composites

This chapter explains how you can provision SOA Artifacts and Composites using Oracle Enterprise Manager Cloud Control (Cloud Control). In particular, this chapter covers the following:

- [Getting Started with SOA Artifacts Provisioning](#)
- [Understanding SOA Artifacts Provisioning](#)
- [Deployment Procedures, Supported Releases, and Core Components Deployed](#)
- [Provisioning SOA Artifacts](#)
- [Deploying SOA Composites](#)

Note:

In Enterprise Manager 12c Release 4, there is support to provision SOA 11g artifacts. Note that you cannot provision SOA 12c artifacts.

Getting Started with SOA Artifacts Provisioning

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 36-1 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 Cloud 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 Deployment Procedures, Supported Releases, and Core Components Deployed .

Table 36-1 (Cont.) Getting Started with Provisioning SOA Artifacts and Composites

Step	Description	Reference Links
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 Provisioning SOA Artifacts. - To learn about provisioning SOA Composites, see Deploying SOA Composites.
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, and 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 <i>Prerequisites</i> 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 <i>Provisioning Procedure</i> subsection.

Understanding SOA Artifacts Provisioning

SOA artifacts deployment procedures 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 (Human Workflow, Decision Service, and Oracle Mediator) that execute the business logic of their respective components within the SOA composite application (for example, a Human Workflow 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 de-

scribed 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.

Human Workflow

Human Workflow component is responsible for managing the lifecycle of human tasks, including creation, assignment, expiration, deadlines, and notifications, as well as its presentation to end users. It supports sophisticated dynamic task routing leveraging declarative patterns and tight integration with business rules. The three main sub-components of Human Workflow are a Task editor, Task Service Engine, and a Worklist application.

Oracle B2B

Oracle B2B provides the secure and reliable exchange of documents between businesses. For example, Retailer, Supplier, and Manufacturer. This type of eCommerce, B2B, represents mature business documents, classic business processes and industry specific messaging services and requires an architecture to manage the complete

end-to-end business process. Together with the Oracle SOA Suite, Oracle B2B meets this challenge and provides an architecture enabling a unified business process platform, end-to-end instance tracking, visibility, auditing, process intelligence, governance, and security.

Deployment Procedures, Supported Releases, and Core Components Deployed

Cloud Control offers the following Deployment Procedures for provisioning SOA Artifacts and Composites:

Deployment Procedure	Supported Releases	Artifacts Migrated
SOA Artifacts Provisioning	Oracle SOA Suite 11gR1 Patch Set 1 to Patch Set 4 (11.1.1.2.0) to 11.1.1.5.0)	• SOA Composites • Oracle WebLogic Server Policies • Assertion Templates • JPS Policy and Credential Stores • Human Workflow • Oracle B2B
Deploy SOA Composites	Oracle SOA Suite 11gR1 Patch Set 1 to Patch Set 4 (11.1.1.2.0) to 11.1.1.5.0)	• 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://docs.oracle.com/docs/cd/E14571_01/index.htm

Provisioning SOA Artifacts

This section describes how you can provision SOA artifacts. In particular, this section covers the following:

- [Provisioning SOA Artifacts from a Reference Installation](#)

- [Provisioning SOA Artifacts from Gold Image](#)

Provisioning SOA Artifacts from a Reference Installation

This section describes how you can provision SOA artifacts from one soa-infra domain to another.

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in [Provisioning SOA Artifacts and Composites](#).
- 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.

To provision SOA artifacts (composites, web service policies, JPS configuration) from a reference installation, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Deployment Procedure Manager page, in the Procedure Library subtab, from the table, select **SOA Artifacts Provisioning** deployment procedure. Select **Launch** and click **Go**. Cloud Control displays the Select Source page of the Deployment Procedure.



Note:

You can also access this deployment procedure as follows:

From the SOA Infrastructure Home page:

- a. From the Targets menu, click **Middleware**.
- b. In the Middleware page, click on a target of type **SOA Infrastructure**.
- c. In the SOA Infrastructure home page, from the SOA infrastructure-specific menu, select **SOA Artifacts Provisioning**.

3. On the Select Source page, 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**.

- 4. 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**.

- 5. 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**.

- 6. 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**.

7. 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**.
8. 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**.
To view a list of providers for the source and target, click **Provider details** link.
 - b. Click **Next**.
9. On the Human Workflow page, select all the workflow artifacts that you want to migrate such as **Views**, **Flex Field Mappings**, and **Attribute Labels**, then click **Next**.
10. On the B2B artifacts page, select all the B2B artifacts that you want to migrate such as **Trading Partners**, **Trading Agreements**, and **Document Protocols**, then click **Next**.
11. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.
12. On the Review page, review the details you have provided for provisioning SOA artifacts, and click **Submit**.

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:

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- 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.

To provision SOA artifacts from a gold image, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then click **Procedure Library**.
2. On the Deployment Procedure Manager page, in the Procedure Library subtab, from the table, select **SOA Artifacts Provisioning** deployment procedure. Select **Launch** and click **Go**. Cloud Control displays the Select Source page of the Deployment Procedure.

 Note:

You can also access this deployment procedure as follows:

- From the SOA Infrastructure Home page:
 - a. From the Targets menu, click **Middleware**.
 - b. In the Middleware page, click on a target of type **SOA Infrastructure**.
 - c. In the SOA Infrastructure home page, from the SOA infrastructure-specific menu, select **SOA Artifacts Provisioning**.

3. On the Select Source 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**.
4. 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**.
5. 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**.
6. 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**.

7. 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**.
8. 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 target, click **Provider details** link.
 - b. Click **Next**.
9. On the Human Workflow page, select all the workflow artifacts that you want to migrate like **Views**, **Flex Field Mappings**, and **Attribute Labels**, then click **Next**.
10. On the B2B artifacts page, select all the B2B artifacts that you want to migrate like **Trading Partners**, **Trading Agreements**, and **Document Protocols**, then click **Next**.
11. On the Review page, review the details you have provided for provisioning SOA artifacts, and click **Submit**.

Deploying SOA Composites

This section explains how you can deploy SOA composites. In particular, this section contains:

Before running the Deployment Procedure, meet the following prerequisites:

- Ensure that you meet the prerequisites described in [Provisioning SOA Artifacts and Composites](#).
- 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.

To provision SOA composites, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Deployment Procedure Manager page, in the Procedure Library subtab, from the table, select **Deploy SOA Composites** deployment procedure. Select **Launch** and click **Go**. Cloud Control displays the Select Destination page of the Deployment Procedure.
3. 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**.
4. 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**.
 5. On the Schedule page, schedule the Deployment Procedure to run either immediately or later.
 6. On the Review page, review the details you have provided for provisioning SOA composites, and click **Submit**.

Provisioning Service Bus Resources

Service Bus is an enterprise-class service bus that connects, manages, and mediates interactions between heterogeneous services. Service Bus accelerates service configuration, integration, and deployment, thus simplifying management of shared services across the Service-Oriented Architecture (SOA).

The resources of 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 Oracle Enterprise Manager Cloud Control (Cloud Control) allows you to discover and monitor these Service Bus targets, it also provides Deployment Procedures that help you provision Service Bus resources.

This chapter explains how you can provision Service Bus resources. In particular, this chapter covers the following:

- [Getting Started with Provisioning Service Bus Resources](#)
- [Supported Releases](#)
- [Provisioning Service Bus Resources from Service Bus Domain](#)
- [Provisioning Service Bus Resources from Oracle Software Library](#)

Getting Started with Provisioning Service Bus Resources

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

Table 37-1 Getting Started with Provisioning Service Bus Resources

Step	Description	Reference Links
Step 1	Knowing About The Supported Releases Know what releases of Service Bus can be provisioned by the Deployment Procedure.	To learn about the releases supported by the Deployment Procedure, see Supported Releases .

Table 37-1 (Cont.) Getting Started with Provisioning Service Bus Resources

Step	Description	Reference Links
Step 2	Selecting the Use Case This chapter covers a few use cases for provisioning Service Bus resources. Select the use case that best matches your requirements.	- To learn about provisioning Service Bus resources from the an Service Bus domain, see Provisioning Service Bus Resources from Service Bus Domain. - To learn about provisioning Service Bus resources from the Software Library, see Provisioning Service Bus Resources from Oracle Software Library.
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, and setting up of Oracle Software Library.	- To learn about the prerequisites for provisioning Service Bus resources from the an Service Bus domain, see Provisioning Service Bus Resources from Service Bus Domain. - To learn the prerequisites for provisioning Service Bus resources from the Software Library, see Provisioning Service Bus Resources from Oracle Software Library.
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Service Bus resources.	- To provision Service Bus resources from the an Service Bus domain, see Provisioning Service Bus Resources from Service Bus Domain. - To provision Service Bus resources from the Software Library, see Provisioning Service Bus Resources from Oracle Software Library.

Supported Releases

Using this Deployment Procedure, you can provision the resources for Service Bus 2.6, 2.6.1, 3.0, and 10gR3 (3.1).

Provisioning Service Bus Resources from Service Bus Domain

This section describes how you can provision Service Bus resources directly from an Service Bus domain.

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- Ensure that the source Service Bus (from where you want to export the resources) is already discovered and monitored in Cloud 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 [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 [Setting Up Credentials](#).

To provision Service Bus resources from a source Service Bus domain, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. From the Deployment Procedures section, select the Service Bus Resource Provisioning procedure from the list and click **Launch**.
3. On the Select Source page, in the Source section, select **Service Bus Domain**.
 - a. For **Domain**, click the torch icon and select the Service Bus domain from where the resources can be exported and deployed to a target 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**.
4. 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 Service Bus domain. The selected projects are exported to a JAR file, and the JAR file is moved to the host where the target Service Bus domain is running.

Note that the resources of the selected projects that exist in the target Service Bus 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 in [Understanding the Export Modes for Service Bus Resources](#).

- 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 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.

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 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 Service Bus domain.

Note that for 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 Service Bus 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 Service Bus 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 Service Bus domains.
- b. Specify the credentials of the hosts where the Management Agents, which are monitoring the administration servers of the Service Bus domains, are running

7. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only pre-

requisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed. Click **Next**.

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.
9. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Understanding the Export Modes for Service Bus Resources

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 37-2 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.

Provisioning Service Bus Resources from Oracle Software Library

This section describes how you can provision Service Bus resources from the Software Library.

Before running the Deployment Procedure, meet the following prerequisites:

Prerequisites for Designers

- Ensure that you meet the prerequisites described in [Setting Up Your Infrastructure](#).
- Export the resources of an Service Bus domain as a JAR file. Use 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 [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 [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 [Customizing Deployment Procedures](#).

To provision Service Bus resources from a source Service Bus domain, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. From the Deployment Procedures section, select the Service Bus Resource Provisioning procedure from the list and click **Launch**.
3. 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 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**. Cloud Control displays Select Target page.
 4. 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 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 Service Bus domain.

Note that for 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 Service Bus 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 Service Bus 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**.
 5. On the Set Credentials page, specify the following and click **Next**.
 - a. Specify the login credentials of the source and target Service Bus domains.
 - b. Specify the credentials of the hosts where the Management Agents, which are monitoring the administration servers of the Service Bus domains, are running.
 6. In the Schedule page, specify a Deployment Instance name. 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. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure to allow me to analyze results after performing prerequisite checks** to pause the procedure execution after all prerequisite checks are performed. Click **Next**.
 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.

8. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

Part VIII

Bare Metal Server Provisioning

This part contains the following chapter:

- [Provisioning Bare Metal Servers](#)

Provisioning Bare Metal Servers

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

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

Tip:

Before you begin provisioning of Linux on bare metal boxes, it is advisable to set preferred credentials for the Stage Server. If not, follow instructions in [Setting Up Credentials](#) to set up preferred credentials. If you want to use a reference host, set credentials for the reference host also. You can also set preferred credentials when configuring the deployment procedure for provisioning Linux.

Note:

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

Getting Started with Provisioning Bare Metal Servers

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

Table 38-1 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 Supported Releases of Linux.
Step 2	Knowing the Use Case This chapter covers provisioning Linux. Understand the use case for Linux provisioning.	To learn about provisioning bare metal boxes, see Provisioning Bare Metal Servers.
Step 3	Setting Up Infrastructure Before you perform Linux provisioning, you must meet the prerequisites, such as setting up of the provisioning environment, applying mandatory patches, and setting up of Oracle Software Library.	To learn about the prerequisites to be met for provisioning bare metal boxes, see Setting Up Infrastructure for Bare Metal Provisioning.
Step 4	Provisioning Linux Provision Linux on bare metal boxes.	To provision Linux on bare metal boxes, follow the steps explained in Provisioning Bare Metal Servers.

Overview Of Bare Metal Provisioning

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

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


This section covers the following:

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

Accessing Bare Metal Provisioning Page

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

Figure 38-1 Bare Metal Provisioning Home Page



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

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

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

Provisioning Environment for Bare Metals

The deployment environment in the data center needs to be setup in a certain manner in order to support the provisioning application. Besides the Oracle Management Serv-

er (OMS) which hosts Cloud Control and Provisioning Application, the following need to be setup and configured before using the provisioning application.

Software Library and its Entities

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

Boot Server

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

Stage Server

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

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

Reference Host

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

RPM Repository

The Provisioning application picks up the RPMs for the operating system from the RPM repository. At least one repository needs to be setup for use by the Provisioning application. From the networking perspective, you are advised to keep the RPM Repository as close to the target machines as possible. It will help in bringing down the installation time drastically by reducing the time taken to transfer RPMs from the RPM Repository to the hardware servers. If you have multiple hardware server groups residing at physically different locations, it would be better to have one RPM Repository for

each of these locations. Refer to section [Setting Up RPM Repository](#) for setting up a RPM repository. Also refer to section [Configuring RPM Repository](#).

Provisioning Bare Metal

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

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

Supported Releases of Linux

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

- SuSE Linux (SLES) 10
- Oracle VM Server (OVS) 3.0.x, Oracle VM Server 3.1.x, Oracle VM Server 3.2.x, Oracle VM Server 3.3.x, Oracle VM Server 3.4.x

Setting Up Infrastructure for Bare Metal Provisioning

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

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

Setting Up Stage Server

This section contains:

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

Prerequisites to Setup a Stage Server

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

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

Note:

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

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

Setting up a Stage Server and Accessing the Management Agent files

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

1. Set up an NFS Stage Server or a HTTP Server.

To setup an NFS Stage Server, see [Setting up an NFS Stage Server](#).

To set up a HTTP Server, see [Setting up a HTTP Stage Server](#).

2. Log in to the stage server running on the Management Agent, and create a top-level directory to store all the Management Agent installation files.

In this section, the variable `STAGE_TOP_LEVEL_DIRECTORY` is used to refer to the top level directory on the stage server.

For example:

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

Note:

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

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

For using the NFS Stage Server:

```
STAGE_TOP_LEVEL_DIRECTORY=/scratch/stage
```

For using the HTTP based Stage Server:

```
#STAGE_TOP_LEVEL_DIRECTORY=/var/www/html/stage
```

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

```
[root@upgps1t12 stage]# pwd
/scratch/stage
```

```
[root@upgps1t12 stage]# ls
1 10 21 6 9 oracle-agt-12.1.0.2.0-1.0.x86_64.rpm
```

Note:

1. If NFS is used then the staging process will automatically discover the agent rpm and there's no requirement for you to provide a URL for the rpm.
2. If HTTP is used then a URL will be required to reference the Agent rpm. The Agent URL is `http://host.example.com/agent_dir/oracle-agt-12.1.0.4.0-1.0.x86_64.rpm`. For more information on setting up HTTP Stage Server, see [Setting up a HTTP Stage Server](#).

Setting up an NFS Stage Server

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

Ensure that you perform the following steps on the stage server:

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


```
rpm --quiet -q nfs || yum -y install nfs
```
2. Run the following commands to configure NFS to export the stage server's top level directory (STAGE_TOP_LEVEL_DIRECTORY):


```
STAGE_TOP_LEVEL_DIRECTORY=/scratch/stage
echo "${STAGE_TOP_LEVEL_DIRECTORY}*(ro,sync)" >>/etc/exports
```
3. To reflect these changes on the NFS daemons, run the following command:


```
service nfs restart
```
4. Ensure NFS starts up on reboot, and is working now:


```
chkconfig nfs on
```
5. Install a Management Agent.

 See Also:

Oracle Enterprise Manager Cloud Control Basic Installation Guide to install a 12.1.0.1.0 or higher version of Management Agent.

Setting up a HTTP Stage Server

To setup a HTTP Stage Server, follow these steps:

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

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

2. Create a HTTP stage directory as follows:

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

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

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

4. Install a Management Agent.

 See Also:

Oracle Enterprise Manager Cloud Control Basic Installation Guide to install a 12.1.0.1.0 or higher version of Management Agent.

 Note:

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

Setting Up Boot Server and DHCP Server

 Note:

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

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

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

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

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

2. Configure the TFTP server:

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

3. Configure DHCP Server:

Edit the **dhcpcd.conf** (*/etc/dhcpcd.conf*) file. A sample **dhcpcd.conf** file for PXE set-up is shown below:

```
allow booting;
allow bootp;

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

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

group {

# PXE-specific configuration directives...

    next-server <TFTP_server_IP_address>;

    filename "linux-install/pxelinux.0";

    host <hostname> {
        hardware ethernet <MAC address>;
        fixed-address <IP address>;
    }
}
```

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

The *filename* option specifies the boot loader location on the TFTP server. The location of the file is relative to the main TFTP directory.

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

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

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

6. Install Oracle Management Agent. This step is not necessary if the DHCP and Boot servers are installed on the Cloud Control server.



Note:

See *Oracle Enterprise Manager Cloud Control Basic Installation Guide* to install a 12.1.0.1.0 or higher version of Management agent on the boot server.

Setting Up RPM Repository




Note:

It is recommended that you use RAM of 2 GB.

Setting Up RHEL 5/Oracle Linux 5 RPM Repository

Oracle Linux RPM repository should have the Install tree structure shown below:



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

1. Download Oracle Linux from <http://edelivery.oracle.com/linux>.
2. Copy all the contents of the first CD to a directory say Root Directory.

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

Exposing RPM Repository through HTTP or FTP

To expose RPM Repository through HTTP, follow these steps:

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

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


Configuring Stage Server

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

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

Follow these steps:

1. Log in to Cloud Control as an administrator.
2. From the **Enterprise** menu, select **Provisioning and Patching** and then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the Stage Servers section, click **Add Server**.




4. In the Add Staging Server dialog, select a **Stage Server**, specify a **Stage Directory**, for example, /scratch/stage, and **Base URL**, for example, file://stgserv.example.com/scratch/stage. Click **OK**.

Configuring Boot Server

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

Follow these steps:

1. Read about DHCP, PXE, and Redhat Kickstart technology before going through the boot server setup.
2. Ensure that you have administrator privileges.
3. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
4. In the Infrastructure tab, in the Boot Servers section, click **Add**.



5. In the Add Boot Server dialog, select a **Boot Server** and specify a **TFTP Boot Directory**, for example, /tftpboot/linux-install/. Click **OK**.

Configuring DHCP Server

Follow these steps:

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

Configuring RPM Repository

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

Follow these steps:

1. Ensure that you have administrator privileges.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
3. In the Infrastructure tab, in the RPM Repositories section, click **Add**.
4. In the Add RPM Repository Server dialog, specify a **Repository Name** and **URL**. For RPM repository either accessible by HTTP or on a local server, specify the URL in the HTTP format, for example, http://example.com/OEL5/. For NFS location, specify the URL as file://example/OEL5/.

Click **OK**.

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

Ensure that the following criteria are met before provisioning:

Table 38-2 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 present in the same subnet where the target machines to be provisioned are present or will be added. Management Agent is installed. Boot server machine is visible as a managed target in Cloud Control. A brand new PXE-bootable box actually detects the boot server and starts to boot it (even if no image is installed yet)
Stage Server	Large storage, High Memory and Sufficient Memory. If NAS server is used for storage then it should have NFS support. Management Agent is installed. Boot server machine is visible as a managed target in Cloud Control. The required agent rpm is staged for installing agents on targets. Preferred Credentials are set. Stage server is reachable from the box to be provisioned (or the same subnet)
RPM Repository	RPM Repository is as close as possible to the target servers. Install tree structure is as indicated in Configure RPM repository section. RPM repository is available via HTTP. Provide the exact URL and test the RPM repository access over HTTP
Reference Host	Agent is installed on local disk and not on NFS mounted directory. Preferred Credentials are set.
Software Library	Shared storage used for Software Library is accessible through NFS mount points to all OMS servers.

Configuring Software Library Components

To set up and configure the Software Library, see [Setting Up Oracle Software Library](#).

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

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

Creating Operating System Component

Follow these steps to create an operating system component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the Software Library Home, from the **Actions** menu, select **Create Folder**.
3. In the Create Folder popup, specify a **Name** and **Description** for the folder and select the folder location. For example, create a folder BMP-OL56 to represent the components you will use to provision a bare metal server of Oracle Linux 5 Update 6 Click **Save**.
4. From the **Actions** menu, select **Create Entity** and then **Bare Metal Provisioning Components**.
5. In the Create Entity: Bare Metal Provisioning Components dialog box, select **Operating System Component** and click **Continue**.
6. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

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

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

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

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

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

Click **Next**.

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

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

Click **Next**.

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

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

Table 38-3 Agent Settings

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

Table 38-4 Additional OS Configuration

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

Table 38-5 Boot Configuration and Configuration Scripts

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

Creating Disk Layout Component

Follow these steps to create a disk layout component:

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

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

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

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

7. In the Configure page, specify the hard disk, RAID, partition, and logical configurations.

To specify the hard disk profile, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**.

To specify the Partition Configuration, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.

To specify RAID configuration, click **Add**. Specify the **Device Name** and **Capacity**.

To specify the Logical Volume Group Configuration, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.

To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

Click **Next**.

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

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

Creating an Oracle Virtual Server Component

Follow these steps to create an operating system component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the Software Library Home, from the **Actions** menu, select **Create Folder**.
3. In the Create Folder popup, specify a **Name** and **Description** for the folder and select the folder location.

4. From the **Actions** menu, select **Create Entity** and then **Bare Metal Provisioning Components**.
5. In the Create Entity: Bare Metal Provisioning Components dialog box, select **Oracle Virtual Server Component** and click **Continue**.

6. On the Describe page, enter the **Name**, **Description**, and **Other Attributes** that describe the entity.

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

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

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

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

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

Click **Next**.

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

Click **Next**.

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

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

Table 38-6 Additional OS Details

Element	Description
Mount Point Settings	Specify entries for the /etc/fstab file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
NIS Settings	Specify entries for the /etc/yp.conf file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
NTP Settings	Specify entries for the /etc/ntp.conf file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for/etc/inittab file. All processes are started as part init operation in boot process. Init operation decides the processes that will start on booting of a machine or when runlevel changes.

Table 38-6 (Cont.) Additional OS Details

Element	Description
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 38-7 Boot Configuration and Configuration Scripts

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

Prerequisites For Provisioning Bare Metal Servers and Oracle VM Servers

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

Provisioning Bare Metal Servers

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

Follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In the Server Image section, from the **Provision** menu, select **Operating System**.
3. In the General/Target Selection page, in the General section, specify the **Deployment Name**. Select the **Operating System** you want to provision and provide a description. Select the **Patching Groups** and **Monitoring Templates** you want to associate with the system.

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

- **MAC Addresses** if you want to provision the bare metal systems by specifying MAC addresses. Click **Add** to specify the list of MAC Address. In the Add MAC dialog box, specify the MAC addresses. Click **OK**.

Optionally, click **Add from File** to add the MAC address from a file. In the Add from File dialog box, click **Browse** and select the file from the location where you have stored it.

- **Subnet** to specify the subnet for the bare metal provisioning. In the Subnet to be Provisioned section, specify the **Subnet IP**, **Netmask**, **Number of Network Interfaces**, and **Bootable Network Interface**.
- **Re-image EM Host Targets** to re-provision an existing Cloud Control host target. In the Enterprise Manager Hosts to be Provisioned section, click **Add** to search and select the host target. Click **OK**. Select the **Bootable Network Interface**.

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

Click **Next**.

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

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

Click **Next**.

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

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

You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password and select the Run Privilege. Choose to **Save Credentials** to use these credentials in future deployments.

Click **Next**.

6. In the Additional OS Details page, specify agent settings, configuration scripts, package selection, and additional operating system configuration, and boot configuration as explained in [Table 38-8](#), [Table 38-9](#), and [Table 38-10](#).

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

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

Click **Next**.

7. In the Disk Layout page, specify hard disk profile, partition configuration, RAID configuration, Logical Volume Group configuration, and Logical Volume configuration.

To specify the hard disk profile, click **Add**. Specify the **Device Name** and **Capacity**.

To specify the Partition Configuration, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.

To specify RAID Configuration, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**. To configure RAID, ensure that your hard disk has two partitions at the minimum.

To specify the Logical Volume Group Configuration, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.

To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

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

Click **Next**.

8. In the Network page, the network properties for the MAC Address or Subnet as specified during target selection, is displayed.

Click **Add** to configure the network interfaces. In the Input Network Interface Properties dialog box, specify the Interface name. Select the **Configuration Type** as:

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

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

Click **Next**.

9. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Specify the Stage Server and Boot Server credentials. You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password and select the Run Privilege. Choose to **Save Credentials** to use these credentials in future deployments.

Click **Next**.

10. In the Review page, verify that the details you have selected are correctly displayed and submit the job for the deployment. If you want to modify the details, click **Back** repeatedly to reach the page where you want to make the changes. Click **Save As Plan** to save the configuration details you have specified. Specify a

name and description and click **OK** to save the plan. You can later use the saved plan to provision bare metal boxes. For more information, see [Using Saved Plans for Provisioning Linux Operating Systems on Bare Metal Servers](#)

Click **Submit**.

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

Table 38-8 Agent Settings

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

Table 38-9 Additional OS Configuration

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

Table 38-10 Boot Configuration and Configuration Scripts

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

**Note:**

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

Provisioning Oracle VM Servers

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

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In the Server Image section, from the **Provision** menu, select **Oracle VM Server**.
3. In the General/Target Selection page, in the General section, specify a unique **Deployment Name**.

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

- **MAC Addresses** if you want to provision the bare metal systems by specifying MAC addresses. Click **Add** to specify the list of MAC Address. Alternately, to add the addresses from a file, click **Add from File**. In the Add from File dialog box, select the file that contains the addresses and click **OK**.
- **Subnet** to specify the subnet for the bare metal provisioning. In the Subnet to be Provisioned section, specify the **Subnet IP**, **Netmask**, **Number of Network Interfaces**, and **Bootable Network Interface**.

The Oracle VM Registration section allows you to select an OVM Manager registered in cloud to manage the Oracle VM servers you are provisioning. To do so, click the search icon. From the Select Target dialog box, select a target VM machine, and click **Select**.

Click **Next**.

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

- c. Select DHCP Server, and select **Run DHCP Server Pre-requisite checks** to check if the DHCP server is configured properly
- d. Select **Local RPM Repository** from the available list.

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

Click **Next**.

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

Click **Next**.

- 6. In the Additional OS Details page, do the following:

In the Dom0 Configuration section, you can provide the memory and power requirements for the target provisioned.

In the Additional OS details, you can click the configure icon to add certain other configuration details such as: **Mount Point Settings**, **NIS Settings**, **NTP Settings**, **Kernel Parameter Settings**, **Inittab Settings**, **Firewall Settings**.

In the Boot Configuration section, click **First Boot** to add commands/scripts that must be run on the system when it boots for the first time after installation. Click **Post Install** to provide commands to run on the system once the installation is complete.

Click **Next**.

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

To specify the hard disk profile, click **Add**. Specify the **Device Name** and **Capacity**.

To specify the Partition Configuration, click **Add**. Specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.

To specify RAID Configuration, click **Add**. Specify the **Mount Point**, **RAID Level**, **Partitions**, and **File System Type**. To configure RAID, ensure that your hard disk has two partitions at the minimum.

To specify the Logical Volume Group Configuration, click **Add**. Specify the **Group Name**, **Partitions**, and **RAIDs**.

To specify the Logical Volume Configuration, click **Add**. Specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.

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

Click **Next**.

8. In the Network page, the network properties for the MAC Address or Subnet as specified during target selection, is displayed.

Click **Add** to configure the network interfaces. In the Add Network Interface dialog box, specify the Interface name. Select the **Configuration Type** as:

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

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

Click **Next**.

9. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Specify the Stage Server and Boot Server credentials. You can choose to use preferred credentials, named credentials, or enter your credentials. If you choose to enter your credentials, specify the user name and password and select the Run Privilege. Choose to **Save Credentials** to use these credentials in future deployments.

Click **Next**.

10. In the Review page, verify that the details you have selected are correctly displayed and submit the job for the deployment. If you want to modify the details, click **Back** repeatedly to reach the page where you want to make the changes.

Click **Submit**.

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

Table 38-11 Additional OS Configuration

Element	Description
Mount Point Settings	Specify entries for the /etc/fstab file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
NIS Settings	Specify entries for the /etc/yp.conf file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
NTP Settings	Specify entries for the /etc/ntp.conf file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
Kernel Parameter Settings	Specify scripts for Kernel Parameters.
Initab Settings	Specify settings for/etc/inittab file. All processes are started as part init operation in boot process. Init operation decides the processes that will start on booting of a machine or when runlevel changes.

Table 38-11 (Cont.) Additional OS Configuration

Element	Description
Firewall Settings	Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Make sure that the port used by Management Agent is open for its communication with the Management Service. By default, Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.

Table 38-12 Boot Configuration and Configuration Scripts

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

Viewing Saved Plans

To view saved plans, follow these steps:

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



Note:

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

Using Saved Plans for Provisioning Linux Operating Systems on Bare Metal Servers

To edit the saved plans, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Bare Metal Provisioning**.
2. In Server Image section, from Provision menu, select Using Saved Plan.
3. From the Saved Plans dialog box, select any template to pre-populate the provisioning wizard with the saved values, and click **Continue**.
4. Update the Deployment Name, the Provisioning Category information in the General/Target selection page.
5. Follow steps 4 to step 10 listed in the section [Provisioning Bare Metal Servers](#).

6. In the Schedule/Credentials page, provide a schedule for the job, either immediately or at a later date. Also, update the Stage Server and Boot Server credentials.
7. In the Review page, verify all the details you have selected, and click **Submit**.

Part IX

Host Management

This part includes the following chapters:

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

Overview of Host Management

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

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

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

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

 Note:

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

Host Statistics

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

Using the host UI, you can:

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

Diagnosing Host Problems

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

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

Viewing Targets on the Host

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

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

When working with targets on a host:

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

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

When removing multiple targets:

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

Storage Statistics and History

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

By default the storage history feature is not activated. Enabling storage history is expensive in regards to database resources. The amount of database resources used to

calculate history data depends on the amount of storage data associated with the host target.

Setting Up the Environment to Monitor Hosts

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

- [Required Installations](#)
- [For Linux Hosts - Installing YAST \(Only if the host is not patched with latest 13.2\)](#)
- [Setting Up Credentials](#)
- [Setup Needed for Host Monitoring](#)
- [Target Setup Needed for Host Administration](#)

 Note:

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

Required Installations

Note: These required installations are only applicable to hosts running Oracle Linux, Red Hat Linux, and SUSE Linux Operating Systems (x86 and x64 architectures only).

To administer a host through Enterprise Manager, you need to install scripts. To determine which scripts you need to install for your host, follow these steps:

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

 Note:

The credential user must have access to the Agent installation.

YaST and EM wrapper scripts must be installed for agents running on earlier versions like Oracle Linux 5, Oracle Linux 6, and SuSE Enterprise Linux 9 and 10. The scripts can be downloaded from <http://oss.oracle.com/projects/yast>

For Linux Hosts - Installing YAST (Only if the host is not patched with latest 13.2)

YAST is an operating system setup and configuration tool that comes as a standard tool as part of SUSE Linux distribution. The Linux administration feature uses YAST to run scripts. For Oracle Linux and RHEL4 (Red Hat), YAST rpm contains the Enterprise Manager scripts. Therefore installing YAST rpm from the following location will also install the Enterprise Manager scripts:

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

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

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

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

1. Determine the version of Linux on your machine. For example, the `uname -a` command lists the RHEL (RedHat), Oracle Linux, or SUSE versions, and the bits of the machine, for example 32-bit versus 64-bit. YAST is supported on RHEL4 and later and on Oracle Linux.
2. Verify that you have root privileges.

To install YAST, perform the following steps:

1. Go to <http://oss.oracle.com/projects/yast>. The Project: Yast page appears.
2. Click the 'here' link. The Project Downloads: Yast page appears. Click the link that coincides with your version of Linux, for example, EL5.
3. On the EL5 page, click the link associated with the bits on your machine, either i386 for 32 bits or x86-64 for 64 bits.
4. Click `yast_el5_x86_64.tar` and download the tar file.
5. Once the tar is downloaded, go to the directory where the tar file is available.
6. untar the file using `tar -xvf yast_el5_x86_64.tar`
7. `cd` to the `yast_el5_x86_64` directory.
8. Type `sudo ./install.sh`
9. To verify that YAST is installed, type `/sbin/yast2`. This should display the YAST control center. If it does not, the YAST installation has failed.
10. When you return to the Administration menu, the options should now display the available Linux administration features.

For a demonstration of how to install YAST, see the YouTube video *Oracle Enterprise Manager 12c: Install YAST* located at <http://www.youtube.com/watch?v=7ZiwmxZV-mAw>.

Setting Up Credentials

Credentials are needed to manage target instances.

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

- **Named Credentials** are used for the Management Agent install. Named credentials explicitly grant you privileges on the host.
- **Preferred Credentials**

If a target has preferred credentials set, applications that log in to that target will automatically use the preferred credentials. Using preferred credentials simplifies access to managed targets.

Default credentials can be set for each target type. Default credentials are used for any targets that do not have preferred credentials explicitly set.

- **Privilege Delegation Setting** enables you to configure the Management Agent to use Sudo or PowerBroker so you can run privileged scripts.

See the online help for additional information.

Setup Needed for Host Monitoring

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

This section explains the required steps for these tasks.

Viewing Monitoring Configuration

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

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

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

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

Setting Up Monitoring Credentials

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

To edit Monitoring Credentials, perform the following steps:

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

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

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

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

Target Setup Needed for Host Administration

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

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

Customizing Your Host Monitoring Environment

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

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

Customizing the Host Home Page

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

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

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

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

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

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

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

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

Using Groups

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

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

To create a group, follow these steps:

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

The Enterprise Manager Console displays a set of Create Group pages that function similarly to a wizard. See the online help for a description of the fields on each page.

3. On the General tab of the Create Group page or the Create Dynamic Group page, enter the **Name** of the Group you want to create. If you want to make this a privilege propagating group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target privileges granted on the group to an administrator (or a role) are propagated to the member targets.

Note: The Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner *must* be able to view any target that can potentially match the membership criteria.

4. Configure each page, then click **OK**. You should configure all the pages before clicking **OK**.

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

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

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

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

Monitoring Hosts

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

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

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

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

Overall Monitoring

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

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

CPU Details

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

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

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

Memory Details

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

Disk Details

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

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

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

Program Resource Utilization

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

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

Log File Alerts

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

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

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

Metric Collection Errors

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

Storage Details

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

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

- Summary attributes presented are rolled up for one or multiple associated hosts.
- A host is associated with a group either through:
 - Explicit membership, or
 - Implicit *hosted by* association which is inherited through a group member target

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

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

Storage Utilization

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

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

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

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

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

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

Overall Utilization

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

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

Provisioning Summary

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

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

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

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

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

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

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

Consumption Summary

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

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

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

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

ASM

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

Databases

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

Disks

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

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

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

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

File Systems

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

Local File Systems

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

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

NFS

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

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

Suggestions for Monitoring NFS Mounts

The following are suggestions on monitoring NFS mounts.

- Monitor the remote host if NFS exports are coming from another host supported by Enterprise Manager. The Filesystems metric will monitor the local file systems on the remote host.
- Monitor the Netapp Filer if NFS exports are coming from a remote Netapp Filer. Volumes and Qtress metrics will monitor the exports from the remote Netapp Filer.
- Use the 'File and Directory Monitoring' metric if any of the previous choices do not meet the need. Set the threshold against the 'File or Directory Size' metric to monitor specific remote mounts.

Volumes

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

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

Types of Entities

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

Top-Level Entities

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

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

Bottom-Level Entities

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

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

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

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

Intermediate Entities

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

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

The following vendor products are instrumented:

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

Vendor Distribution

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

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

Storage History

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

In the case of a group, history is not enabled by default. The user interface allows you to enable or disable the history for each group. Computation of history for a group consumes resources in the Enterprise Manager Repository database. It is not anticipated

that a given deployment would find it useful to have the history for all instances of groups, so the control is given to you to choose for which groups it is worth keeping the history.

Storage Layers

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

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

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

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

Storage Refresh

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

Storage Refresh in Context of Host Target

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

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

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

Storage Refresh in Context of Group Target

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

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

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

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

Ksplice for Oracle Linux

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

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

Administering Hosts

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

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

Configuration Operations on Hosts

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

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

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

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

Follow the instructions for each configuration explanation.

Configuring File and Directory Monitoring Criteria

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

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

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

1. On the Metric and Collection Settings page, select **All metrics** in the View menu. Locate the File and Directory Monitoring metrics. The metrics are:

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

Notes about Specifying Monitored Object

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

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

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

Configuring Generic Log File Monitor Criteria

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

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

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

Perform the following steps using the Enterprise Manager console:

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

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

1. By default, matching number of lines is reported through log file monitoring. To enable upload of matching content for a specific file, add the absolute path for the file to the \$ORACLE_HOME/sysman/admin/lfm_ifiles file. The \$ORACLE_HOME/sysman/

admin/lfm_ifiles.template file is a template needed for creating the \$ORACLE_HOME/sysman/admin/lfm_ifiles file.

2. For security purposes, you may want to disable monitoring of sensitive files by Enterprise Manager permanently by adding the names of the sensitive files to the \$ORACLE_HOME/sysman/admin/lfm_efiles file. The \$ORACLE_HOME/sysman/admin/lfm_efiles.template file is a template needed for creating the \$ORACLE_HOME/sysman/admin/lfm_efiles file.

Notes about Specifying Monitored Object

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

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

Configuring Program Resource Utilization Monitoring Criteria

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

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

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

1. On the Metric and Collection Settings page, select **All metrics** in the View menu. Locate the Program Resource Utilization metrics. The metrics are:
 - Program's Max CPU Time Accumulated (Minutes)
 - Program's Max CPU Utilization (%)
 - Program's Max Process Count
 - Program's Max Resident Memory (MB)
 - Program's Min Process Count
 - Program's Total CPU Time Accumulated (Minutes)
 - Program's Total CPU Utilization (%)
2. After reviewing each metric, decide which metrics need to change. Click the pencil icon to navigate to the corresponding Edit Advanced Settings page.
3. You can edit or remove existing criteria by selecting the row from the Monitored Objects table and clicking **Edit** or **Remove**. Refer to Notes about Specifying Monitored Objects for details on configuring the criteria.

Notes about Specifying Monitored Object

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

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

Administration Tasks

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

Using the Administration tab, you can manage:

- Services

View statistics of individual service and edit their services.

Note: This feature is only available on hosts running Oracle Linux, Red Hat Linux and SUSE Linux Operating Systems (x86 and x64 architectures only).

- Default System Run Level
- Network Cards

Manage routing configuration, view configuration statistics, and view network file system clients.

- Host Lookup Table
- NFS Client
- User and Group Administration

Manage user and group settings.

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

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

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

Note: To see a video showing the navigation in the Administration menu option, see http://www.youtube.com/watch?v=R0fqR2GhQ_E.

Services

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

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

Run Level	Description
0	System halt
1	Single user mode Only one user can be logged on at any point in time. Additional users will not be allowed to log on until the user using the system logs off.
2	Basic multiuser mode without network
3	Full multiuser mode with network

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

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

Default System Run Level

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

Run Level	Description
0	System halted
1	Single user mode
2	Basic multiuser mode without network
3	Full multiuser mode with network
4	Unused. This run level is for future Oracle use.
5	Full multiuser mode with network and X display manager
6	Reboot the system.

Note the following:

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

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



Note:

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

Network Card

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

Using the Network Card page, you can:

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

Notes

Note the following:

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

Configuring the Network Card

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

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

Adding Routing Configuration

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

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

Host Lookup Table

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

- Edit hostname and aliases.
- Delete a lookup table entry for a host.
- Add a lookup table entry for a host by accessing the Add Host Configuration page.

- Refresh updates the **Host Lookup Table** page with latest host Look up data from host. This is helpful to fetch updates which are done directly on host, instead through Enterprise Manager.

Note the following:

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

NFS Client

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

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

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

Note the following:

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

Adding and Editing an NFS Client

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

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

Options

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

Note the following:

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

User and Group Administration (Users)

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

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

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

Adding or Editing a Local User

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

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

When editing a local user, you can:

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

User and Group Administration (Groups)

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

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

Group Statistic	Description
Group Name	Name of the group.
Group ID	Group identifier. This identifier is unique to the group.

Group Statistic	Description
Group Members	Groups that belong to the group. Group shares the permissions given to the subordinate groups.

Note: This feature is only available on Linux.

Adding or Editing a Local Group

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

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

When editing a local group, you can:

- Change the group ID
- Add, delete, or change group members

Using Tools and Commands

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

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

Enabling Sudo and Power Broker

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

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

To enable Sudo or PowerBroker, perform these steps:

1. From the **Setup** menu located at the top-right of the page, select **Security** then select **Named Credentials**.
2. On the Named Credentials page, click **Create**.

3. On the Create Credential page, provide host credentials with root privileges.
4. In the Credential Properties section, select **Sudo** or **PowerBroker** from **Run Privilege**.
5. Provide the details for Sudo or PowerBroker and the system performs the administrative task.

Using Sudo or Power Broker

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

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

Executing the Host Command Using Sudo or PowerBroker

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

1. Navigate to the Host target page.
From the **Targets** menu, select **Hosts**. On the Hosts page, highlight the row containing the name of the host in which you are interested.
2. Click the **Run Host Command** button.
3. Provide host target credentials. Provide information for Sudo (runas) support or PowerBroker (profile) support.
Note: On the target host, the `/etc/sudoers` file needs to be present with the target user information inserted.
4. Type the specific command to be run on the system and view the command output.

Using Remote File Editor

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

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

Accessing Remote File Editor

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

1. From the **Targets** menu, select **Hosts**. On the Hosts page, click the name of the host in which you are interested.
2. On the resulting Host home page, select **Remote File Editor** from the Host menu (located at the top-left of the page).
3. If the preferred credentials are not set for the host target, the Host Credentials page appears. Three options are available: Preferred, Named, and New.

You must have permissions to perform operations on a target host.

4. Once credentials are set, you can perform the following on the Remote File Editor page:
 - Perform operations on files, for example, listing of files in a directory, opening a file for reading, editing, and saving.
 - Provide host target credentials.
 - Provide details of the file and type of operation to be done.

Note the following:

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

Adding Host Targets

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

Detailed information is available in:

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

Running Host Command

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

Accessing Host Command

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

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

Executing Host Command Using Sudo or Power Broker

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

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

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

Execute Host Command - Multiple Hosts

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

On this page, you can:

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

- Add targets and modify the targets list.

Note the following:

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

Target Properties

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

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

Name	Description	Target Type
%MachineName%	Listener Machine Name	Database Instance

Note the following:

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

Examples:

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

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

Execute Host Command - Group

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

On this page, you can:

- Choose the target type. You can choose hosts directly, or choose hosts by way of database instance targets.
- Refine the command, re-execute the command, and view the execution results, all without leaving the page.
- Either type operating system commands, load the commands from a script on the browser machine or on the host, or load host commands from a job defined in the job library.
- Select hosts individually or through the use of a group. You can also switch to Single Target Mode where only one host target is acted upon.
- Interactively view command execution results or hide the results to be viewed at a later time.
- Use preferred credentials or override preferred credentials.
- Add targets and modify the targets list.

Note the following:

- Re-execute the command by clicking **Execute**.
- Cancel execution of the command by clicking **Cancel** on the Processing: Executing Host Command page.
- If the current target type is Host, clicking **Add** launches the Target Selection page with the target type list limited to host targets and any groups that contain host targets.

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

Execute Host Command - Single Host

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

On this page, you can:

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

Note the following:

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

Load OS Script

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

On this page, you can:

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

Load From Job Library

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

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

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

Execution History

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

On this page, you can:

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

Execution Results

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

On this page, you can:

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

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

Miscellaneous Tasks

This section explains the following:

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

- [Adding and Editing Host Configuration](#)

Enabling Collection of WBEM Fetchlet Based Metrics

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

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

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

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

Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts

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

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

Note: This link was accurate at time of writing. If you find this link to be out of date, contact Dell support.

To identify the latest OMSA software, perform the following steps:

- Search for the latest `om??_lnx_managed_system*.gz` file on the website
- As of June 13, 2004, the latest version is:

http://ftp.dell.com/sysman/om38_lnx_managed_system_A01.tar.gz

2. Install the software using the instructions provided by Dell.
3. Verify that the installation was successful by performing the following steps:
 - Verify that snmp daemon is up and running.

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

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

```
% /usr/bin/omreport system version
```
4. Verify the Dell OMSA software is functioning correctly by verifying that the Dell OpenManage Server Administrator website is up and running.
 - Using your web browser, access the URL - `https://target_hostname:1311`.

- Log in using an operating system account. Check that you are able to successfully log in and navigate in the website.
5. If the SNMP Community String for the SNMP daemon running on the Linux host is not *public*, set the SNMP Community String property in Enterprise Manager using the following steps:
 - a. Login into Enterprise Manager Cloud Control.
 - b. Navigate to the home page of the specific host.
 - c. From the **Host** menu, select **Target Setup**, then **Monitoring Configuration**.
 - d. Set the SNMP Community String property to the correct value on this page.
 6. Restart the Management Agent on the host.
% login into host as the owner account of emagent software
% emctl stop agent
% emctl start agent
 7. **Note:** The following step is not required if the Dell OMSA software was functional prior to the previous startup of the Management Agent.
Verify that hardware monitoring is working correctly using the following steps:
 - a. Log in to Enterprise Manager Cloud Control.
 - b. Navigate to the home page of the specific host.
 - c. From the **Host** menu, select **Monitoring**, then **All Metrics**.
 - d. Verify that the following metrics are present on this page: Fans, Memory Devices, PCI Devices, Power Supplies, Processors, Remote Access Card, System BIOS, and Temperature.
 - e. You can navigate to the metric data page by clicking on one of the metrics listed in the previous step and view the data that Enterprise Manager is able to fetch.

Adding and Editing Host Configuration

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

On the Add Host Configuration page, you can:

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

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

When editing a host configuration, you can:

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

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

Part X

Patch Management

This part contains the following chapters:

- [Patching Software Deployments](#)
- [Patching Linux Hosts](#)
- [Oracle Linux Home](#)
- [Database Fleet Maintenance](#)
- [Performing Engineered System Software Updates](#)

Patching Software Deployments

Patching is one of the important phases of the product lifecycle that enables you to keep your software product updated with bug fixes. Oracle releases several types of patches periodically to help you maintain your product. Patching is complex and involves downtime. You can use several approaches to identify and apply the patches.

This chapter describes how Oracle Enterprise Manager Cloud Control's (Cloud Control) patch management solution addresses these patch management challenges. In particular, this chapter covers the following:

- [Overview of the New Patch Management Solution](#)
- [Setting Up the Infrastructure for Patching](#)
- [Identifying the Patches to Be Applied](#)
- [Applying Patches](#)
- [Diagnosing and Resolving Patching Issues](#)
- [Additional Patching Tasks You Can Perform](#)
- [End-to-End Use Case: Patching Your Data Center](#)
- [Patching Database as a Service Pools](#)

Overview of the New Patch Management Solution

This section describes the following:

- [Overview of the Current Patch Management Challenges](#)
- [About the New Patch Management Solution](#)
- [Overview of Patch Plans](#)
- [Overview of Patch Templates](#)
- [Supported Targets, Releases, and Deployment Procedures for Patching](#)
- [Overview of Supported Patching Modes](#)
- [Understanding the Patching Workflow](#)

Overview of the Current Patch Management Challenges

Before you understand the new patch management solution offered by Cloud Control, take a moment to review some of the tools you might be using currently to patch your databases, other Oracle software, and the challenges you might be facing while using them ([Table 44-1](#)).

Table 44-1 Current Patch Management Tools and Challenges

Approach	Description	Challenges
OPatch	Oracle proprietary tool that is installed with Oracle products like Oracle Database, Management Agent, SOA, and so on.	• Difficult to identify the patches to be rolled out • Patches only one Oracle home at a time • Offers limited support to handle pre and post-patching scripts
Custom Scripts	User-created scripts developed around OPatch, SQLPlus, and so on.	• Difficult to identify the patches to be rolled out • Can be used only on a single server • Requires significant maintenance overhead to meet the new version and configuration needs
Deployment procedures	Default procedures offered by Cloud Control for automating the patching operations	• Confusion over which deployment procedure to select • Limited scope for validating the patches and targets selected in a deployment procedure • Separate deployment procedures for patching in rolling and parallel mode • Difficult to handle patch conflicts

About the New Patch Management Solution

Cloud Control addresses the challenges described in [Overview of the Current Patch Management Challenges](#) with its much-improved patch management solution that delivers maximum ease with minimum downtime. The new patch management solution offers the following benefits:

- Integrated patching workflow with My Oracle Support (MOS), therefore, you see recommendations, search patches, and roll out patches all using the same user interface.
However, Cloud Control does not upload any data to MOS. It only uses MOS to download the latest updates.
- Complete, end-to-end orchestration of patching workflow using *Patch Plans*, including analysis of the patch conflicts, therefore, there is minimal manual effort required. For more information on patch plans, see [Overview of Patch Plans](#).
- Clear division of responsibilities between designers and operators. Designers can focus on creating patch plans, testing them on a test system, and saving them as patch templates. Operators can focus on creating patch plans out of the template for rolling out the patches on a production system.
- Easy review of patches for applicability in your environment, validation of patch plans, and automatic receipt of patches to resolve validation issues.
- Saving successfully analyzed or deployable patch plans as patch templates, which contain a predetermined set of patches and deployment options saved from the source patch plan.
- Out-of-place patching for standalone (single-instance) database, Oracle Restart, Oracle Real Application Clusters (RAC), Oracle Grid Infrastructure (that may or may not be a part of Oracle Exadata) targets, and targets in Oracle Data Guard configuration.


 Note:

- Out-of-place patching is supported for RAC targets and Oracle Grid Infrastructure targets that are not a part of Oracle Exadata, only if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed.
- Out-of-place patching is supported for Oracle Data Guard targets, only if you have the 12.1.0.6 Enterprise Manager for Oracle Database plug-in deployed.
- Out-of-place patching is supported by Enterprise Manager for Oracle Restart 11.2.0.4 and 12.1.0.2 on Linux x64 and AIX platforms only if you have the 13.1.0.1 Enterprise Manager for Oracle Database plug-in deployed.

- Flexible patching options such as rolling and parallel, both in offline and online mode.

Figure 44-1 shows you how you can access the Patches and Updates screen from within Cloud Control console.

Figure 44-1 Accessing the Patches & Updates Screen



The screenshot displays the Oracle Enterprise Manager Cloud Control console. The top navigation bar includes 'Enterprise', 'Targets', 'Favorites', and 'History'. The 'Enterprise' menu is expanded, showing options like 'Summary', 'Monitoring', 'Job', 'Reports', 'Configuration', 'Compliance', 'Provisioning and Patching', 'Quality Management', 'My Oracle Support', and 'Cloud'. The 'Provisioning and Patching' option is highlighted, and a sub-menu is visible with 'Patches & Updates' selected. A green arrow points from this menu item to the 'Patches & Updates' screen below.

The 'Patches & Updates' screen is divided into several sections:

- Patching Quick Links:** Includes links for 'Target Patchability Report', 'Software and Patch Search Sites', 'Oracle E-Business Suite', 'Oracle Server and Tools', and 'Patch Recommendations'.
- Patch Search:** Features a search bar with filters for 'Number/Name or Bug Number (Simple)', 'Product or Family (Advanced)', and 'Recommended Patch Advisor'. It also includes a 'Search' button and a 'Clear' button.
- Plans:** A table listing patch plans with columns for Name, Status, Target Type, Type, Created By, and Last Updated Date.

Name	Status	Target Type	Type	Created By	Last Updated Date
a1	Deployed Successfully	Database Ins...	Patch	SYSMAN	Nov 14, 2015 9:57 PM
a2	Review Analysis	Database Ins...	Patch	SYSMAN	Nov 15, 2015 2:25 PM
db1	New	Database Ins...	Patch	SYSMAN	Nov 10, 2015 8:31 PM
db1203	Review Analysis	Database Ins...	Patch	SYSMAN	Nov 22, 2015 3:56 PM
GL_RAC_1	Analysis Scheduled	Cluster	Patch	SYSMAN	Nov 12, 2015 6:25 PM
temp1	Not Applicable	Database Ins...	Patch Template	SYSMAN	Nov 22, 2015 3:49 PM

Overview of Patch Plans

This section describes the following:

- [About Patch Plans](#)
- [About Types of Patch Plans](#)
- [About the Create Plan Wizard](#)

About Patch Plans

Patch plans help you create a consolidated list of patches you want to apply as a group to one or more targets. Patch plans have states (or status) that map to key steps in the configuration change management process. Any administrator or role that has view privileges can access a patch plan.

Patch plan supports the following types of patches:

- Patches (One-Off)
 - Interim Patches that contain a single bug fix or a collection of bug fixes provided as required. It can also include customer specific security bug fixes.
 - Diagnostic Patches, intended to help diagnose or verify a fix or a collection of bug fixes.
 - Patch Set Updates (PSU), contain a collection of high impact, low risk, and proven fixes for a specific product or component.
 - Critical Patch Updates (CPU), contain a collection of security bug fixes.
- Patch Sets

 Note:

- Patch Sets are available for Oracle Database 10g Release 2 (10.2.0.x) and Oracle Database 11g Release 1 (11.1.0.x). You can apply these patches using a patch plan. However, Patch Sets for Oracle Database 11g Release 2 (11.2.0.x) are complete installs (full releases), and you must use the Database Upgrade feature to apply them. The Database Upgrade feature follows the out-of-place patching approach.
- Using Oracle Enterprise Manager Cloud Control to apply Patch Sets to any Oracle Fusion Middleware software (including Oracle WebLogic Server, Oracle SOA Infrastructure, and Oracle Service Bus) and Oracle Fusion Application software is not supported. Similarly, using Oracle Enterprise Manager Cloud Control to apply Bundle Patches to Oracle Identity Management software is not supported.
- Patch Sets are supported for only one specific version of Oracle Database.
- You cannot add both patch sets and patches to a patch plan. Instead, you can have one patch plan for patch sets, and another patch plan for patches.
- You cannot add patches that require different database startup modes (upgrade or normal) in a single patch plan.

A patch can be added to a target in a plan only if the patch has the same release and platform as the target to which it is being added. You will receive a warning if the product for the patch being added is different from the product associated with the target to which the patch is being added. The warning does not prevent you from adding the patch to the plan.

You can patch any target by using a patch plan. The plan also validates Oracle Database, Fusion Middleware, and Cloud Control patches against your environment to check for conflicts with installed patches.

The patch plan, depending on the patches you added to it, automatically selects an appropriate deployment procedure to be used for applying the patches. For information on the patching deployment procedures used for various database target types, see [Supported Targets, Releases, and Deployment Procedures for Patching](#).

 Note:

- Patch plans are currently not available for hardware or operating system patching.
- Any administrator or role that has view privileges can access a patch plan. For information on roles and privileges required for patch plans and patch templates, see [Creating Administrators with the Required Roles for Patching](#).

About Types of Patch Plans

This section describes the types of patch plans.

A patch plan can be either deployable or non-deployable. A patch plan is deployable when:

- It contains only patches of the same product family and platform.

Note:

- You cannot add both patch sets and patches to a patch plan. Instead, you can have one patch plan for patch sets, and another patch plan for patches.
- You cannot add patches that require different database startup modes (upgrade or normal) in a single patch plan.

- It contains targets that are supported for patching, similarly configured, and are of the same product type, platform, and version (homogeneous targets).

Error Plans

If your patch plan consists of Oracle Management Agent (Management Agent) targets and patches, and the analysis fails on some Management Agents, then the patch plan is split into two plans. The Management Agents and their associated patches on which the analysis was successful are retained in the original plan. The failed targets and their associated patches are moved to a new error plan. The deployment options are also copied into the error plan. The error plan is accessible from the Patches & Updates page.

About the Create Plan Wizard

Figure 44-2 shows the Create Plan Wizard that enables you to create, view, and modify patch plans.

Figure 44-2 Create Plan Wizard

Step 1: Plan Information

Plans can be saved and referenced while you work out any details (especially if you have to wait on patches to be created). Use the Plan to coordinate the patch needs of team members. Then analyze and review to determine what you want to keep.

Overview

Name:

Description:

Created By:

Make Favorite:

Allow access for

Name	Role	Permission
		Add...

* Required Field

The wizard has the following screens:

Screen 1: Plan Information

Enables you to provide basic information about the plan, such as a unique name for the plan and a brief description. Also enables you to add an administrator or a role that can access the patch plan.

Screen 2: Patches

Enables you to view the patches already part of the patch plan, and manually add additional patches to the plan and corresponding targets that need to be patched.

Screen 3: Deployment Options

Enables you to configure the patch plan with deployment options that suit your needs. Although this step is common for all target types, the deployment options offered by this step depend on the target types selected in the patch plan.

For all target types, you can select a customized deployment procedure for deploying the patches, and specify the credentials to be used. For database targets, you can specify a nondefault staging location for storing patches and skip the staging process if the patches are already staged.

For dynamic deployment procedure, you can specify the custom steps that can be added in the generated deployment procedure. However, you cannot choose an existing procedure from the Procedure Library and customize and save it.

For standalone (single-instance) database, Oracle Restart, Oracle RAC database, Oracle Grid Infrastructure (that may or may not be a part of Oracle Exadata), and targets in Oracle Data Guard configuration, and you can choose between out-of-place patching and in-place patching. For Oracle RAC database targets and Oracle Grid Infrastructure targets, you can choose to apply the patches in rolling or parallel mode in order to control the downtime of the system.

Screen 4: Validation

Enables you to validate the patch plan and determine whether the patches can be rolled out without any problems. Essentially, it enables you to perform the following checks using the patch information from Oracle, the inventory of patches on your system (gathered by the configuration manager), and the information from candidate patches.

- Patch Conflict Checks
 - Conflict between the patches added to the patch plan and the patches already present in the Oracle home
 - Conflict among patches within the patch plan
- Target Sanity Checks
 - Target status and configuration checks
 - OPatch and OUI checks
 - SmartUpdate checks
 - Inventory sanity checks, such as locks, access, and so on
 - Hard disk space checks
 - Cluster verification checks (cluvfy, srvctl config)

- SQLPlus checks (with sample SQL)

In addition to checking for conflicts, it enables you to check for patch conflicts between the patches listed in the plan.

Screen 5: Review & Deploy

Enables you to review the details you have provided for the patch plan, and then deploy the plan. After you click the **Deploy** button, a dialog box appears which allows you to confirm the deployment of patches at that time or in future. The page also enables you to review all the impacted targets so that you understand what all targets are affected by the action you are taking.

Overview of Patch Templates

This section describes the following:

- [About Patch Templates](#)
- [About the Edit Template Wizard](#)

About Patch Templates

Patch templates are another important aspect of the patch management solution. Patch templates help you create predesigned plans based on an existing successfully analyzed or deployable patch plan, however without any targets selected. A patch template contains a predetermined set of patches and deployment options saved from the source patch plan, and enables you to select a completely new set of targets.

This way, as a *Patch Designer*, you can create a patch plan with a set of patches, test them in your environment, save the successfully analyzed patch plan as a patch template, and publish them to *Patch Operators*. As a *Patch Operator*, who can create patch plans out of the templates, add another set of targets, and roll out the patches to the production environment in a recursive manner.

This way, you reduce the time and effort required to create new patch plans, and as a *Patch Designer*, you expose only the successfully analyzed and approved plans to *Patch Operators*.



Note:

An administrator or role that has the privileges to create a patch template and view a patch plan, which is being used to create a template, can create a patch template.

About the Edit Template Wizard

[Figure 44-3](#) shows the Edit Template Wizard that enables you to view the contents of a patch plan template and modify its description.

Figure 44-3 Edit Template Wizard

Step 1: Plan Information

Plans can be saved and referenced while you work out any details (especially if you have to wait on patches to be created). Use the Plan to coordinate the patch needs of team members. Then analyze and review to determine what you want to keep.

Overview

* Name:

Description:

Created By: SYSMAN

Make Favorite

Allow access for

Name	Role	Permission
DESIGNER	Admin	View
EMUSER_ADMIN	Admin	Full
EM_ALL_ADMINISTRATOR	Role	Full
EM_ALL_DESIGNER	Role	View
FU_ALL_OPERATOR	Role	View

Exit Wizard Export Plan Back Next Review

When you view or modify a patch plan template, the Edit Template Wizard opens. This wizard has the following screens:

Screen 1: Plan Information

Enables you to view general information about the template, and modify the description and the deployment date.

Screen 2: Patches

Enables you to view a list of patches part of the patch plan template. The patches listed here are the patches copied from the source patch plan that you selected for creating the template.

Screen 3: Deployment Options

Enables you to view the deployment options configured in the patch plan template.

Supported Targets, Releases, and Deployment Procedures for Patching

Table 44-2 and Table 44-3 list the targets and their releases you can patch on different platforms using the new patch management solution, and the default deployment procedures that the patch plans automatically select depending on the target type. The deployment procedures are supported only through patch plans. Although they are exposed in the Deployment Procedure Manager page, you cannot select and run them independently; you must always create a patch plan to run them.

Table 44-2 Supported Targets and Releases for Patching Oracle Database

Supported Target Type	Supported Targets and Releases	Patching Mode	Deployment Mode	Supported Platforms	Deployment Procedure
Oracle Database	Oracle Database (Stand-alone) 10g to 12c ⁴ Multitenant Database (Container and Pluggable Databases)	In-Place	N/A	All Platforms	Patch Oracle Database
	Oracle Database (Stand-alone) 10g to 12c ⁴ Multitenant Database (Container and Pluggable Databases)	Out-of-Place	N/A	All Platforms	Clone and Patch Oracle Database
	Oracle RAC One Node Database 10g to 12c ⁴	In-Place	Rolling	All Platforms except for Microsoft Windows	Patch Oracle RAC Database - Rolling
	Oracle RAC One Node Database 10g to 12c ⁴	In-Place	Parallel	All Platforms except for Microsoft Windows	Patch Oracle RAC Database - Parallel
	Oracle RAC Database 10g to 11.2.0.1	In-Place	Rolling	All Platforms except for Microsoft Windows	Patch Oracle RAC Database - Rolling
	Oracle RAC Database 10g to 11.2.0.1	In-Place	Parallel	All Platforms except for Microsoft Windows	Patch Oracle RAC Database - Parallel
Oracle Database	Oracle RAC Database 11.2.0.2 to 12.1.0.2 ⁴	In-Place	Rolling / Parallel	All Platforms except for Microsoft Windows	Dynamic Deployment Procedures ¹
	Oracle RAC Database 11.2.0.2 to 12.1.0.2 ⁴	Out-of-Place	Rolling / Parallel	All Platforms except for Microsoft Windows	Dynamic Deployment Procedures
	Oracle Grid Infrastructure* 10g to 11.2.0.1	In-Place	Rolling	All Platforms except for Microsoft Windows	Patch Oracle Clusterware - Rolling
	Oracle Grid Infrastructure* 10g to 11.2.0.1	In-Place	Parallel	All Platforms except for Microsoft Windows	Patch Oracle Clusterware - Parallel
Oracle Database	Oracle Grid Infrastructure* 11.2.0.2 to 11.2.0.4	In-Place	Rolling / Parallel	All Platforms except for Microsoft Windows	Dynamic Deployment Procedures ²

Table 44-2 (Cont.) Supported Targets and Releases for Patching Oracle Database

Supported Target Type	Supported Targets and Releases	Patching Mode	Deployment Mode	Supported Platforms	Deployment Procedure
	Oracle Grid Infrastructure* 11.2.0.2 to 11.2.0.4	Out-of-Place	Rolling / Parallel	All Platforms	Dynamic Deployment Procedures
	Oracle Grid Infrastructure* 12.1.0.1 to 12.1.0.2 ⁴	In-Place	Rolling / Parallel	All Platforms	Dynamic Deployment Procedures ³
	Oracle Grid Infrastructure* 12.1.0.1 to 12.1.0.2 ⁴	Out-of-Place	Rolling / Parallel	All Platforms	Dynamic Deployment Procedures
	Oracle Restart 10g to 11.2.0.4	In-Place	N/A	All Platforms	Patch Oracle Restart
	Oracle Restart - Release 12c ⁴ onwards	In-Place	N/A	All Platforms	Patch Oracle Restart(12.1.0.1.0 onwards)
	Oracle Restart - Release 11.2.0.4 and 12.1.0.2 ⁴	Out-of-Place	N/A	Only Linux.X64 and AIX	Dynamic Deployment Procedures

 Note:

Symbol	Description
¹	To override and use the following static deployment procedures for In-Place patching mode. Set (update/insert) the <i>use_static_dp</i> to true in <i>mgmt_patching_properties</i> table. • Patch Oracle RAC Database - Parallel • Patch Oracle RAC Database - Rolling
²	To override and use the following static deployment procedures for In-Place patching mode. Set (update/insert) the <i>use_static_dp</i> to true in <i>mgmt_patching_properties</i> table. • Patch Oracle Clusterware - Parallel • Patch Oracle Clusterware - Rolling
³	To override and use the following static deployment procedures for In-Place patching mode, set (update/insert) the <i>use_static_dp</i> to true in <i>mgmt_patching_properties</i> table. • Patch Oracle Clusterware(12.1.0.1.0 onwards) - Rolling • Patch Oracle Clusterware(12.1.0.1.0 onwards) - Parallel
⁴	Multitenant Database (Container and Pluggable Databases) is supported.
Clusters*	Grid Infrastructure*
N/A	Not applicable.

Table 44-3 Supported Targets and Releases for Patching other target types

Supported Target Type	Supported Targets and Releases	Supported Platform	Supported Default Deployment Procedure
Oracle Service Bus	Oracle Service Bus Infrastructure 13c Release 2	All Platforms	• Patch Oracle Service Bus Infrastructure In Parallel Mode • Patch Oracle Service Bus Infrastructure In Rolling Mode
Oracle WebLogic Server	Oracle WebLogic Server 10g (10.3.1), (10.3.2), (10.3.3), (10.3.4), (10.3.5), (10.3.6), and 12c Release 1 (12.1.1), (12.1.2), (12.1.3), and 12c Release 2 (12.2.1.x)	All Platforms	• Patch Oracle WebLogic Server In Parallel Mode • Patch Oracle WebLogic Server In Rolling Mode
Oracle Service Oriented Architecture	Oracle SOA Infrastructure 11g Release 1 (11.1.1.1.0 - 11.1.1.7.0), 12c Release 1 (12.1.3.0.0), and 12c Release 2 (12.2.1.0.0)	All Platforms	• Patch Oracle SOA Infrastructure In Parallel Mode • Patch Oracle SOA Infrastructure In Rolling Mode
Oracle Identity Management (only Oracle Access Management Server and Oracle Identity Management Server targets are supported)	Oracle Identity Management 11g Release 2 (11.1.2.2.0)	All Platforms	Patching IDM
Oracle Siebel (only Siebel Server and Siebel Gateway Server targets are supported)	Oracle Siebel 8.1.1.9, 8.1.1.10, 8.1.1.11, and 8.2.x	All Platforms	Patching Siebel Targets
Oracle Fusion Applications	Oracle Fusion Applications 11g Release 1 (11.1.1.5.1) and (11.1.2.0.0) (RUP1)	All Platforms	Patch Oracle Fusion Applications
Oracle Database	Oracle Automated Storage Management (Oracle ASM) 10g Release 1 to 11g Release 2	All Platforms	Patch Standalone Oracle ASM
Oracle Database	Oracle Software Only Patching (Grid Home / DB Home)	All Platforms	Patch Oracle Home

 Note:

- You can also patch primary and standby databases configured with Oracle Data Guard. For information on how to patch these targets, see [Patching Oracle Data Guard Targets](#).
- Patching an Oracle database whose listener was started by a user different from the Oracle database home owner is not supported. To patch such a database, stop the database listener, restart it as the database home owner, then apply the required patch using a patch plan.
- From the 12.2 release, classic patching using Lifecycle Management Graphical User Interface is deprecated. You must use the Database Fleet Maintenance feature to patch 12.2 databases and GI environments.
- Shared home patching is not feasible from EM console, because patch plan uses the OPlan.

Overview of Supported Patching Modes

This section describes the following patching modes:

- [Overview of Patching in Online and Offline Mode](#)
- [Overview of Patching in In-Place and Out-of-Place Mode](#)
- [Overview of Patching in Rolling and Parallel Mode](#)

Overview of Patching in Online and Offline Mode

You have the flexibility to choose between Online and Offline modes of patching.

Online Mode

Online mode is useful when Cloud Control can connect to My Oracle Support (MOS) using an Internet connection. Using this mode, you can see recommendations from Oracle for the patches to be applied, and manually search patches directly on MOS and add them to your patch plan. In addition, you can access community information, knowledge articles, service requests, and so on, and also automatically resolve patch conflicts with a merge patch directly from MOS.

Note that Cloud Control does not upload any data to MOS. It only uses MOS to download the latest updates.

Offline Mode

Offline mode is useful when Cloud Control cannot connect to My Oracle Support. Using this mode, you can search patches that were manually uploaded to the Software Library, and add them to your patch plan. In offline mode, you cannot do the following:

- Search and download patches from My Oracle Support
- View additional information about the patch
- Access community information, knowledge articles, service requests

- View the Related Activity region



Note:

- By default, the patching mode is set to online. If you want to switch the mode to offline, then from the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**. For connection, select **Offline**.
- Ensure to run the necessary OPatch in online mode prior to applying the patches in offline mode. This can be accomplished either manually uploading the relevant version of OPatch to the software library or execute the out of box OPatch update job.

Overview of Patching in In-Place and Out-of-Place Mode


You have the flexibility to choose between in-place and out-of-place patching modes.

In-Place Mode

In-place mode of patching is a patching mechanism where you directly patch the database home. In this mode, before applying the patch, you bring down all the database instances running out of that database home. Thus, there is downtime. After applying the patch, you restart all the database instances. All the database instances are patched in the same manner. Note that in this mode, recovery takes longer if there is a problem, because the Oracle home that you patched is the original database home, and not a copy.

Figure 44-4 illustrates how multiple database instances running from an Oracle home get patched in in-place patching mode.

Figure 44-4 In-Place Mode of Patching



Out-of-Place Mode

Out-of-place mode of patching is a patching mechanism that clones the existing database home, and patches the cloned home instead of the original home. Once the cloned home is patched, you can migrate the database instances to run from the cloned home, which ensures minimal downtime.

 Note:

- Out-of-place patching is supported for standalone (single-instance) database targets, Oracle Restart targets, Oracle RAC Database targets, Oracle Data Guard targets, and Oracle Grid Infrastructure targets (that may or may not be a part of Oracle Exadata).

The 13.1.0.1 Enterprise Manager for Oracle Database plug-in is required for patching Oracle Restart targets in out-of-place mode.

The 12.1.0.5 Enterprise Manager for Oracle Database plug-in is required for patching Oracle RAC Database targets and Oracle Grid Infrastructure targets that are not a part of Oracle Exadata, in out-of-place mode. The 12.1.0.6 Enterprise Manager for Oracle Database plug-in is required for patching Oracle Data Guard targets in out-of-place mode.

- If the cloned home contains certain additional patches (that are not added to the patch plan) that include SQL statements, the SQL statements for these patches are not executed automatically when the database instance is migrated to the cloned home. For these patches, you must execute the SQL statements manually.

While migrating the database instances, you can choose to migrate all the instances, or only some of them, depending on the downtime you can afford to have in your data center. If you choose to migrate only a few instances in one session, then ensure that you migrate the rest in the next session. This way, you can control the downtime in your data center as you divide the migration activity. This is particularly useful when you have multiple database instances running out of an Oracle home.

 Note:

You select the database instances that you want to migrate, while creating the patch plan using the Create Plan Wizard. The selected database instances are migrated when the patch plan is in the *Deploy* state. If you selected only a few database instances for migration, create another patch plan to migrate the remaining instances. On the Deployment Options page, select the existing home, then select the remaining database instances that need to be migrated.

In case of a system patch such as a Grid Infrastructure patchset update, edit and provide the candidate Oracle Home location. This location is validated during analysis.


Oracle always recommends out-of-place patching because downtime is minimal and recovery in case of a problem is easy; you can always switch back to the original database home in case of a problem with the clone home.

 Note:

If you patched an Oracle RAC target, an Oracle single-instance database target, Oracle Restart target, or an Oracle Grid Infrastructure target (that may or may not be a part of Oracle Exadata), in out-of-place patching mode, then you can switch back to the original home directly from the Create Plan Wizard—the wizard provides a **Switchback** button that enables you to perform the operation. However, you cannot perform this operation for any other target type.

Figure 44-5 illustrates how multiple database instances running from an Oracle home get patched in out-of-place patching mode.

Figure 44-5 Out-of-Place Mode of Patching



 Note:

Alternatively, to obtain a new Oracle home that has the required patches, you can provision it directly from a software image (that has the required patches) using provisioning deployment procedures, and then use a patch plan for the analysis and post patching steps. For information on how to do this, see [Analyzing, Preparing, and Deploying Patch Plans](#).

Overview of Patching in Rolling and Parallel Mode

While patching Oracle Real Application Cluster (Oracle RAC) targets, Oracle Grid Infrastructure targets (whether or not they are part of Oracle Exadata), Oracle Data

Guard targets, Oracle WebLogic Server targets, Oracle Fusion Application targets, or Oracle SOA Infrastructure targets you can choose to patch the instances of the cluster either in rolling or parallel mode.

Rolling Mode

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 target 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.




Note:

The ReadMe of the patch clearly states whether or not you can use the *Rolling Mode* to apply your patches. Therefore, use this mode only if it is stated in the ReadMe. This is validated during analysis.

Figure 44-6 illustrates how a two-node Oracle RAC target gets patched when rolling mode is used.

Figure 44-6 Rolling Mode of Patching




Parallel Mode

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.


Figure 44-7 illustrates how a two-node Oracle RAC target gets patched when parallel mode is used.

Figure 44-7 Parallel Mode of Patching



Understanding the Patching Workflow

The following illustration describes the overall workflow of the patch management solution offered through the integrated functionality within the Cloud Control console.



Step	Step Name	Description	Reference Links
Step 1	Set Up Infrastructure	Meet the prerequisites and set up the infrastructure for rolling out patches. Essentially, create admin roles for creating <i>Patch Plans</i> and <i>Patch Templates</i> , meet other mandatory and optional prerequisites, make online or off-line patching settings.	Setting Up the Infrastructure for Patching
Step 2	Identify the Patches	View the recommendations made by Oracle on the patches to be applied, and identify the ones you want to apply. Access community information (from innumerable customers).	Identifying the Patches to Be Applied
Step 3	Apply Patches	Create patch plans with patches and associated targets, perform prerequisite checks, analyze the patches for conflicts and resolve the issues, and then save the successfully analyzed plan as a patch template. Then, create a new patch plan out of the patch template and use that to deploy the patches in your environment.	Applying Patches

Setting Up the Infrastructure for Patching

This section describes how you can set up the infrastructure for patching. Meet these prerequisites before you start rolling out the patches.

NOT_SUPPORTED:



This section is mainly for Patch Administrators or Patch Designers who want to keep the infrastructure ready for rolling out patches. This is mostly a one-time task if you have decided on the way (online or off-line) you want to patch.

This section covers the following:

- [Meeting Basic Infrastructure Requirements for Patching](#)
- [Creating Administrators with the Required Roles for Patching](#)
- [Setting Up the Infrastructure for Patching in Online Mode \(Connected to MOS\)](#)
- [Setting Up the Infrastructure for Patching in Offline Mode \(Not Connected to MOS\)](#)
- [Analyzing the Environment and Identifying Whether Your Targets Can Be Patched](#)

Meeting Basic Infrastructure Requirements for Patching

Meet the basic infrastructure requirements as described in [Setting Up Your Infrastructure](#). The chapter describes both mandatory and optional requirements.

Also, when patching a domain with the admin port and/or the SSL port enabled, the patching deployment procedure assumes that the SSL certificates are consistent across all nodes in the domain. Therefore, ensure that the SSL certificates are in the same location on all hosts. If the SSL certificates are in different locations, then the deployment procedure will fail as it will not be able to find the certificates in some hosts.

Note:

If the targets that you want to patch are deployed on Microsoft Windows hosts, you must ensure that Cygwin is installed on the target hosts, before patching the targets.

For information on how to install Cygwin on a host, see *Enterprise Manager Cloud Control Basic Installation Guide*.

Creating Administrators with the Required Roles for Patching

[Table 44-4](#) describes the roles and the minimum privileges required for using patch plans and patch templates. These roles are default roles available in Cloud Control.

You need not create them, but you must explicitly create administrators based on these roles. For instructions, see [Creating Enterprise Manager User Accounts](#)

Table 44-4 Roles and Privileges for Using Patch Plans and Patch Templates

Role	Patch Plan Scope	Patch Template Scope	Patch Plan and Patch Templates Privileges	Target Privileges	Resource Privileges	Implementation Recommendation
Patch Administrator	Create, View, Modify, Delete	Create, View, Modify, Delete	FULL_ANY_PATCH_PLAN FULL_ANY_PATCH_TEMPLATE GRANT_PRIV_PATCH_PLAN	Operator on selected Target (You can select a particular target you want to patch and grant operator privilege to it)	Resource Type: Deployment Procedure Privilege: Create Resource Type: Job System Privilege: Create Resource Type: Named Credential Privilege: Create Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity	Required if you want to create and manage <i>Patch Designer</i> and <i>Patch Operator</i> roles. You can also create these roles directly as a SUPER ADMIN or SYSMAN.
Patch Designer	Create, View	Create, View	FULL_PATCH_PLAN FULL_PATCH_TEMPLATE	Operator on selected Target (You can select a particular target you want to patch and grant operator privilege to it)	Resource Type: Deployment Procedure Privilege: Create Resource Type: Job System Privilege: Create Resource Type: Named Credential Privilege: Create Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity	Required when you want to grant the access and also have restrictions. Alternatively, you can create an EM user with <i>Patch Designer</i> role.
Patch Operator	Create	View	CREATE_PATCH_PLAN VIEW_ANY_PATCH_TEMPLATE	Operator on selected Target (You can select a particular target you want to patch and grant operator privilege to it)	Resource Type: Deployment Procedure Privilege: Create Resource Type: Job System Privilege: Create Resource Type: Named Credential Privilege: Create Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity	Required when you want to grant the access and also have restrictions. Alternatively, you can create an EM user with <i>Patch Operator</i> role.

Setting Up the Infrastructure for Patching in Online Mode (Connected to MOS)

If you choose to patch your targets when Cloud Control is online, that is, when it is connected to MOS, then meet the following setup requirements:

- [Enabling Online Mode for Patching](#)
- [Registering the Proxy Details for My Oracle Support](#)

Enabling Online Mode for Patching

Note:

- This is the default mode for patching in Cloud Control. Therefore, you do not have to manually set this up the first time. However, if you have set it to Offline mode for a particular reason, and if you want to reset it to Online mode, then follow the steps outlined in this section.
- Cloud Control does not upload any data to MOS. It only uses MOS to download the latest updates.

To patch the targets in Online mode, you must set the connection setting in Cloud Control to Online mode. To do so, log in as a user that has the Patch Administrator role, then follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. For **Connection**, select **Online**.

Registering the Proxy Details for My Oracle Support

In Online mode, Cloud Control connects to MOS to download patches, patch sets, ARU seed data such as products, platforms, releases, components, certification details, and patch recommendations. For this purpose, Cloud Control uses the Internet connectivity you have on the OMS host to connect to MOS. However, if you have a proxy server set up in your environment, then you must register the proxy details. You can register the proxy details for MOS using the My Oracle Support Proxy Settings page.

Note:

Beginning with Enterprise Manager Cloud Control 12c Release 3 (12.1.0.3), My Oracle Support accesses support.oracle.com directly. This means that you must provide network access to this URL, or grant proxy access to it from any client that will access My Oracle Support.

To register the proxy details for My Oracle Support (MOS), follow these steps:

1. From the **Setup** menu, select **Proxy Settings**, then select **My Oracle Support**.
2. If you want the OMS to connect to MOS directly, without using a proxy server, follow these steps:
 - a. Select **No Proxy**.
 - b. Click **Test** to test if the OMS can connect to MOS directly.
 - c. If the connection is successful, click **Apply** to save the proxy settings to the repository.
3. If you want the OMS to connect to MOS using a proxy server, follow these steps:
 - a. Select **Manual proxy configuration**.
 - b. Specify the proxy server host name for **HTTPS** and an appropriate port value for **Port**.
 - c. If the specified proxy server has been configured using a security realm, login credentials, or both, select **Password/Advanced Setup**, then provide values for **Realm**, **User Name**, and **Password**.
 - d. Click **Test** to test if the OMS can connect to MOS using the specified proxy server.
 - e. If the connection is successful, click **Apply** to save the proxy settings to the repository.

 Note:

- If you are using a proxy server in your setup, ensure that it allows connectivity to `aru-akam.oracle.com`, `ccr.oracle.com`, `login.oracle.com`, `support.oracle.com`, and `updates.oracle.com`.

NTLM (NT LAN Manager) based Microsoft proxy servers are not supported. If you are using an NTLM based Microsoft proxy server, to enable access to the above sites, add the above URLs to the Unauthenticated Sites Properties of the proxy server.
- The MOS proxy server details specified on the MOS Proxy Settings page apply to all OMSes in a multi-OMS environment.

Setting Up the Infrastructure for Patching in Offline Mode (Not Connected to MOS)

If you choose to patch your targets when Cloud Control is offline, that is, when it is not connected to My Oracle Support, then meet the following setup requirements:

- [Enabling Offline Mode for Patching](#)
- [Downloading Enterprise Manager Catalog Zip File From Another Host With Internet Connectivity](#)
- [Uploading Enterprise Manager Catalog Zip File from your Host With No Internet Connectivity](#)

- [Uploading Patches to Oracle Software Library](#)

Enabling Offline Mode for Patching

To patch the targets in offline mode, you must set the connection setting in Cloud Control to offline mode. To do so, follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.

Note:

Ensure to run the necessary OPatch in online mode prior to applying the patches in offline mode. This can be accomplished either manually uploading the relevant version of OPatch to the software library or execute the out of box OPatch update job.

2. For **Connection**, select **Offline**.

Note:

Once Cloud Control is running in offline mode, you must download the latest Enterprise Manager catalog file from a host that has internet connectivity, transfer the catalog file to your local host, then upload the catalog file to the Management Repository. For information on how to do this, see [Downloading Enterprise Manager Catalog Zip File From Another Host With Internet Connectivity](#) and [Uploading Enterprise Manager Catalog Zip File from your Host With No Internet Connectivity](#).

Downloading Enterprise Manager Catalog Zip File From Another Host With Internet Connectivity

In Offline mode, you must use another host that has an Internet connection, and manually download the `em_catalog.zip` file from My Oracle Support. Use the following URL to download the latest catalog file:

`https://updates.oracle.com/download/em_catalog.zip`

Information about the targets affected by the latest patches, the patches that you have to download manually, and so on is available in the catalog zip file.

Uploading Enterprise Manager Catalog Zip File from your Host With No Internet Connectivity

After downloading the catalog zip file as described in the preceding section, ensure that you transfer the latest downloaded `em_catalog.zip` file back to your local host using FTP or any other file transfer methodology. Then, from your local host, you can log in to Cloud Control to upload the zip file. To do so, follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. Click **Browse** to specify the location of the latest `em_catalog.zip` file.
3. Click **Upload** to upload the file to the Management Repository.

On clicking **Upload**, a Refresh From My Oracle Support job is created and submitted to the Enterprise Manager job system.

Uploading Patches to Oracle Software Library

For patching targets in offline mode, you must have already stored the patches and their metadata files in Software Library so that they can be searched, selected, and added to the patch plans from Software Library. For this purpose, you must first manually download the patches and their metadata files from My Oracle Support, and then upload them to the Software Library.

This section describes the following:

Downloading a Patch from My Oracle Support

To download a patch and its metadata file from My Oracle Support, follow these steps:

1. Log in to My Oracle Support (<https://support.oracle.com/>), then click the **Patches & Updates** tab.
2. On the Patches & Updates page, in the Patch Search section, enter the patch number you want to search for as shown in [Figure 44-8](#), then click **Search**.

Figure 44-8 Searching for Patches


3. On the Patch Simple Search Results page, select the row that has the patch that you want to download. Click **Download**. In the File Download dialog, click the name of the patch zip file to download it to your local host. Click **Download Patch Metadata**, and then in the Download Patch Metadata dialog, click **Download** to download the patch metadata file. This step is described in [Figure 44-9](#).



Note:

Oracle recommends that you transfer the patch ZIP file and the metadata XML file to the Management Agent host, where the Management Agent could be an agent on an OMS machine, or on the target host. Upload these files from the Management Agent host to Software Library.

Figure 44-9 Downloading Patches from My Oracle Support



Uploading Patches to Software Library Using the Cloud Control Console

Using this method, you can upload only a single patch at a time. Therefore, use this method only when you want to upload a few patches. Also, use this method when the sizes of the patches that you want to upload are small.

To upload a patch to Software Library using the Cloud Control console, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Saved Patches**.
2. Click **Upload**.
3. For **Patch Zip File**, specify the location of the patch zip file you downloaded onto your local host. If the patch zip file you downloaded contains the `PatchSearch.xml` file (a file containing patch metadata information such as patch ID, product, platform, language etc.), you do not need to specify a value for **Patch Metadata**. However, if the patch zip file you downloaded does not contain the `PatchSearch.xml` file, and you downloaded the patch metadata file onto your local host separately, for **Patch Metadata**, specify the location of the patch metadata file.

On a Unix based operating system, run the following command to verify whether the `PatchSearch.xml` file is contained within a patch zip file:

```
unzip -l <patch zip file path> | grep PatchSearch.xml
```

For information on how to download the patch metadata file of a patch, see [Uploading Patches to Oracle Software Library](#).

4. Click **Upload** to upload the patch to Software Library.

 Note:

If you encounter an error mentioning that the patch could not be uploaded as it is too large, either use EM CLI to upload the patch (as described in [Uploading Patches to Oracle Software Library](#)), or run the following command, re-start the OMS, then retry the patch upload:

```
emctl set property -name "oracle.sysman.emSDK.ui.trinidad.uploadedfilemaxdiskspace" -sysman_pwd sysman -value 2589934592
```

Ensure that the value you specify for `-value` is in bytes, and is larger than the size of the patch that you want to upload.

Uploading Patches to Software Library Using EM CLI

Using this method, you can perform a batch upload of multiple patches. Also, this method is faster than using the Cloud Control console to upload patches. Hence, use this method when you want to upload multiple patches at one time, or the sizes of the patches that you want to upload are large.

To upload patches to Software Library using EM CLI, follow these steps:

1. Set up EM CLI on the host on which the downloaded patch files that you want to upload are located.

EM CLI is set up by default on every OMS host. For information on how to set up EM CLI on a host that is not running the OMS, refer the Command Line Interface Concepts and Installation chapter of *Oracle Enterprise Manager Command Line Interface*.

2. From the EM CLI install location, log in to EM CLI:

```
<emcli_install_location>/emcli login -username=<username> -password=<password>
```

For example,

```
<emcli_install_location>/emcli login -username=sysman -password=2be-not2be
```

 Note:

Ensure that the EM CLI log in user has the ADD_TARGET privilege.

3. Synchronize EM CLI:

```
<emcli_install_location>/emcli sync
```

4. Run the `upload_patches` verb to upload the required patches to Software Library:

```
<emcli_install_location>/emcli upload_patches
    -location=<patch_location> | -patch_files=<meta-data_file_path;ZIP_file_path;second_part_of_ZIP_file_path;>
    -from_host=<host_name>
    [-cred_name=<credential_name> [-cred_owner=<credential_owner>]]
```


The parameters mentioned in [] are optional.

Use the `-from_host` option to specify the host on which the downloaded patch zip files and metadata files are present. You can use the `-location` option to specify the location of the downloaded patch files on the host you specified using `-from_host`. When you specify `-location`, all the patch zip files and metadata files present at the specified location are uploaded to Software Library. Hence you can use this option to perform a batch upload of multiple patches to Software Library. For example:

```
./emcli upload_patches -location=/scratch/aime/patches -from_host=h1.example.com
```

This example uploads all the patch zip files and patch metadata files present at `/scratch/aime/patches` on the `h1.example.com` host to Software Library.

Use the `-patch_files` option to provide the absolute path of a patch zip file and its patch metadata file. If you use this option, you can specify only one patch zip file. Hence, you can use this option to upload only a single patch at a time. Also, use the `-cred_name` option to specify the named credentials that must be used to access the specified host, and the `-cred_owner` option to specify the owner of the specified named credential. If you do not specify the `-cred_name` option, the preferred normal credentials are used to access the host. If you do not specify the `-cred_owner` option, the credential owner is assumed to be the current user. For example:

```
./emcli upload_patches -patch_files="/scratch/p13741363_112310_Linux-x86-64_M.xml;/scratch/p13741363_112310_Linux-x86-64.zip" -from_host=h1.example.com -cred_name=AIMECRED -cred_owner=SYSMAN
```

This example uploads the `p13741363_112310_Linux-x86-64.zip` patch zip file and the `p13741363_112310_Linux-x86-64_M.xml` metadata file present on the `h1.example.com` host to Software Library, using the `AIMECRED` named credential which is owned by `SYSMAN`.

 Note:

Ensure that you specify either the `-location` option or the `-patch_files` option with the `upload_patches` verb, but not both. Specifying the `-location` option enables you to perform a batch upload of multiple patches, and is hence the recommended option.

To view more information on the syntax and the usage of the `upload_patches` verb, run the following command:

```
$<OMS_HOME>/bin/emcli help upload_patches
```

Analyzing the Environment and Identifying Whether Your Targets Can Be Patched

Before creating a patch plan to patch your targets, Oracle recommends that you view the patchability reports to analyze the environment and identify whether the targets you want to patch are suitable for a patching operation. These reports provide a summary of your patchable and non patchable targets, and help you create deployable

patch plans. They identify the problems with the targets that cannot be patched in your setup and provide recommendations for them.

Patchability reports are available for Oracle Database, Oracle WebLogic Server, and Oracle SOA Infrastructure targets.

To view the patchability reports, follow these steps:

1. From the **Enterprise** menu, select **Reports**, then select **Information Publisher Reports**.
2. Click **Expand All** to view all the branches under Information Publisher Reports.
3. To view the patchability report for Oracle Database targets, under the Deployment and Configuration branch, and the Patching Automation Reports sub branch, select **EM Target Patchability Report**.
4. To view the patchability report for Oracle Fusion Middleware (Oracle WebLogic Server and Oracle SOA Infrastructure) targets, under the Deployment and Configuration branch, and the Patching Automation Reports sub branch, click **EM FMW Target Patchability Report**.

 Note:

If you see any missing property errors, then resolve the errors using the workarounds described in [Workaround for Missing Property Errors](#). If you see any unsupported configuration errors, then resolve the errors using the workarounds described in [Workaround for Unsupported Configuration Errors](#).

Identifying the Patches to Be Applied

This section describes how you can identify the patches to be applied.

NOT_SUPPORTED:



This section is mainly for Patch Designers who want to keep track of the various patch releases, look for recommendations from Oracle, and identify the ones they want to roll out in the environment.

This section covers the following:

- [About Patch Recommendations](#)
- [About Knowledge Articles for Patching](#)
- [About Service Requests for Patching](#)
- [Searching for Patches on My Oracle Support](#)
- [Searching for Patches in Oracle Software Library](#)

About Patch Recommendations

Patch recommendations are proactive notifications of potential system issues and recommendations that help you improve system performance and avert outages.

The 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).

Enterprise Manager Cloud Control 13c supports patch recommendations for the following target types:

- Oracle Database
- Enterprise Manager Cloud Control (Management Agent/Oracle Management Service)
- Oracle Fusion Applications
- SOA Infrastructure
- Oracle Real Applications Clusters
- Oracle Real Applications Clusters with Clusterware
- Oracle Data Guard
- Oracle Exadata (QFSDP recommendations only)
- Oracle Exalytics (QFSDP recommendations only)
- Oracle Coherence, Oracle WebLogic Server, Oracle JRockit for non-Exalogic systems
- Oracle Siebel
- Oracle Business Process Management
- Application Development Framework Runtime
- Oracle Service Bus
- WebLogic Server

 Note:

The pre-requisite for patching WebLogic domains is that the underlying Oracle Home must be discovered as a target. For SOA 11g, 12c and Service Bus domains this is automatically handled. However, for Service Bus 11g when the domain is discovered the OH is not discovered as a target. If you attempt to patch such a domain then the following warning is displayed. To overcome this, ensure that the Middleware Home and the underlying Oracle Homes are discovered as targets and the configuration collections for Oracle Home targets are the current ones.

For more information on patching Oracle Service Bus and respective recommendations, see [Patching Oracle Service Bus](#).

Patches mentioned in the Patch Recommendation section are a collection of patches offered within MOS which can be applied as a group to one or more targets. To keep the Patch Recommendation section updated with the latest patches for your environment, there is a step called the Config Collection step that runs as a part of the patch plan when a patching job is submitted. Essentially, Config Collection enables to collect the changes that happen due to application of a patch or a rollback. These updates are communicated to the OMS through the Management Agents, which ultimately help in updating the patch recommendation region.

 Note:

Starting with Enterprise Manager 12c, the Config Collection that is triggered at the end of patch application happens asynchronously, which means that collection may not complete when the plan completes execution. In such cases, you might need to recalculate the patch recommendations for your enterprise. Also, if the target collection has not happened properly, then too you might have to recalculate the patch recommendations.

To do so, follow these steps:

1. Run the following EM CLI commands to get the target information of a patch plan:

`emcli get_patch_plan_data -name=<name of the plan>`
2. Perform the following steps to determine the time when the plan was deployed on the targets:
 - a. From **Enterprise** menu, select **Provisioning and Patching**, and then click **Patches and Updates**.
 - b. On the Patches and Updates page, from the Plans section, select the patch plan.
 - c. From the Create Plan Wizard, select **Review and Deploy**.
 - d. Click the link to track the details of the deployment.
 - e. Note down the start time of the job from the Job UI page.
3. On the Config Management side, use the following attributes to calculate the collection time stamp using MGMT\$ECM_CURRENT_SNAPSHOTS view as follows:

```
start_timestamp - timestamp of the collection (in target timezone)
last_upload_timestamp (in DB timezone) when a collection was processed
for this snapshot type.
```

A combination of the start timestamp and last uploaded timestamp attributes help you determine when the collection has happened, and thereby lets you determine if the information available in the patch recommendation region is up-to-date.

The Recommended Patches region appears by default on the Patch & Updates page. You can edit this region to filter its contents.

To view details of a recommended patch, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Patch Recommendations region, click on the bar graph pertaining to the desired patches.

**Note:**

If you do not see the Patch Recommendations region, click **Customize Page** from the top-right corner, and then drag the Patch Recommendations region to the page.

Alternatively, click **All Recommendations** to display all recommendations in the Patch Recommendation page. The Patch Recommendations page also displays the recommended Patch Set Updates and Composite Patch Updates currently available for the targets.

3. On the Patch Recommendations page, select a recommended patch to view the context menu appears. From the context menu, click **Full Screen**.

About Knowledge Articles for Patching

Knowledge articles are documents published on My Oracle Support. These articles can either be announcements or workarounds to known issues.

Some of the knowledge articles that describe workarounds to known issues have patch numbers mentioned. You can choose to make a note of this patch number and search it on My Oracle Support or Oracle Software Library as described in [Searching for Patches on My Oracle Support](#) or [Searching for Patches in Oracle Software Library](#), respectively.

Alternatively, you can click the patch number in the knowledge article. This takes you to the Patch Search page. On the Patch Search page, from the context menu of the patch, click **Add to Plan**, and select either **Add to New** if you want to add the patch to a new patch plan, or **Add to Existing** if you want to add the patch to an existing patch plan.

About Service Requests for Patching

Service requests are support requests raised on My Oracle Support to report and track issues, and receive online assistance from Oracle in resolving those issues.

Some of the service requests that describe workarounds to known issues have patch numbers mentioned. You can choose to make a note of this patch number and search it on My Oracle Support or Oracle Software Library as described in [Searching for Patches on My Oracle Support](#) or [Searching for Patches in Oracle Software Library](#), respectively.

Alternatively, you can click the patch number in the knowledge article. This takes you to the Patch Search page. On the Patch Search page, from the context menu of the patch, click **Add to Plan**, and select either **Add to New** if you want to add the patch to a new patch plan, or **Add to Existing** if you want to add the patch to an existing patch plan.

Searching for Patches on My Oracle Support

If you already know about the existence of a patch from external sources such as blogs, Oracle technology forums, or from colleagues, then use the search functionality to search for those patches. The search functionality enables you to perform more

flexible and advanced searches, and offers capabilities such as saving a search that is used routinely, and searching based on existing saved searches. All of this lets you perform searches more quickly and efficiently.

To search a patch on My Oracle Support, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Patch Search region, enter the search parameters you want to use and click **Search**.

 Note:

If you do not see the Patch Search region, click **Customize Page** from the top-right corner, and then drag the Patch Search region to the page.

Alternatively, you can use the **Saved** tab to search any previously save searches. You can also use the **Recent** tab to access any recently performed identical (or similar) search.

After the patch search is complete, the results appear in the **Patch Search Results** page. On this page, you can select a patch and download it either to the local host (desktop) or to the Software Library.

Searching for Patches in Oracle Software Library

By default, when you search a patch on the Patches & Updates screen, the Cloud Control connects to My Oracle Support using the Internet connectivity available on that host, and searches the requested patch in My Oracle Support. This is because the search functionality is set to perform in online mode by default.

However, if your host does not have Internet connectivity, then you must switch over to offline mode so that the search can be performed in Oracle Software Library (Software Library).

 Note:

Ensure to run the necessary OPatch in online mode prior to applying the patches in offline mode. This can be accomplished either manually uploading the relevant version of OPatch to the software library or execute the out of box OPatch update job.

To switch over to offline mode, follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. For **Connection**, select **Offline**.

 Note:

In offline mode, you cannot:

- Search and download patches from My Oracle Support
- Resolve patch conflicts with merge patches
- View the Related Activity region
- Access Quicklinks
- View or create upgrade plans

To search a patch in the Software Library, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Software Library Patch Search region, enter the search parameters you want to use and click **Search**.

 Note:

If you do not see the Patch Search region, click **Customize Page** from the top-right corner, and then drag the Patch Search region to the page.

Once the patch search is complete, the results appear in the **Patch Search Results** page.

Applying Patches

This section describes how you can create a patch plan with patches, save the patch plan as a patch template, create a new patch plan out of that template, and then apply the patches.

 Note:

All the targets running on a host must be registered as Enterprise Manager targets for patching to be successful. Otherwise, it could result in partial patching wherein the binary bits are patched but the SQL is not applied and so on.

For example, if you have a RAC system with four Oracle homes and five instances of the Database running, then all four Oracle homes and five databases must be registered Enterprise Manager targets for the patching process to be successful.

This section covers the following:

- [Creating a Patch Plan](#)

- [Accessing the Patch Plan](#)
- [Analyzing, Preparing, and Deploying Patch Plans](#)
- [Switching Back to the Original Oracle Home After Deploying a Patch Plan](#)
- [Saving Successfully Analyzed or Deployed Patch Plan As a Patch Template](#)
- [Creating a Patch Plan from a Patch Template and Applying Patches](#)
- [Patching Oracle Grid Infrastructure Targets](#)
- [Patching Oracle Exadata](#)
- [Patching Oracle Data Guard Targets](#)
- [Patching Oracle Identity Management Targets](#)
- [Patching Oracle Siebel Targets](#)
- [Patching Oracle Service Bus](#)
- [Rollback of Oracle Service Bus Patch](#)
- [Deploying WebLogic Patches Along with SOA or Oracle Service Bus Patches In A Single Patch Plan](#)

Creating a Patch Plan

NOT_SUPPORTED:



This section is mainly for Patch Designers who want to create patch plans.

To create a patch plan, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, identify the patches you want to apply as described in [Identifying the Patches to Be Applied](#).
3. On the Patch Recommendations page or on the Patch Search page (depending on how you identified the patch), select a patch you want to add to the Patch Plan.
4. From the context menu, click **Add to Plan**, and select **Add to New**.

**Note:**

If you have already created a patch plan, and if you want to add patches to that patch plan, then you can select **Add to Existing**.

5. In the Create a New Plan window, enter a unique name for your Patch Plan, and click **Create Plan**.

The patch you select and the target it is associated with get added to the plan.

**Note:**

- If the patch you selected impacts other targets, then you are prompted to add the impacted targets as well.
- When you create a Patch Plan, you can add a target that is a member of any system target in Cloud Control. When doing so, the other member targets of that system are automatically added to the Patch Plan. A system is a set of infrastructure components (hosts, databases, application servers, and so on) that work together to host your applications. In Cloud Control, a system and each of the components within it are modeled as individual target types that can be monitored.
- If the WebLogic domain target that you add to a patch plan is a shared domain, then all the Administration Servers and the Managed Servers that are running on the domains that are shared with the domain being patched are automatically added into the same patch plan.
- For Oracle SOA Infrastructure targets, all the SOA WebLogic domains that must be shutdown to patch the SOA targets are added to the patch plan as impacted targets. Therefore, the Administration Server and the Managed Servers running in each of these domains also are affected, and form the *Other impacted targets* when creating a patch plan.

Accessing the Patch Plan

NOT_SUPPORTED:



This section is mainly for Patch Designers who want to access the patch plans they have created.

To access the patch plan you created in [Creating a Patch Plan](#), use one of the following approaches.

Approach 1: Accessing Patch Plan from Plans Region

To access the patch plan from the Plans region, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, click the Patch Plan you want to view. Alternatively, select the Patch Plan, and in the context menu, click **View**. The Create Plan Wizard appears.

To filter the plans table, select **All Plan Types** or **Patch** depending on your preference. To search for a plan, enter a plan name or partial plan name in the search box, then click the search button.

Note:

- If you do not see the Plans region, click **Customize Page** from the top-right corner, and then drag the Plans region to the page.
- To view only the plans that you created, click the icon of a person in the Plans region.

Approach 2: Accessing a Patch Plan from the Patch Recommendations Region

To access the patch plan from the Patch Recommendations region, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Patch Recommendations region, click **All Recommendations**.

 Note:

If you do not see the Patch Recommendations region, click **Customize Page** from the top-right corner, and then drag the Patch Recommendations region to the page.

3. On the Patch Recommendations page, in the table, a patch that is already part of a plan is shown with the plan icon in the **In Plan** column. Click the icon to display all plans with which the patch is associated, and then click a plan to open the Create Plan Wizard.

Analyzing, Preparing, and Deploying Patch Plans

NOT_SUPPORTED:



This section is mainly for Patch Designers who want to analyze the patch plans and deploy them to roll out the patches.

To analyze the patch plan you created in [Creating a Patch Plan](#) and deploy it (or save it as a patch template), follow these steps:

1. Access the patch plan using one of the approaches described in [Accessing the Patch Plan](#).
Cloud Control displays the Create Plan Wizard.
2. On the Plan Information page, do the following:
 - a. In the Overview section, validate the Patch Plan name. You can choose to edit it if you want.
 - b. (Optional) Enter a short description for the patch plan.
 - c. (Optional) In the Allow Access For section, click **Add** to grant patch plan access permissions to administrators or roles, for the current patch plan.

In the Add Privileges to Administrators window, select an administrator or a role, the access permission that you want to grant, then click **Add Privilege**.

 Note:

- From within a patch plan, you can only grant patch plan permissions to administrators or roles that have previously been created in Cloud Control. You cannot create administrators or roles from within a patch plan. For information on the roles and privileges required for patch plans and patch templates, see [Creating Administrators with the Required Roles for Patching](#).
- To remove an administrator or a role, in the Allow Access For section, select the administrator or role that you want to remove, then click **Delete**.

d. Click **Next**.

3. On the Patches page, do the following:

- a. Review the patches added to the patch plan. Any recommended patches for your configuration are automatically added to the patch plan. In addition, any patches that you added manually are listed.

If you are patching an Oracle Grid Infrastructure target that is part of Oracle Exadata, then you can add one Exadata Bundle Patch, and any number of one-off Grid Infrastructure and Oracle Database patches to a single patch plan, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed. In this scenario, ensure that you add the Exadata Bundle Patch while creating the patch plan, and then add the one-off Grid Infrastructure and Oracle Database patches as additional patches.

If you do not have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed, then you cannot add one-off patches to a patch plan that already contains an Exadata Bundle Patch.

If you are patching an Oracle Grid Infrastructure target that is not part of Oracle Exadata, then you can add one Grid Infrastructure Patch Set Update (PSU), and any number of one-off Grid Infrastructure and Oracle Database patches to a single patch plan, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed. In this scenario, ensure that you add the Grid Infrastructure PSU while creating the patch plan, and then add the one-off Grid Infrastructure and Oracle Database patches as additional patches.

If you do not have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed, then you cannot add one-off patches to a patch plan that already contains a Grid Infrastructure PSU.

To associate additional targets to a patch that is already in your patch plan, follow the instructions outlined in [Associating Additional Targets to a Patch in a Patch Plan](#).

To view the details of a patch, select the patch, then from its context menu, click **View**. To temporarily remove the patch from analysis and deployment, click **Suppress**. This leaves the patch in the patch plan, but does not consider it for analysis and deployment.

- b. Click **Next**.

**Note:**

Patching an Oracle database whose listener was started by a user different from the Oracle database home owner is not supported. To patch such a database, stop the database listener, restart it as the database home owner, then apply the required patch using a patch plan.

4. On the Deployment Options page, do the following:

- a. In the How to Patch section, select the mode of patching that you want to use.

For standalone (single-instance) database targets, Oracle Restart targets, Oracle RAC targets, Oracle Data Guard targets, and Oracle Grid Infrastructure targets (that may or may not be a part of Oracle Exadata), you can choose between in-place patching and out-of-place patching. Out-of-place patching is available for Oracle RAC targets and Oracle Grid Infrastructure targets that are not a part of Oracle Exadata, only if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed. Out-of-place patching is available for Oracle Data Guard targets only if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed.

For Oracle WebLogic Server, Oracle Fusion Applications, and Oracle SOA Infrastructure targets, the only patching mechanism offered is in-place patching. Out-of-place patching is not offered for these targets, see [Overview of Patching in In-Place and Out-of-Place Mode](#).

If you want to clone the existing Oracle home of the database and patch the cloned Oracle home instead of the original Oracle home, select **Out of Place**. If you want to patch the existing original Oracle home of the target directly, without cloning the original Oracle home, select **In Place**.

Also, you can choose to patch Oracle RAC targets, Oracle Grid Infrastructure targets (whether or not they are part of Oracle Exadata), Oracle Data Guard targets, Oracle WebLogic Server targets, Oracle Fusion Application targets, and Oracle SOA Infrastructure targets in rolling mode, or in parallel mode, see [Overview of Patching in Rolling and Parallel Mode](#).

If you want to patch a single node of the target at a time, select **Rolling**. It is the default option, and it involves very little downtime. While patching your targets in rolling mode, you can choose to pause the execution of the patching deployment procedure after each node is patched. For information on how to do so, see [Pausing the Patching Process While Patching Targets in Rolling Mode](#).

If you want to patch all the nodes of the target simultaneously, select **Parallel**. It involves downtime, as your entire system is shut down for a significant period. However, this process consumes less time, as all the target nodes are patched simultaneously.

- b. *(Appears only for standalone database targets, Oracle RAC targets, Oracle Data Guard targets, and Oracle Grid Infrastructure targets)* In the What to Patch section, do the following:

If you have selected in-place patching, then review the details of the Oracle homes that will be patched. By default, all of the database instances are migrated.

For out-of-place patching, select the Oracle home that will be cloned, click **Create New Location** against the Oracle home you want to clone, and enter the full path name to a new directory that will be automatically created during the cloning process. The directory cannot already exist, and the specified operating system user must have write permission for its parent directory.

For standalone database targets, Oracle RAC targets, Oracle Data Guard targets, and Oracle Grid Infrastructure targets (that may or may not be a part of Oracle Exadata), you have the option of migrating either all, or only some of the database instances created from the specified Oracle home. Select the ones that you want to migrate.

 Note:

- After the cloned Oracle home is patched, and after the database instances are migrated to the cloned Oracle home, you can choose to retain the original Oracle home, or delete it if there are no other targets running from that Oracle home.
- If the cloned home contains certain additional patches (that are not added to the patch plan) that include SQL statements, the SQL statements for these patches are not executed automatically when the database instance is migrated to the cloned home. For these patches, you must execute the SQL statements manually.

- c. In the Where to Stage section, do the following:

In the Stage Patches section, select **Yes** for **Stage Patches** if you want the wizard to stage the patches from Software Library to a temporary location, before the patch is applied on the target. Specify the stage location. Select **No** for **Stage Patches** if you have already manually staged the patches that you want to apply. To manually stage a patch, download the patch, navigate to the location (parent directory) where you want to stage the patch, create a subdirectory with the same name as the patch zip file, then extract the contents of the patch zip file into this subdirectory. Specify the parent directory for **Stage Location**. If the stage location is a shared location, select **Shared Location**.

For example, if you downloaded patch 699099.zip, and the stage location, which is the parent directory, is /u01/app/oracle/em/stagepatch, then, in this parent directory, create a subdirectory titled 699099 and extract the contents of the zip file. Specify /u01/app/oracle/em/stagepatch as the stage location.

 Note:

If you are patching a WebLogic Server 10.3.6 target, or an earlier version, and you provide a custom stage location, the location you provide is disregarded, and the selected patches are staged to the default directory configured with SmartUpdate (which is <WEBLOG-IC_HOME>/utils/bsu/cache_dir).

In the Stage Root Component section, specify whether or not you want the wizard to stage the root component. The root component is a set of scripts (which are a part of the directive steps of deployment procedures) which requires *root* privileges for execution. The root component is required for patching Oracle database targets.

Select **Yes** for **Stage Root Component**, then specify the location where you want the root component to be staged (the patch dispatcher location), if you want the wizard to stage the root component. `root_dispatcher.sh` is copied to this location. However, if you have already staged the root component, select **No** for **Stage Root Component**, then specify the location where you staged the root component, that is, the patch dispatcher location. If the patch dispatcher location is shared, select **Dispatcher Location Shared**.

For information on how to stage the patching root component manually, see [Manually Staging the Patching Root Component](#).

- d. In the Credential Information section, provide the required credentials. You can choose to use preferred credentials, or override them with different credentials.

For patching Oracle WebLogic Server targets and Oracle SOA Infrastructure targets, you need the following sets of credentials:

- **Oracle WebLogic Domain Administrator Credentials:** These credentials are required to connect to the Admin Server of the domain which monitors all the Managed Servers present in that domain.

- **Oracle WebLogic Server Home Credentials:** These credentials are required to connect to all the Oracle homes present on different hosts.

You can also choose to override the existing credentials by selecting **Override Preferred WebLogic Domain Administrator Credentials** and **Override Preferred WebLogic Home Credentials**. However, if you choose to override the preferred credentials, then you must validate the credentials. For Oracle WebLogic Server targets, you can validate only the Oracle WebLogic Server Home credentials, and not the administrator credentials.

**Note:**

The validation of credentials fails when the Management Agent is down, or when the credentials are incorrect.

If the Management Agent targets that you want to patch are not secure, then you must set the preferred Management Agent host credentials for all the Management Agent targets that you want to patch. To set the preferred host credentials for Management Agent targets, from the **Setup** menu, select **Security**, then select **Preferred Credentials**. Select the **Agent** target type, then click **Manage Preferred Credentials**. Set the preferred host credentials for the required Management Agent targets.

- e. (Optional) Use the Customization section to customize the default deployment procedure used by the patch plan, and use the customized deployment procedure to patch your targets.

By default, a patch plan uses a static, Oracle-supplied deployment procedure, or an OPlan based, dynamically generated deployment procedure to apply patches.

 Note:

OPlan is an independent tool that creates and executes dynamic deployment procedures to patch your targets. It supports the patching of Oracle Database targets of version 11.2.0.2.0 and above deployed on the Linux x64, Solaris x64, Solaris SPARC, and AIX platforms.

If you have the Enterprise Manager for Oracle Database 12.1.0.6 plug-in (or a higher version) deployed in your system, dynamically generated deployment procedures are used to patch Grid Infrastructure targets (those that are a part of Oracle Exadata, as well as those that are not a part of Oracle Exadata), Oracle RAC database targets, and Oracle Data Guard targets, in both, the in-place and out-of-place patching modes. However, if you do not have the Enterprise Manager for Oracle Database 12.1.0.6 plug-in deployed in your system, dynamically generated deployment procedures are only used to patch Grid Infrastructure targets (those that are a part of Oracle Exadata, as well as those that are not a part of Oracle Exadata) and Oracle RAC database targets in out-of-place patching mode. Static deployment procedures are used in all other patching operations. You can choose to customize both these types of deployment procedures, and use the customized deployment procedure to patch your targets.


If the patching procedure consists of a static deployment procedure, click **Create Like and Edit** to customize the default deployment procedure. To use a customized deployment procedure, select the customized procedure from the list displayed in the Customization section of the Deployment Options page. For more information, see [Customizing a Static Patching Deployment Procedure](#).

If the patching procedure consists of a dynamically generated deployment procedure based on OPlan, select **Specify custom steps to be added to generated patching deployment procedure**, then edit the deployment procedure. For information on how to edit the dynamically generated deployment procedure, see [Customizing a Dynamic Patching Deployment Procedure](#).

If you are patching Oracle WebLogic Server or Oracle SOA Infrastructure targets, you can set a timeout value after which the server is forcefully shut down, if it was not shut down by default. By default, the shutdown time is 30 minutes. You can change this by entering a value in the **Timeout for Shutdown (in minutes)** field. Oracle recommends that you set a timeout value, and ensure that it is based on monitoring of real time systems. If the SOA Infrastructure patch that you are applying involves SQL application functionality, then you must provide the absolute path to the SQL scripts used, in the **Post-Install SQL Script Metadata** field. For information about the SQL script location, refer to the respective readme documents of each patch. Ensure that the SQL scripts that you provide are JDBC-compliant.

To patch SOA Infrastructure targets that are running on a Windows host, ensure that you use the %FMW_ORACLE_HOME% environment variable to provide a

relative path to the SQL files present in the SOA patch, as displayed in the following graphic:



Providing an absolute path, or using the %ORACLE_HOME% environment variable will result in an error. Also, before patching SOA Infrastructure targets that are running on a Windows host, ensure that you shut down all the servers running from the SOA instance home being patched. Stopping just the SOA servers running out of the SOA instance home will result in an error.

- f. In the Notification section, specify whether or not you want to enable email notifications when the patch plan is scheduled, starts, requires action, is suspended, succeeds, and fails.

To enable email notifications, select **Receive notification emails when the patching process**, then select the required options. If a warning message, mentioning that the sender or the receiver email address is not set up, is displayed, perform the action mentioned in the warning.

- g. In the Rollback section, select **Rollback patches in the plan** to roll back the patches listed in the plan, rather than deploy them.

The roll back operation is supported only for Management Agent, Oracle WebLogic Server, Oracle SOA Infrastructure, Oracle Restart, single-instance Oracle database, Oracle RAC database, Oracle RAC One node database, Oracle Grid Infrastructure, and Oracle Grid Infrastructure targets that are part of Oracle Exadata.

For more information on how to roll back a patch, see [Rolling Back Patches](#).

 Note:

- For Oracle WebLogic Server targets, patching and rollback happens at domain level. When Oracle WebLogic Server targets are selected for rollback, the domain along with the Administration Server and the Managed Servers are rolled back. You cannot select the instances you want to rollback, and deselect the ones you do not want to rollback from the domain.
- For SOA Infrastructure targets, patching and rollback happens at instance home level, which means, you can select the SOA Oracle home instances that you want to patch or from where you want to rollback the existing patches. If there are other servers running from the SOA home being rolled back, then all of these servers and their corresponding domains will also be rolled back along with the SQL metadata scripts that can be rolled back. Some of the SQL metadata scripts cannot be rolled back, in which case, Cloud Control will rollback the patch (bits in the Oracle Home) but the SQL remains unchanged.

While patching a SOA setup, if the patch application fails on one of the Managed Servers, and succeeds on other Managed Servers, there will **not** be an automatic rollback operation to remove the patch from all Managed Servers where the patch was successfully applied. However, you will be notified about the failure, so you can manually rollback the patch.

Typically, you will rollback a patch for the following reasons:

- If you had forcefully applied an incoming patch that conflicted with a patch in the Oracle home, and now you want to uninstall that applied patch.
- If the applied patch does not meet your requirements satisfactorily; the patch might have fixed a bug, but at the same time, introduced other bugs in the process.

- h. In the Conflict Check section, specify whether you want to enable or disable ARU Conflict Check, a check that uses Oracle Automated Release Updates (ARU) to search for patch conflicts within the patch plan during the analysis stage. Also, specify the action that the patching procedure must take when a patch conflict is encountered during deployment.

For **Conflicts**, select **Stop at Conflicts** if you want the patching procedure to stop the deployment of the plan when a conflict is encountered, select **Force Apply** if you want the patching procedure to roll back the conflicting patches and apply the incoming patches when a conflict is encountered, or select **Skip conflicts** if you want the patching procedure to apply only the non-conflicting patches, and skip the application of the conflicting patches, when a conflict is encountered.

- i. Click **Next**.
5. On the Validation page, click **Analyze** to check for conflicts. For information about what checks are performed in the validation screen, see [About the Create Plan Wizard](#).
6. On the Review & Deploy page, review the details and do one of the following:

- If you are patching your database targets in out-of-place patching mode, then click **Prepare**. This operation essentially clones the source Oracle home, and patches it. While this happens, the source Oracle home and the database instances are up and running.

Once you click **Prepare**, a Deploy Confirmation dialog box appears, which enables you to schedule the Prepare operation. Select **Prepare**. If you want to begin the Prepare operation immediately, select **Immediately**. If you want to schedule the Prepare operation such that it begins at a later time, select **Later**, then specify the time. Click **Submit**.

After the Prepare operation is successful, click **Deploy**. This operation essentially switches the database instances from the source Oracle home to the cloned and patched Oracle home. The Prepare and Deploy operations enable you to minimize downtime.

Once you click **Deploy**, a Deploy Confirmation dialog box appears, which enables you to schedule the Deploy operation. Select **Deploy**. If you want to begin the Deploy operation immediately, select **Immediately**. If you want to schedule the Deploy operation such that it begins at a later time, select **Later**, then specify the time. Click **Submit**.

**Note:**

Instead of patching your Oracle homes in out-of-place patching mode, you can provision an Oracle home directly from a software image stored in Software Library (that has the required patches). To do this, follow these steps:

- a. Use the Provision Oracle Database deployment procedure to provision the software image stored in Software Library to the required location on the host.
- b. Create a patch plan. Ensure that you add all the patches that are a part of the provisioned Oracle home to this plan.
- c. On the Deployment Options page, in the What to Patch section, specify the location of the provisioned Oracle home.
- d. Analyze and deploy the plan.

The database instance is switched from the old Oracle home to the newly provisioned Oracle home, and the post patching steps are performed.

- If you are patching any other target in any other mode, click **Deploy**.

Once you click **Deploy**, a Deploy Confirmation dialog box appears, which enables you to schedule the Deploy operation. Select **Deploy**. If you want to begin the Deploy operation immediately, select **Immediately**. If you want to schedule the Deploy operation such that it begins at a later time, select **Later**, then specify the time. Click **Submit**.

 Note:

- After scheduling a Prepare or Deploy operation, the **Prepare or Deploy** button on the Review and Deploy page is renamed to **Reschedule**. If you want to reschedule the Prepare or Deploy operation, click **Reschedule**, specify the time, then click **Submit**.
- After scheduling a Prepare or Deploy operation, if you want to discard the schedule and bring the patch plan back to its last valid state, click **Stop Schedule**.
- The Prepare or Deploy operation schedule is discarded if you edit a patch plan deployment option or patch target. In this case, you must validate the patch plan again.

Many patching operations use OPlan, an independent tool that creates and executes dynamic deployment procedures to patch your targets. The OPlan readme file for a patch plan contains detailed information about the step wise execution of the patch plan. If you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in deployed, you can view this file. To view this file, on the Review & Deploy page, click the **Download** link present beside **Patching Steps**.

 Note:

To save a successfully analyzed or deployed patch plan as a patch template, see [Saving Successfully Analyzed or Deployed Patch Plan As a Patch Template](#).

Conflict Patches

When applying WebLogic Server, SOA, or Service Bus patches, there may be conflicts among the patches. This is encountered during the analysis phase. We can select the appropriate patch conflict resolution action in the **Deployment Options** page. Refer to the following table for details.

Patch Conflict resolution is supported from WebLogic Server 12.1.2.0.0 onwards.

The following tasks or actions can be performed by the user when the analysis results display any patch conflicts:

Selected action	Description
Stop at Conflicts	This can be selected if you want the patching procedure to stop the deployment of the plan when a conflict is encountered.
Force Apply	This can be selected if you want the patching procedure to roll back the conflicting patches and apply the incoming patches when a conflict is encountered.
Skip Conflicts	This can be selected if you want the patching procedure to apply only the non-conflicting patches and skip applying the conflicting patches when a conflict is encountered.

For WebLogic Server versions earlier than 12.1.2.0.0, the patch conflicts can be removed using the following steps.

1. Create a Patch Plan with a new Patch Set Update.
2. Analyze the patch. In case there are any conflicting patches, the **Run WebLogic Home Prerequisite Checks** indicates the conflicting patches with the new patches that are part of the new Patch Plan just created.
3. Make a note of the patch number that is conflicting with the newly created Patch Plan and that needs to be rolled back.
4. Create a new Rollback plan that includes the conflicting patches for all the targets.
5. Analyze the Rollback plan.
6. Apply the Rollback plan.
7. All the conflicting patches are now removed from the WebLogic Server Home. Now the Patch Plan that was initially created with the Patch Set Update can be analyzed.
8. Apply the Patch Plan.

Switching Back to the Original Oracle Home After Deploying a Patch Plan

If you had patched an Oracle RAC target, Oracle single-instance database target, Oracle Restart target, or an Oracle Grid Infrastructure target (that may or may not be a part of Oracle Exadata), in out-of-place patching mode, and you now want to switch back to the original home for some reason, then you can use the *Switchback* option available in the Create Plan Wizard. The advantage of using this option is that you do not actually roll back the patches from the cloned and patched Oracle home; you only switch back to the old, original Oracle home that exists without the patches.

Note:

- The *Switchback* option is available only for Oracle RAC, Oracle single-instance database, Oracle Restart, and Oracle Grid Infrastructure targets (that may or may not be a part of Oracle Exadata), and only when these targets are patched in out-of-place patching mode.
- You can switch back only if you have successfully analyzed and deployed a patch plan.

To switch back to the original Oracle home after a patch plan is deployed, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, click the successfully analyzed and deployed patch plan you used for patching the Oracle RAC, Oracle Restart, Oracle single-instance database, or Oracle Grid Infrastructure targets. Alternatively, select the patch plan, and in the context menu, click **View**. The Create Plan Wizard appears.

3. In the Create Plan Wizard, on the Review & Deploy page, click **Switchback**.

Saving Successfully Analyzed or Deployed Patch Plan As a Patch Template

NOT_SUPPORTED:



This section is mainly for Patch Designers who want to save the successfully analyzed or deployed patch plans as patch templates so that operators can consume them to create fresh patch plans with the approved patches and predefined deployment options.

To save a patch plan as a patch template, follow Step (1) to Step (5) as outlined in [Analyzing, Preparing, and Deploying Patch Plans](#), and then for Step (6), on the Review & Deploy page, click **Save as Template**. In the Create New Plan Template dialog, enter a unique name for the patch template, and click **Create Template**.

! Important:



Oracle recommends you to follow this as a best practice if you have to roll out in a mass scale over a period of time involving more administrators. If you have a large data center, then as a Patch Designer, create a patch plan and apply the patches on a few databases, test if the patches are being applied successfully, then save the plan as a template. Later, have your Patch Operators create patch plans out of these templates so that they can roll out the same set of patches to the rest of the databases in the data center.

Creating a Patch Plan from a Patch Template and Applying Patches

Once a successfully analyzed or deployed patch plan is saved as a patch template, you can create patch plans out of the template, associate targets you want to patch, and deploy the newly created patch plan.

This is purely an optional step. You do not have to save your patch plans as patch templates to roll out patches. You can roll out patches directly from a patch plan as described in [Analyzing, Preparing, and Deploying Patch Plans](#).

NOT_SUPPORTED:



This section is mainly for Patch Operators who want to create patch plans from patch templates for rolling out the patches.

Approach 1

To create patch plans out of the patch templates, use one of the following approaches:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, select the patch template you want to use to create a patch plan out of it.
3. From the context menu, select **Create Plan**.
4. In the Create Plan from Template dialog, enter a unique name for the patch plan, select the targets on which you want to patch, and click **Create Plan**.
5. Return to the Patches & Updates page, and in the Plans region, click the patch plan you want to use. Alternatively, select the patch plan, and in the context menu, click **View**. The Create Plan Wizard appears.
6. In the Create Plan Wizard, go to the Validation page, and click **Re-Analyze** to analyze the patch plan with the newly associated targets.
7. After successfully analyzing the patch plan, on the Validation page, click **Next**.
8. On the Review & Deploy page, click **Deploy**.

Approach 2

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.

2. On the Patches & Updates page, in the Plans region, do one of the following:
 - Select a patch template. From the context menu, select **View**. The Edit Template Wizard appears.
 - Click the name of a patch template. The Edit Template Wizard appears.
3. In the Edit Template Wizard, click **Create Plan**.
4. In the Create Plan from Template dialog, enter a unique name for the patch plan, select the targets on which you want to patch, and click **Create Plan**.
5. Return to the Patches & Updates page, and in the Plans region, click the patch plan you want to use. Alternatively, select the patch plan, and in the context menu, click **View**. The Create Plan Wizard appears.
6. In the Create Plan Wizard, go to the Validation page, and click **Re-Analyze** to analyze the patch plan with the newly associated targets.
7. After successfully analyzing the patch plan, on the Validation page, click **Next**.
8. On the Review & Deploy page, click **Deploy**.

Patching Oracle Grid Infrastructure Targets

You can patch Oracle Grid Infrastructure targets using patch plans. To do so, follow these steps:

1. Identify the Oracle Grid Infrastructure Patch Set Update (PSU) or one-off patches that you want to apply, as described in [Identifying the Patches to Be Applied](#).
2. Create a patch plan as described in [Creating a Patch Plan](#).
3. Analyze and deploy the patch plan as described in [Analyzing, Preparing, and Deploying Patch Plans](#).

 Note:

- You can add one Grid Infrastructure PSU, and any number of one-off Grid Infrastructure and Oracle Database patches to a single patch plan, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed. In this scenario, ensure that you add the Grid Infrastructure PSU while creating the patch plan, and then add the one-off Grid Infrastructure and Oracle Database patches as additional patches.

If you do not have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed, then you cannot add one-off patches to a patch plan that already contains a Grid Infrastructure PSU.

- You can apply a Grid Infrastructure PSU on a Grid Infrastructure target when the RAC databases being managed are of different versions, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed. If you do not have this plug-in deployed, and the RAC databases being managed are of different versions, then you cannot apply a Grid Infrastructure PSU on the Grid Infrastructure target.

For example, if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed, and a Grid Infrastructure target consists of Grid Infrastructure of version 12.1.0.2, and manages RAC databases of versions 11.2.0.4 and 12.1.0.2, you can apply a 12.1.0.2 Grid Infrastructure PSU on the Grid Infrastructure target. In this case, the Grid Infrastructure binaries and the 12.1.0.2 RAC databases are patched, but not the 11.2.0.4 RAC databases. You can also apply a 11.2.0.4 Grid Infrastructure PSU on the Grid Infrastructure target. In this case, the 11.2.0.4 RAC databases are patched, but not the Grid Infrastructure binaries and the 12.1.0.2 RAC databases.

Patching Oracle Exadata

Using patch plans, you can patch the Oracle RAC database targets and the Oracle Grid Infrastructure targets (Oracle Clusterware) running on an Exadata Database Machine.

Exadata Patching can be performed in two modes: In Place Patching, and Out-of-place Patching, though Oracle recommends you to use the Out-of-place patching mechanism as the downtime involved is much lesser.

For information about the supported Exadata releases, see [Supported Targets, Releases, and Deployment Procedures for Patching](#).

However, note that patch plans do not patch the Exadata Database Machine's compute node entities such as the operating system and firmware, and also its storage cells. They patch only the associated Oracle RAC database targets and the Oracle Grid Infrastructure (Oracle Clusterware) targets running on the machine.

**Note:**

Oracle Exadata Database Machine recommended bundle patches that apply to Oracle RAC and the Oracle Grid Infrastructure targets (Oracle Clusterware) are tested and certified with Cloud Control.

Therefore, when you create a patch plan with an Exadata Database Machine recommended bundle patch, make sure you add the cluster or the Oracle RAC database target running on the Exadata Database machine. The patch plan automatically recognizes their association with the Exadata Database machine, and prompts you to add all the impacted targets running on that machine. For example, if you select the cluster target, it prompts you to add all the Oracle RAC database targets and the Oracle Grid Infrastructure targets that are part of that cluster running on the Exadata Database Machine.

To patch an Exadata Database machine, follow these steps:

1. Identify the Exadata Database Machine recommended bundle patch you need to apply, as described in [Identifying the Patches to Be Applied](#).
2. Create a patch plan as described in [Creating a Patch Plan](#).
3. Add the cluster or the Oracle RAC database target running on the Exadata Database machine, and analyze and deploy the patch plan as described in [Analyzing, Preparing, and Deploying Patch Plans](#).

 Note:

- You can add one Exadata Bundle Patch, and any number of one-off Grid Infrastructure and Oracle Database patches to a single patch plan, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed. In this scenario, ensure that you add the Exadata Bundle Patch while creating the patch plan, and then add the one-off Grid Infrastructure and Oracle Database patches as additional patches.

If you do not have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed, then you cannot add one-off patches to a patch plan that already contains an Exadata Bundle Patch.

- You can apply an Exadata Bundle Patch on a Grid Infrastructure target (that is a part of Oracle Exadata) when the RAC databases being managed are of different versions, as long as you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed. If you do not have this plug-in deployed, and the RAC databases being managed are of different versions, then you cannot apply an Exadata Bundle Patch on the Grid Infrastructure target.

For example, if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed, and a Grid Infrastructure target (that is a part of Oracle Exadata) consists of Grid Infrastructure of version 12.1.0.2, and manages RAC databases of versions 11.2.0.4 and 12.1.0.2, you can apply a 12.1.0.2 Exadata Bundle Patch on the Grid Infrastructure target. In this case, the Grid Infrastructure binaries and the 12.1.0.2 RAC databases are patched, but not the 11.2.0.4 RAC databases. You can also apply a 11.2.0.4 Exadata Bundle Patch on the Grid Infrastructure target. In this case, the 11.2.0.4 RAC databases are patched, but not the Grid Infrastructure binaries and the 12.1.0.2 RAC databases.

Exadata Out-Of-Place Patching Of Oracle Grid Infrastructure and Oracle RAC Targets

Exadata Out-of-place patching mechanism allows you patch the Oracle Grid Infrastructure and Oracle RAC targets by making a copy of their existing Oracle homes, and patching the cloned Oracle homes. Once the cloned homes are patched, you can migrate the instances to run from the cloned home, which means there will be minimal downtime while switching over the instances, but not a significant downtime.

[Figure 44-10](#) illustrates how Oracle Grid Infrastructure and Oracle RAC targets running on an Exadata Database Machine get patched in Out-of-place patching mode:

Figure 44-10 Out Of Place Patching Of Clusters



The migration of these instances can be performed in the following modes:

- **Full Migration:** If you choose to migrate all the database instances running in your data center in one session, then it is termed as Full Migration.
- **Partial Migration:** If you choose to migrate only some of the instances depending on the downtime you can afford in your data center in one session, then it is termed as Partial Migration. However, you must ensure that you migrate the remaining instances in the following sessions. This approach is particularly useful when you have multiple instances running out of an Oracle home.

Note:

For steps on how to perform Full Migration or Partial Migration of Oracle Grid Infrastructure Targets and Oracle RAC Targets running on Exadata Database Machine, see [Analyzing, Preparing, and Deploying Patch Plans](#).

Switch Back is an option available for Oracle Grid Infrastructure targets and Oracle RAC targets running on Exadata machine, that enables you to switch the instances from the newly cloned Oracle homes back to the original Oracle homes in case of any problems.

For more information on how to perform a Switch Back, see [Switching Back to the Original Oracle Home After Deploying a Patch Plan](#).

Patching Oracle Data Guard Targets

This section describes how to apply an Oracle Grid Infrastructure PSU or an Oracle Database PSU on various configurations of Oracle Data Guard targets. You can choose to patch your Oracle Data Guard targets in the in-place patching mode, or the out-of-place patching mode.

It consists of the following sections:

- [Oracle Data Guard Patching Workflow](#)
- [Oracle Data Guard Patching Scenarios](#)

Oracle Data Guard Patching Workflow

If the primary and the standby databases are running from the same cluster, follow these steps:

1. Search for the required patch, then create a patch plan by selecting the cluster target.
2. Specify the required deployment options.
3. Analyze and deploy the plan.

However, if the primary and the standby databases are running from different clusters, follow these steps:

1. Search for the required patch, then create a patch plan by selecting the cluster target that contains the standby database.
2. Specify the required deployment options.
3. Analyze and deploy the plan.
4. *(Optional)* Perform a database switchover, such that the standby database becomes the new primary database, and vice versa. Performing a switchover can save database downtime.

For information on how to perform a database switchover, see *Oracle® Data Guard Broker*.

5. Search for the required patch, then create another patch plan by selecting the cluster target that contains the primary database (or the new standby database, in case you have performed a switchover).
6. Specify the required deployment options.
7. Analyze and deploy the plan.

**Note:**

- Oracle recommends that you patch the standby database first, then patch the primary database.
- If your environment consists of multiple standby databases, ensure that you patch all the standby databases first, then patch the primary database. In case you want to perform a switchover, ensure that you patch a single standby database first, perform the switchover, then patch the remaining standby databases and the new standby database.

Oracle Data Guard Patching Scenarios

This section describes the steps to patch your Oracle Data Guard targets in various scenarios.

Scenario 1: The primary and standby databases are RAC databases.

Table 44-5 describes the steps to apply an Oracle Grid Infrastructure PSU or an Oracle Database PSU, when the primary and standby databases are RAC databases:

Table 44-5 Oracle Data Guard Patching (RAC - RAC)

Primary and Standby Database Configuration	How to Patch?
The primary and the standby RAC databases are running from the same cluster.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target. 2. Specify the required deployment options. 3. Analyze and deploy the plan.
The primary and the standby RAC databases are running from different clusters.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target that contains the standby RAC database. 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. <i>(Optional)</i> Perform a database switchover, such that the standby RAC database becomes the new primary database, and vice versa. For information on how to perform a database switchover, see <i>Oracle® Data Guard Broker</i>. 5. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the cluster target that contains the primary RAC database (or the new standby database, in case you have performed a switchover). 6. Specify the required deployment options. 7. Analyze and deploy the plan.

Table 44-5 (Cont.) Oracle Data Guard Patching (RAC - RAC)

Primary and Standby Database Configuration	How to Patch?
Two primary RAC databases, P1 and P2, are running from two different clusters, C1 and C2, respectively. Their standby RAC databases, S1 and S2, are running from C2 and C1, respectively.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the first cluster target (C1). 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the second cluster target (C2). 5. Specify the required deployment options. 6. Analyze and deploy the plan. Note: You can patch the clusters in any order, that is, C1 first and then C2, or C2 first and then C1.

Scenario 2: The primary database is a RAC database, and the standby database is a single-instance database that is managed by a Cluster Ready Services (CRS) target or a Single Instance High Availability (SIHA) target.

Table 44-6 describes the steps to apply an Oracle Grid Infrastructure PSU or an Oracle Database PSU, when the primary database is a RAC database, and the standby database is a single-instance database that is managed by a CRS target or a SIHA target.

Table 44-6 Oracle Data Guard Patching (RAC - SIDB)

Primary and Standby Database Configuration	How to Patch?
The primary RAC database and the standby single-instance database are running from the same cluster.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target. 2. Specify the required deployment options. 3. Analyze and deploy the plan.

Table 44-6 (Cont.) Oracle Data Guard Patching (RAC - SIDB)

Primary and Standby Database Configuration	How to Patch?
The primary RAC database and the standby single-instance database are running from different clusters.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target that contains the standby single-instance database. 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. (Optional) Perform a database switchover, such that the standby single-instance database becomes the new primary database, and vice versa. For information on how to perform a database switchover, see <i>Oracle® Data Guard Broker</i>. 5. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the cluster target that contains the primary RAC database (or the new standby database, in case you have performed a switchover). 6. Specify the required deployment options. 7. Analyze and deploy the plan.
Two primary RAC databases, P1 and P2, are running from two different clusters, C1 and C2, respectively. Their standby single-instance databases, S1 and S2, are running from C2 and C1, respectively.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the first cluster target (C1). 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the second cluster target (C2). 5. Specify the required deployment options. 6. Analyze and deploy the plan. Note: You can patch the clusters in any order, that is, C1 first and then C2, or C2 first and then C1.

Scenario 3: The primary and standby databases are single-instance databases that are managed by CRS or SIHA targets.

Table 44-7 describes the steps to apply an Oracle Grid Infrastructure PSU or an Oracle Database PSU, when the primary and standby databases are single-instance databases that are managed by CRS or SIHA targets.

Table 44-7 Oracle Data Guard Patching (SIDB - SIDB)

Primary and Standby Database Configuration	How to Patch?
The primary and the standby single-instance databases are running from the same cluster.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target. 2. Specify the required deployment options. 3. Analyze and deploy the plan.
The primary and the standby single-instance databases are running from different clusters.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the cluster target that contains the standby single-instance database. 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. (Optional) Perform a database switchover, such that the standby single-instance database becomes the new primary database, and vice versa. For information on how to perform a database switchover, see <i>Oracle® Data Guard Broker</i>. 5. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the cluster target that contains the primary single-instance database (or the new standby database, in case you have performed a switchover). 6. Specify the required deployment options. 7. Analyze and deploy the plan.
Two primary single-instance databases, P1 and P2, are running from two different clusters, C1 and C2, respectively. Their standby single-instance databases, S1 and S2, are running from C2 and C1, respectively.	1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the first cluster target (C1). 2. Specify the required deployment options. 3. Analyze and deploy the plan. 4. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the second cluster target (C2). 5. Specify the required deployment options. 6. Analyze and deploy the plan. Note: You can patch the clusters in any order, that is, C1 first and then C2, or C2 first and then C1.

Scenario 4: A primary RAC database, P1, and a primary single-instance database, P2 (that is managed by a CRS or a SIHA target), are running from two different clusters, C1 and C2, respectively. Their standby single-instance databases, S1 and S2, that are managed by CRS or SIHA targets, are running from C2 and C1, respectively.

Follow these steps to apply an Oracle Grid Infrastructure PSU or an Oracle Database PSU on this configuration:

1. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create a patch plan by selecting the first cluster target (C1).
2. Specify the required deployment options.
3. Analyze and deploy the plan.
4. Search for the required Oracle Grid Infrastructure PSU, or the Oracle Database PSU, then create another patch plan by selecting the second cluster target (C2).
5. Specify the required deployment options.
6. Analyze and deploy the plan.

**Note:**

You can patch the clusters in any order, that is, C1 first and then C2, or C2 first and then C1.

Patching Oracle Identity Management Targets

Using patch plans, you can patch those Oracle Access Management Server and Oracle Identity Management Server targets that were provisioned using Identity Management Lifecycle tools. Other Oracle Identity Management targets are not supported for patching.

To patch Oracle Access Management Server and Oracle Identity Management Server targets (that were provisioned using Identity Management Lifecycle tools) using patch plans, you must have an Identity Management Pack Plus license. Also, the 12.1.0.6 Enterprise Manager for Oracle Fusion Middleware plug-in must be deployed on the OMS, and on all the Management Agents running on the hosts on which the Oracle Access Management Server and Oracle Identity Management Server targets are deployed.

To patch Oracle Access Management Server and Oracle Identity Management Server targets using a patch plan, follow these steps:

1. Identify the patches that you want to apply, as described in [Identifying the Patches to Be Applied](#).
2. Create a patch plan, as described in [Creating a Patch Plan](#).
3. Analyze and deploy the patch plan, as described in [Analyzing, Preparing, and Deploying Patch Plans](#).

Patching Oracle Siebel Targets

Using patch plans, you can patch Siebel Gateway Server and Siebel Server targets. Other Oracle Siebel targets are not supported for patching. To patch Siebel Gateway Server and Siebel Server targets, you must have the 12.1.0.5 Enterprise Manager for Oracle Siebel plug-in deployed in your system.

To patch Siebel Gateway Server and Siebel Server targets using a patch plan, follow these steps:

1. Identify the patches that you want to apply, as described in [Identifying the Patches to Be Applied](#).

2. Create a patch plan, as described in [Creating a Patch Plan](#).
3. Analyze and deploy the patch plan, as described in [Analyzing, Preparing, and Deploying Patch Plans](#).

Patching Oracle Service Bus

Fusion Middleware patching through Enterprise Manager is supported for targets like WebLogic Domain, SOA, and Oracle Service Bus (OSB). This topic contains the procedure to patch an Oracle Service Bus using the Enterprise Manager.

A single node of Oracle Service Bus Domain consists of Service Bus Domain Target, Service Bus Cluster Target, Service Bus Target, WebLogic Cluster Target, and WebLogic Server Targets.

The Service Bus Cluster Target comprises of all the Service Bus targets. The WebLogic Cluster Target comprises all the WebLogic Server targets. The WebLogic Server Targets comprise of a single Administration Server and a Managed Server.

This procedure of patching an Oracle Service Bus through the Enterprise Manager patching framework comprises of creating a patch plan by selecting a Service Bus Target or a Service Bus Cluster Target. The user can select a single target and the associated targets are included automatically during patching. The user can also edit the plan to add more patches, analyze, rollback, and create a template.

For more information on creating a patch plan, see [Creating a Patch Plan](#).

Before you begin patching of Oracle Service Bus, ensure that you meet the following prerequisites:

- The user needs to ensure that Oracle Home collection is complete. Patching throws an error if one or more Oracle Homes are not discovered.



Note:

The pre-requisite for patching WebLogic domains is that the underlying Oracle Home must be discovered as a target. For SOA 11g, 12c and Service Bus domains this is automatically handled. However, for Service Bus 11g when the domain is discovered the OH is not discovered as a target. If you attempt to patch such a domain then the following warning is displayed. To overcome this, ensure that the Middleware Home and the underlying Oracle Homes are discovered as targets and the configuration collections for Oracle Home targets are the current ones.

There is no Oracle Home (OH) target associated for the patch being applied. To rectify this issue the user should promote the OH being patched, by choosing the "Discover Promote Oracle Home Target" job in Enterprise Manager. Ensure the OH target is available after this operation is completed, and the configuration metrics for the OH target is collected.

- The user needs to ensure to set the Oracle Homes and Domain administrator credentials.

- For offline patching, the user needs to ensure to have the necessary OPatch downloaded from MOS and manually upload to the saved patches.
For patching WebLogic Server 12.2.1.x Enterprise Manager expects OPatch release 13.3.0.0.0 to be available in the software library. At the time of releasing Enterprise Manager, this version of OPatch was not available for download from MOS. Hence, when patching WebLogic Server 12.2.1.x, Enterprise Manager uses the OPatch installed in the WebLogic Server Oracle Home instead of trying to download the latest OPatch from MOS.
- The target selector needs metadata from **My Oracle Support** to filter the targets based on the patch types. The user must run **Refresh from My Oracle Support** job in online mode. In case of offline mode the user must upload the *catalog.zip* file.
- Patching on Oracle Service Bus is supported on the following versions:
 - 11.1.1.5.0
 - 11.1.1.6.0
 - 11.1.1.7.0
 - 11.1.1.9.0
 - 12.1.3.0.0
 - 12.2.1.0.0
- The SOA patch containing SQL script may fail if the metadata is not specified in the right format. The metadata has to be specified in the **Deployment Options** page in the field **Location of Post Install SQL Script**.

Ensure that the SQL scripts are specified in the following format:

```
<patch_number1>:<sql_script1_location>:<sql_script2_location>:<patch_number2>:<sql_script3_location>:...
```

Example:

```
14082705:%ORACLE_HOME%/rcu/integration/soainfra/sql/createschema_soa_infra_oracle.sql
```

Additionally, this information is available in the **Post-Installation Instructions** in the *README.txt* file. In case the **Location of Post Install SQL Script** field is left blank or the right SQL script is not specified then the following error message is mentioned in the log:

One or more SQL files included in this patch are not selected for post-patch execution. To correct this, update your selection in the "Deployment options" page. To continue without making any changes, run the required SQL scripts manually after the patch is deployed successfully.

- The relevant patch, target, and domain details are displayed in the following table:

Patch	Targets
SOA	SOA Infrastructure
Oracle Service Bus	Service Bus

To apply patching for Oracle Service Bus, follow these steps:

- Create a patch plan.

2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Saved Patches**.
3. Click **Upload**.
4. Provide the location of the patch file and the patch metadata, and click **Upload**.
5. Once the patch has been uploaded, select the patch. This takes you to the patch home page.
6. Click **Create Plan**.
7. Enter the name of the plan in the **Plan Name** field.
8. Enter Oracle Service Bus or Service Cluster Targets As the target type.

**Note:**

In case a single Service Bus component is selected, all the associated targets are automatically picked up for patching. Ensure to select one Service Bus component.

9. Enter the host name and click **Create Plan**.
10. The patch plan is created and can be deployed.
11. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
12. Select the patch plan that you created from the list of patch plans.
13. Click **Next**.
14. Enter the rest of the details in the EM User Interface, validate the credentials, and choose the customization required. Ensure that the data prefilled or selected is accurate.
15. Click **Analyze**. The deployment procedure is submitted.
16. Once analysis is complete click **Review**, for the tool to review, analyze, and validate the details.
17. Click **Next**.
18. Click **Deploy**. The deployment procedure is submitted and patching is in progress.
19. The patch is deployed.
20. The Oracle Service Bus patch is successfully applied.

**Note:**

For Oracle Service Bus 11g targets patch recommendations are not displayed. The user can search for the recommended patches through the **Recommended Patch Advisor** search in the **Patches and Updates**.

Rollback of Oracle Service Bus Patch

To apply patching for Oracle Service Bus, follow these steps:

1. Create a patch plan.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Saved Patches**.
3. Click **Upload**.
4. Provide the location of the patch file and the patch metadata, and click **Upload**.
5. Once the patch has been uploaded, select the patch. This takes you to the patch home page.
6. Click **Create Plan**.
7. Enter the name of the plan in the **Plan Name** field.
8. Enter Oracle Service Bus or Service Cluster Targets As the target type.

 Note:

In case a single Service Bus component is selected, all the associated targets are automatically picked up for patching. Ensure to select one Service Bus component.

9. Enter the host name and click **Create Plan**.
10. The patch plan is created and can be deployed.
11. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
12. Select the patch plan that you created from the list of patch plans.
13. Click **Next**.
14. Choose the **Rollback** option. Enter the rest of the details in the EM User Interface, validate the credentials, and choose the customization required. Ensure that the data prefilled or selected is accurate.
15. Click **Analyze**. The deployment procedure is submitted.
16. Once analysis is complete click **Review**, for the tool to review, analyze, and validate the details.
17. Click **Next**.
18. Click **Rollback**.
19. The patch is rolled back.
20. The Oracle Service Bus patch is successfully rolled back.

Deploying WebLogic Patches Along with SOA or Oracle Service Bus Patches In A Single Patch Plan

This topic contains the procedure to deploy WebLogic Patches as well as SOA or Oracle Service Bus patches using a single patch plan. Using a single patch plan to deploy all the required patches would reduce down time during patching.

In case of SOA or Oracle Service Bus domain the user can apply Web Logic patches or SOA or Oracle Service Bus product patches. These patches can be applied in dif-

ferent sessions by creating separate patch plans. However each patch deployment session would involve down time. To reduce the system down time, you can create one patch plan by selecting the Web Logic platform patches as well as SOA or Oracle Service Bus patches and then apply this patch plan in one session.

 Note:

- The combined patching is supported only from Oracle Service Bus or SOA release 12.1.3.0.0 onwards. This mode of combined patching is supported only from Oracle Service Bus or SOA release 12.1.3.0.0 onwards. For earlier versions of Oracle Service Bus and SOA, create separate patch plans to apply the patches.
- To achieve this, first create the patch plan by selecting the desired Weblogic patches and selecting the desired domain targets. Add the Oracle Service Bus or SOA patches to the same patch plan and ensure to reselect the same domain targets in the target selector window as detailed in the steps below. This will ensure that all required patches and the targets to deploy them against are selected.

The user can apply multiple patches through a single patch plan.

For more information on creating a patch plan, see [Creating a Patch Plan](#).

 Note:

An Oracle Service Bus or SOA target must have an Oracle Home associated with it. Else, the following error message is displayed with the resolution:

There is no Oracle Home(OH) target associated for the patch being applied. To rectify this issue the user should promote the OH being patched, by choosing the "Discover Promote Oracle Home Target" job in Enterprise Manager. Ensure the OH target is available after this operation is completed, and the configuration metrics for the OH target is collected.

To apply patching for Oracle Service Bus, follow these steps:

 Note:

You can also include any SQL scripts that need to be executed while applying patch.

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Saved Patches**.
2. Click **Upload**.

3. Provide the location of the patch file and the patch metadata, and click **Upload**.
4. Once the patch has been uploaded, select the patch. This takes you to the patch home page.
5. Click **Create Plan** by selecting the WebLogic patch and WebLogic domain As the target.
6. Enter the name of the plan in the **Plan Name** field.
7. Enter WebLogic domain As the target type.
8. Click **Create Plan**.
9. The WebLogic Server patch plan is created and can be deployed along with Oracle Service Bus or SOA patch.
10. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
11. Select the patch plan that you created from the list of patch plans.
12. Click **Next**.
13. In the **Patches & Updates**, under the **Patches** tab, click on **Add Patch**.
14. Search and select a SOA or Oracle Service Bus patch and select the same domain target.
15. Enter the rest of the details in the EM User Interface, validate the credentials, and choose the customization required. Ensure that the data prefilled or selected is accurate.
16. Click **Analyze**. The deployment procedure is submitted.
17. Once analysis is complete click **Review**, for the tool to review, analyze, and validate the details.
18. Click **Next**.
19. Click **Deploy**. The deployment procedure is submitted and patching is in progress.
20. The patch is deployed.
21. The WebLogic Server patch is successfully applied.

Diagnosing and Resolving Patching Issues

This section describes the following:

- [Workaround for Errors](#)
- [Common Patching Issues](#)
- [Resolving Patching Issues](#)
- [Rolling Back Patches](#)

Workaround for Errors

This section describes the workaround for the following errors:

- [Workaround for Missing Property Errors](#)
- [Workaround for Unsupported Configuration Errors](#)

Workaround for Missing Property Errors

Table 44-8 describes the possible missing property errors and the workarounds you can use to resolve them.

Table 44-8 Missing Properties Error and Workaround

Problem	Workaround
Empty target property version	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 the Management Repository is updated with the latest configuration of the host. To do so, follow these steps: - To recompute the dynamic properties, do one of the following: Option A: Stop and restart the Management Agent. This option is simpler because you do not have to understand the target model. <pre>\$ emctl stop agent \$ emctl start agent</pre> For a cluster, restart the Management Agent on all the nodes of the cluster. Option B: Reload the dynamic properties of all the targets. This option is better because there is no blackout or downtime for monitoring. To view the list of targets that are monitored by the Management Agent, run the following command: <pre>\$ emctl config agent listtargets</pre> To reload the dynamic properties of a specific target, run the following command: <pre>\$ emctl reload agent dynamicproperties [<Target_name>:<Target_Type>]</pre> For example: <pre>\$ emctl reload agent dynamicproperties oradb:oracle_database \$ emctl reload agent dynamic properties racdb_1:rac_database \$ emctl reload agent dynamicproperties crs:cluster \$ emctl reload agent dynamic properties wls:weblogic_j2eeserver \$ emctl reload agent dynamicproperties server1.xyz.com:host</pre> - To update the Management Repository with the latest configuration information, from the Enterprise menu, select Configuration, and then select Refresh Host Configuration. On the Refresh Host Configuration page, select the hosts for which the configuration must be updated, and click Refresh Hosts.
Targets are not properly discovered because of inadequate or incomplete target information collected during discovery.	To resolve this issue, rediscover the domain so that all the targets in the domain are discovered effectively. To do so, follow these steps: - Log in to the Domain Home page using appropriate credentials. For example, Farm01_base_domain. - On the (Farm01_base_domain) home page, from the Farm menu, select Refresh WebLogic Domain, and click Ok on all the following pages to complete the process. After successful completion of the process the domain home page is refreshed to discover all the targets afresh.

Workaround for Unsupported Configuration Errors

Table 44-9 describes the possible unsupported configuration errors and the workarounds you can use to resolve them.

Table 44-9 Workaround for Unsupported Configuration Errors

Problem	Workarounds
Oracle RAC Instance does not have an associated Oracle RAC Database	Rediscover the Oracle RAC target and add the Oracle RAC instance to the Oracle RAC database.
The database is not mediated by the OMS	The target discovery is not appropriate. Remove the target from Cloud Control, and rediscover on all the Management Agents in the cluster.

Common Patching Issues

While analyzing or deploying patch plans, you might see errors if the following are true:

- If the OMS or the Management Agent is down
- If the Software Library is not properly configured
- If there is no connectivity to My Oracle Support (in online mode)
- If there is no Management Repository
- If there are no collections
- If there are host or Oracle home security issues
- If there are inherent OPatch or SQL errors
- If the patch plan consists of *non-homogenous* patches, for example, a combination of one-off patches and patch sets
- (For Oracle WebLogic Targets only) If there are inherent problems with the SmartUpdate tool.
- (For Oracle WebLogic Targets only) If there is problem discovering Oracle WebLogic targets reporting to the OMS.

You will see these errors in the following places:

- In the header section of the Validation page or the Review page in the Create Plan Wizard (Figure 44-11)
- In the Issues to Resolve section of the Validation page (Figure 44-12)
- In the Plans region of the Patches & Updates page (Figure 44-13).

Figure 44-11 Patch Plan Errors Displayed on the Validation Page




Figure 44-12 Patch Plan Errors Displayed in the Issues to Resolve Section




Figure 44-13 Patch Plan Errors Displayed in the Plans Region

Name	Status	Type	Planned Deploy	Created	Deploys	Plan Privile
st_ui_test3	New	Patch	Not Specified	SYSM...	Yes	Full
try1	Switched back Successfully	Patch	Not Specified	SYSM...	Yes	Full
try2	Review Analysis	Patch	Not Specified	SYSM...	Yes	Full
vt1	Deploy Failed	Patch	Not Specified	SYSM...	Yes	Full
vt2	New	Patch	Not Specified	SYSM...	Yes	Full
vt3	Review Analysis	Patch	Not Specified	SYSM...	Yes	Full

Resolving Patching Issues

Table 44-10 lists the different phases that the patch plan goes through, the different states the phases can have, and how you can diagnose and resolve the errors.



Note:

Also, refer to the troubleshooting tips described in [Troubleshooting Patching Issues](#).

Table 44-10 Diagnosing Patching Issues

Phase	State	Diagnosis and Resolution
Analysis	Analysis In Progress	N/A
	Analysis Error	Review the following log file: <gc_inst>/sysman/log/emoms.log
	Issues Remain	In the Create Plan Wizard, on the Validation page, review the issues listed in the Issues to Resolve section. In the Issues to Resolve section, if an error message states that you must click Show Detailed Results here , then click it. On the Procedure Activity Status page, in the Status Detail table, review the status of each of the steps. Click the status to view the log details. In case the Issues to Resolve table is empty, then exit and re-enter the plan.
	Conflicts Detected	In the Create Plan Wizard, on the Validation page, review the issues listed in the Issues to Resolve section. In the Issues to Resolve section, if an error message states that you must click Show Detailed Results here , then click it. On the Procedure Activity Status page, in the Status Detail table, review the status of each of the steps. Click the status to view the log details.
	Ready for Deployment	N/A

Table 44-10 (Cont.) Diagnosing Patching Issues

Phase	State	Diagnosis and Resolution
Preparation	Preparation In Progress	N/A
	Preparation Error	Review the following log file: <gc_inst>/sysman/log/emoms.log
	Preparation Failed	On the Validation page, click Show Detailed Results here . On the Procedure Activity Status page, in the Status Detail table, review the status of each of the steps. Click the status to view the log details.
	Preparation Successful	N/A
Deploy	Deployment In Progress	N/A
	Deployment Error	Review the following log file: <gc_inst>/sysman/log/emoms.log
	Deployment Failed	On the Validation page, click Show Detailed Results here . On the Procedure Activity Status page, in the Status Detail table, review the status of each of the steps. Click the status to view the log details.
	Deployment Successful	N/A

Rolling Back Patches

If you want to roll back the patches, follow the roll back instructions outlined in the ReadMe that is packaged with the patch you applied.

Roll back functionality is supported on the following targets: Management Agent, Oracle WebLogic Server, Oracle SOA Infrastructure, Oracle Restart, single-instance Oracle database, Oracle RAC database, Oracle RAC One node, Oracle Grid Infrastructure, and Oracle Grid Infrastructure targets that are part of Oracle Exadata. For more information about the steps to rollback, see [Rolling Back Patches](#). For all other targets, Cloud Control does not support the automatic rollback of patches, and therefore, you must roll back the patches manually.

Additional Patching Tasks You Can Perform

This section covers the additional tasks you can perform with patch plans. In particular, it covers the following:

- [Viewing or Modifying a Patch Template](#)
- [Saving a Deployed Patch Plan as a Patch Template](#)
- [Downloading Patches from a Patch Template](#)
- [Deleting a Patch Template](#)
- [Deleting a Patch Plan](#)
- [Converting a Nondeployable Patch Plan to a Deployable Patch Plan](#)
- [Associating Additional Targets to a Patch in a Patch Plan](#)
- [Manually Staging the Patching Root Component](#)

- [Restricting Root User Access for Patching](#)
- [Resolving Patch Conflicts](#)
- [Analyzing the Results of Patching Operations](#)
- [Customizing Patching Deployment Procedures](#)
- [Rolling Back Patches](#)

Viewing or Modifying a Patch Template

To view or modify a patch template, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, do one of the following:
 - Select a patch template. From the context menu, select **View**. The Edit Template Wizard appears.
 - Click the name of a patch template. The Edit Template Wizard appears.

Note:

- An administrator who created the patch template and the super administrator of Cloud Control can modify a patch template.
- You can modify only the description and the deployment date in the patch template.

Saving a Deployed Patch Plan as a Patch Template

If you have already analyzed and deploy a patch plan, and if you want to save that patch plan as a patch template, then use one of the following approaches:

Approach 1

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region select a successfully analyzed deployable Patch Plan.

Note:

You can create a patch template only from one Patch Plan at a time.

3. Select **Create Template**. The Create New Plan Template dialog appears.
4. Enter a unique name for the template, then click **Create Template**.



Note:

When you select a plan, the **Create Template** option is not visible if you:

- Select a nondeployable Patch Plan or a deployable Patch Plan that has either not been analyzed or resulted in errors after being analyzed.
- Do not have the privileges to view the Patch Plan that you selected.
- Do not have the privileges to create a template.

Approach 2

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the **Patches & Updates** page, in the **Plans** region click the name of a successfully analyzed deployable patch plan. The **Create Plan** wizard appears.
3. In the Create Plan Wizard, in the Review & Deploy page, click **Save as Template**.
4. Enter a unique name for the template, then click **Create Template**.



Note:

You must create patch templates only with unique names.

Downloading Patches from a Patch Template

To download a patch from a patch template, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, select a patch template.
3. From the context menu, select **View**. The Edit Template Wizard appears.
4. In the Edit Template Wizard, on the Patches page, click a patch number. The patch details page appears.
5. On the patch details page, click **Download**.

Deleting a Patch Plan

To delete a patch plan, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, click the Patch Plan you want to delete. From the context menu, click **Remove**.

 Note:

If you do not see the Plans region, click **Customize Page** from the top-right corner, and then drag the Plans region to the page.

Deleting a Patch Template

To delete a patch template, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region select one or more patch templates. The context menu appears.
3. From the context menu, select **Remove**.

 Note:

An administrator who created the patch template and the super administrator of Cloud Control can modify a patch template.

Converting a Nondeployable Patch Plan to a Deployable Patch Plan

To make a nondeployable Patch Plan deployable, divide the Patch Plan into smaller deployable plans that contain only homogenous patches and targets.

 Note:

For more information, see [Patch Plan Becomes Nondeployable and Fails](#).

Associating Additional Targets to a Patch in a Patch Plan

To associate additional targets to a patch in a patch plan, use one of the following approaches:

Approach 1

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Plans region, select a Patch Plan to which the patch belongs. From the context menu that appears, select **View**.



Note:

If you do not see the Plans region, click **Customize Page** from the top-right corner, and then drag the Plans region to the page.

3. In the Create Plan Wizard, on the Patches page, Click **Add Patch**. The Edit Search window appears.
4. In the Edit Search window, search the patch to which you want to associate additional targets.
5. Select the patch that you want to add, then click **Add to This Plan**. The Add Patch To Plan window appears.
6. In Add Patch To Plan window, search and select the additional targets that you want to associate with the patch, and then, click **Add Patch to Plan**.



Note:

Ensure that you select only homogeneous targets.

For Oracle WebLogic Server, using a single patch plan, you can patch only one domain. However, if it is a shared domain, then the admin servers and Managed Servers running on different domains, which are shared with the domain being patched, are automatically added into the same plan.

For Oracle WebLogic Server, if you have deployed the Enterprise Manager for Oracle Fusion Middleware 12.1.0.4 plug-in, you can add any number of WebLogic domain targets to a single patch plan. If you have not deployed the Enterprise Manager for Oracle Fusion Middleware 12.1.0.4 plug-in, you can add only a single WebLogic domain target to a single patch plan. For information about how to deploy a new plug-in or upgrade an existing plug-in, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

If the WebLogic domain target that you add to a patch plan is a shared domain, then all the Administration Servers and the Managed Servers that are running on the domains that are shared with the domain being patched are automatically added into the same patch plan.

Approach 2

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Patch Recommendations region, click the graph.

 Note:

If you do not see the Patch Recommendations region, click **Customize Page** from the top-right corner, and then drag the Patch Recommendations region to the page.

3. On the Patch Recommendations page, select a patch.
4. From the context menu, select **Add to Plan**, then **Add to Existing**, and select the plan you want to add the patch to.

The patch you selected and the target it is associated with get added to the existing plan.

Approach 3

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches & Updates**.
2. On the Patches & Updates page, in the Search region, search a patch you want to add to the patch plan.

 Note:

If you do not see the Search region, click **Customize Page** from the top-right corner, and then drag the Search region to the page.

3. On the Patch Search page, click the patch to view its details.
4. On the patch details page, click **Add to Plan**, then **Add to Existing**, and select the plan you want to add the patch to.

The patch you selected and the target it is associated with get added to the existing plan.

Manually Staging the Patching Root Component

For information on how to manually stage the patching root component, see [Root Components for Patching](#).

Restricting Root User Access for Patching

For information on how to restrict the root user access, see [Root Components for Patching](#).

Resolving Patch Conflicts

If the patches in the patch plan conflict among themselves or with patches in the Oracle home, then you can do one of the following to resolve the conflict:

- Request for a merge patch of the conflicting patches. To do so, click **Request Patch** on the Validation page.

- Roll back the conflicting patches in the Oracle home and forcefully apply the incoming patches. To do so, on the Deployment Options page, in the Advanced Patching Options section, from the **Conflict Resolution** list, select **Forcefully apply incoming patches**.
- Skip the conflicting patches. To do so, on the Deployment Options page, in the Advanced Patching Options section, from the **Conflict Resolution** list, select **Skip conflicting patches**.

Analyzing the Results of Patching Operations

If you want to analyze the results of the patching operations you have done over a period of time, then run the *EM Deployable Patch Plan Execution Summary*. The report shows the number of deployable and nondeployable plans you have had in the past, and provides a breakdown of the deployable plans indicating how many succeeded and failed, how many were analyzed and deployed, and so on.

To run the EM Deployable Patch Plan Execution Summary report, follow these steps:

1. From the **Enterprise** menu, select **Reports**, then select **Information Publisher Reports**.
2. On the Information Publisher Reports page, in the table, expand **Deployment and Configuration**, then expand **Patching Automation Reports**, and then select **EM Deployable Patch Plan Execution Summary**.

Customizing Patching Deployment Procedures

This section describes how you can customize patching deployment procedures. For information about customizing deployment procedures in general, see [Customizing Deployment Procedures](#). This chapter explains in detail what customization means, the different ways of customizing a deployment procedure, and so on.

By default, a patch plan uses a static, Oracle-supplied deployment procedure, or an OPlan based, dynamically generated deployment procedure to apply patches. You can customize both these kinds of deployment procedures. This section consists of the following:

- [Customizing a Static Patching Deployment Procedure](#)
- [Customizing a Dynamic Patching Deployment Procedure](#)

Customizing a Static Patching Deployment Procedure

To customize a default, static deployment procedure, and use it to patch your targets, follow these steps:

1. Access the Deployment Options page.
2. In the Customization section, click **Create Like and Edit**.
3. Edit the deployment procedure (you can even add manual steps to the deployment procedure), then save it with a unique, custom name.

For information on how to edit and save a deployment procedure, see [Customizing a Deployment Procedure](#).

4. To use the customized deployment procedure, return to the Customization section of the Deployment Options page, then select the customized deployment procedure.

If you want to edit a customized deployment procedure, select the customized deployment procedure in the Customization section, then click **Customize**.

Customizing a Dynamic Patching Deployment Procedure

To customize a dynamically generated deployment procedure, and use it to patch your targets, follow these steps:

1. Access the Deployment Options page.
2. In the Customization section, select **Specify custom steps to be added to generated patching deployment procedure**.

For each placeholder step in the deployment procedure, you can add three custom steps, a directive step (which enables you to run a directive stored in Software Library), a host command step (which enables you to run a command script on the patch targets), and a manual step (which enables you to provide manual instructions to a user).

3. Select all the custom steps that you want to add under the placeholder steps, then click **Enable**.

To disable an enabled custom step, select the custom step, then click **Disable**.

4. Select a custom step that you want to add, then click **Edit**.

If the custom step you selected is a directive step, follow the instructions outlined in [Adding Steps](#).

If the custom step you selected is a host command step, an Edit Custom Step dialog box appears, which enables you to specify the details of a command script that you want to run on the patch targets. For **Script**, specify the commands that you want to run on the patch targets. For **Interpreter**, specify the complete path of the interpreter that must be used to execute the specified command script. For **Credential Usage**, select the credentials that must be used to run the custom host command step. Click **OK**.

If the custom step you selected is a manual step, an Edit Custom Step dialog box appears, which enables you to provide manual instructions to the user during the execution of the deployment procedure. For **Instructions**, specify the instructions for the manual tasks that a user must perform in this custom step. Click **OK**. The specified instructions are displayed when this custom step of the deployment procedure is executed.

5. Repeat Step 4 for all the custom steps that you have enabled.

Pausing the Patching Process While Patching Targets in Rolling Mode

While patching your targets in rolling mode, you can choose to pause the execution of the patching deployment procedure after each node is patched. This feature is useful when you want to perform a clean up on the host having the patched node, verify whether the patch is applied successfully on the node, and so on, before starting the patching process on the next node.

 Note:

This feature is available only if you have the 12.1.0.5 Enterprise Manager for Oracle Database plug-in, or a higher version, deployed.

To use this feature, follow these steps:

1. Create a custom static deployment procedure as described in [Customizing a Static Patching Deployment Procedure](#). On the Create Like Procedure page, select the **Procedure Steps** tab, select the **Wait for user confirmation** option, then click **Save**.
2. Use the custom deployment procedure while creating a patch plan to patch your targets in rolling mode (by specifying this custom deployment procedure in the Customization tab of the Deployment Options page).

The patching deployment procedure is paused after a node is patched, and you can resume its execution from the Procedure Activity page.

Rolling Back Patches

To rollback patches, you must create a new patch plan, select the relevant patches to rollback, and then select the rollback check box. To do so, perform the following steps:

 Note:

The roll back functionality is supported on the following targets:

- Oracle Management Agents
- Oracle WebLogic Server
- Oracle Single Instance Databases
- Oracle SOA Infrastructure
- Oracle Restart
- Oracle RAC Database
- Oracle RAC One Node
- Oracle Grid Infrastructure
- Oracle Grid Infrastructure targets that are part of Oracle Exadata

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, and then click **Patches and Updates**.
2. On the Patches and Updates page, in the Patch Search region, enter the patch that you want to rollback.
3. In the **Add Patch to Plan** dialog box, enter a unique name for the plan, and select all the targets from where you want to rollback the patches.
4. Click **Create Plan**.

5. In the Create Plan Wizard, select **Deployment Options**.
6. On the Deployment Options page, in the How to Patch section, select **In Place**, and in the Rollback section, select **Rollback patches in the plan**. Click **Next**.
7. On the Validation page, click **Analyze** to validate the plan. After a successful analysis, click **Next**.
8. On the Review and Deploy page, review the details you have provided for the patch plan, and then click **Rollback** to rollback the selected patch from the corresponding targets mentioned in the plan.

End-to-End Use Case: Patching Your Data Center

For an end-to-end use case that demonstrates how Enterprise Manager can be used to roll out patches across a data center, see [End-to-End Use Case: Patching Your Data Center](#) .

Patching Database as a Service Pools

To patch Database as a Service pools, you must use the Pool Maintenance option. For detailed information on how to patch Database as a Service pools, see *Enterprise Manager Cloud Administration Guide*.

Also, to standardize software and have a homogeneous software configuration across your enterprise, and perform subscription based software maintenance, ensure that you use the Software Standardization Advisor feature.

Patching Linux Hosts

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

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

Overview of Patching Linux Hosts

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

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

The following are concepts related to Linux patching:

Type	Description
Linux Host	A host target in Cloud Control that is running the Linux operating system.
Linux Patching Group	A set of managed Linux hosts that are associated with a common list of RPM repositories. Every group is configured with an update schedule, according to which a recurring job is triggered, that will update the hosts of the group with the associated RPM repositories.

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

About the Deployment Procedure for Patching Linux Hosts

Cloud Control provides the following deployment procedures for Linux patching:

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

Supported Linux Releases

The following releases are supported for Linux patching:

- Oracle Linux 5
- Oracle Linux 6
- Oracle Linux 7
- Red Hat Enterprise Linux 5
- Red Hat Enterprise Linux 6

Setting Up Infrastructure for Linux Patching

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

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

- [Setting Up Linux Patching Group for Compliance Reporting](#)

Prerequisites for Using the Linux Patching Feature

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

1. Meet the basic prerequisites described in [Setting Up Your Infrastructure](#).
2. Install yum on all your Oracle Linux 6 target hosts. Install yum and up2date on all your Oracle Linux 5 target hosts.
3. Enable the following commands through SUDO:
 - /bin/cp
 - /bin/rm
 - /bin/chmod
 - /sbin/chkconfig
 - yum
 - up2date
 - sed
 - rpm

Setting Up the RPM Repository for Linux Patching

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

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

Note:

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

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

Prerequisites for Setting Up the RPM Repository

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

- Identify a Redhat or Oracle Linux host, install a Management Agent, and point to the OMS. This host must have the *sudo* package installed.
- Obtain a valid Customer Support Identifier (CSI) number from your Oracle sales representative.

After obtaining a valid CSI number, ensure that you create a ULN account. To create a ULN account, access the following URL:

<https://linux.oracle.com/register>

- Download the up2date packages (Oracle Linux 5 only) from the following URL:
<https://linux.oracle.com/switch.html>

Upload the downloaded packages to Software Library if the host on which you plan to set up the RPM repository is running on one of the following platforms:

- Red Hat Enterprise Linux 5 (i386)
- Red Hat Enterprise Linux 5 (x86_64)
- Red Hat Enterprise Linux 5 (ia64)

 Note:

You do not need to upload the up2date packages to Software Library if the host on which you plan to set up the RPM Repository is running on an Oracle Linux platform.

Follow these steps to upload up2date packages to the Software Library:

 Note:

For a multi-OMS setup, the following steps only need to be performed on one OMS.

1. Compress up2date and up2date-gnome into a zip file, and name it as up2date_comp.zip.
2. Copy the zip file to the <ORACLE_HOME>/sysman/metadata/swlib/patch/stageServerComponents directory present in the Oracle home of the OMS.
3. Edit the Patch Software Library entities metadata file swlib.xml present in the Oracle home of the OMS to upgrade the ExternalID of the Software Library entity **Up2date Package Component**.

To do so, follow these steps:

(1) Open the swlib.xml file present at the following location: \$ORACLE_HOME/sysman/metadata/swlib/patch/

(2) Search for the tag <Entity name="Install up2date RPM">, which in turn has a subtag ExternalID.

(3) Increase the values of the ExternalID by 0.1.

For example, if the original value of the entity in the software library's ExternalID is 2.0, then update the value by 0.1 to upgrade the ExternalID to 2.1.

4. Upload the zip file to Software Library by running the following command:

```
$ emctl register oms metadata -service swlib -file $ORACLE_HOME/
sysman/metadata/swlib -core
```

- Ensure that the `/var/www/html/` directory on the host on which you plan to set up the RPM repository has at least 60 GB of free disk space per channel.
- Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.

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

- Ensure that the `createrepo` package is installed on the RPM Repository host. To obtain this package, subscribe to the `el*_addon` or the `ol*_addon` channel.
- Ensure that the `yum-arch`, `uln-yum-proxy` (for Oracle Linux 5) or `uln-yum-mirror` (for Oracle Linux 6 and 7), and `yum-utils` packages are installed on the RPM Repository host. To obtain the `yum-arch` and the `uln-yum-proxy/uln-yum-mirrors` packages, subscribe to the `add ons` channel. To obtain the `yum-utils` package, subscribe to the `latest` channel.
- If the RPM Repository host is not running on Oracle Linux 6 (OL6), but is subscribed to an OL6 channel whose name is of the format `ol6_*`, then you must import the OL6 public key manually. To do so, follow these steps:

1. Download the OL 6 key from:

<http://public-yum.oracle.com/RPM-GPG-KEY-oracle-ol6>

2. Store it under the following directory on your host:

`/usr/share/rhn`

3. Run the following command:

```
rpm --import /usr/share/rhn/RPM-GPG-KEY-oracle-ol6
```

- Ensure that the Enterprise Manager user has the `EM_LINUX_PATCHING_ADMIN` role and the `FULL_LINUX_PATCHING_SETUP` privilege. If the Enterprise Manager user does not have these, ensure that the super user grants them.
- Ensure that the Oracle GPG keys are installed on the host on which you plan to set up the RPM Repository.

To install the Oracle GPG keys on a host running on the Oracle Linux 5 or Oracle Linux 6 platforms, run the following command:

```
rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY
```

Setting Up the RPM Repository for Patching

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

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

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the **Linux Patching Setup** tab, click **Setup RPM Repository**.
3. On the Setup RPM Repository page, in the RPM Repository Server section, select the RPM Repository server by clicking the search icon. Select the host assigned for subscribing to ULN.

4. In the Credentials section, ensure that the **Normal Host Credential** user has write access to the stage location, and the **Privileged Host Credential** user can sudo with root privilege. Click **Apply**.
5. In the Deployment Procedure submission confirmation, click **Linux RPM Repository Server Setup**. The deployment procedure starts a job to download latest RPM packages and Advisories from the subscribed ULN channels.
6. (Optional) If you want to change the refresh mode to 30 seconds, then from the **View Data** list, select **Real Time: 30 Second Refresh**.
7. In the Steps tab of the Status Detail section, check the status of this step. Wait till the step **Installing Up2date** is completed or skipped.
8. Click the status of the manual step **Register with ULN** to verify if your host has been registered to ULN.

If you have registered your host to ULN, then select the target and click **Confirm**, and then click **Done** to go to the main flow.

If you have not registered your host to ULN, then perform the following steps on your Linux host:

- a. Log in to the RPM Repository server machine.
- b. Check if your host can connect to ULN. If your host cannot connect to the ULN directly, you can configure **up2date** (for Oracle Linux 5) or **uln_register** (for Oracle Linux 6 or 7) to use a proxy server. To configure access to ULN using a proxy server, follow these instructions:

https://linux.oracle.com/uln_faq.html#9

- c. Register the host to ULN by following the steps at:

https://linux.oracle.com/uln_faq.html#2

Note:

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

9. Click the status of the step **Subscribe to ULN channels**.

When you register a Linux server to ULN, it will be subscribed to a channel that has the latest Oracle Linux packages for the appropriate architecture. If no additional channels are needed to be subscribed to your host, then select the target and click **Confirm**, and then click **Done** to go to the main flow.

If some additional channels are needed to be subscribed to your host, then perform the following steps:

- a. Log in to ULN:
<http://linux.oracle.com/>
- b. Click on the **Systems** tab to manage subscriptions for each subscribed server.
- c. Subscribe to all the additional channels you need.

 Note:

- If the createrepo package is not installed on your Linux host, subscribe to the el*_addon or the ol*_addon channel.
- Ensure that the yum-arch, uln-yum-proxy, (for Oracle Linux 5) or uln-yum-mirror (for Oracle Linux 6 or 7) and yum-utils packages are installed on your Linux host. To obtain the yum-arch and the uln-yum-proxy/uln-yum-mirror packages, subscribe to the add ons channel. To obtain the yum-utils package, subscribe to the latest channel.

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

Setting Up Linux Patching Group for Compliance Reporting

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

In particular, this section describes the following:

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

Prerequisites for Setting Up Linux Patching Group

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

- Set up RPM Repository server or set a custom RPM Repository as a channel in Cloud Control.
- Install yum on all your Oracle Linux 6 target hosts. Install yum and up2date on all your Oracle Linux 5 target hosts.
- Install Sudo on the target hosts.
- Ensure that the Enterprise Manager user logs in to the OMS with super user privileges.
- Ensure that the Enterprise Manager user has the EM_LINUX_PATCHING_ADMIN role and the FULL_LINUX_PATCHING_SETUP privilege. If the Enterprise Manager user does not have these, ensure that the super user grants them.

Setting Up a Linux Patching Group

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

To set up a Linux patching group, follow these steps:

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

In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum or up2date performs a GPG signature check on the packages obtained from the specified repositories. Sometimes, yum or up2date may require a public GPG key to verify the packages obtained from the repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.

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

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

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

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

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

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

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

**Note:**

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

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

[Table 45-1](#) describes the jobs that are submitted for setting up a Linux patching group.

Table 45-1 Jobs Submitted for Setting Up Linux Patching Group

Job	Description
Patching Configuration	This job configures all the hosts for patching. It creates configuration files to be used by the yum and up2date tools on each host. This job is executed just once on all the hosts contained in the Linux Patching group immediately.
Compliance Collection	Compares the versions of the packages already installed in each machine contained in the Linux Patching group with the package versions in the selected RPM Repositories, and generates Compliance Reports for indicating which packages are outdated. This job is executed once every 24 hours (after the group is set up) on all the hosts contained in the Linux Patching group.
Package Information	Collects the metadata information of each package contained in the selected RPM Repositories. This job is executed daily.
Packages Update	Updates non-compliant packages. This job will update the packages installed on the hosts in the group to ensure that they are up-to-date with respect to the package repositories for that group. This job will be submitted only when the option "Update Hosts" is selected in the step "Package Repositories" of the Linux Patching group wizard, and its schedule can be customized in the step "Schedule"

Patching Linux Hosts

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

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

**Note:**

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

Applying Patches on a Linux Patching Group Based on Compliance

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

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Linux Patching page, in the Compliance Report section, select the Linux patching group that you want to patch, then click **Schedule Patching**.
3. On the Package Repository page, in the LINUX Distribution section, select the tool that you want to use to update the RPMs. YellowDog Updater modified (yum), and up2date are two commonly used tools to patch Linux hosts.

**Note:**

If the Linux host to be patched is running on Oracle Linux 6 (OL6), then you must use the yum tool for patching. The up2date patching tool is not supported for this Linux version. If you do not use the yum tool in this scenario, the patching process fails on the *Configure Host For Patching* step with the following error:

You are not selecting 'yum' as the tool to update the RPMs in this system. 'yum' is the only supported tool for updating RPMs in Oracle Linux 6 operating system

(Only if you have selected yum as the patching tool) Ensure that you select the patching mode that you want to use. Select **Package update and new package installation** if you plan to update the existing packages, as well as install new packages. Select **Package update only** if you plan to only update the existing packages, and not install any new packages.

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

In the Package Repository section, select the RPM repositories that you want to use.

In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum or up2date performs a GPG signature check on the packages obtained from the specified repositories. Sometimes, yum or up2date may require a public GPG key to verify the packages obtained from the repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.

In the Advanced Options section, by default, the **Hide obsolete updates** option is selected. Selecting this option hides the obsolete packages on the Select Updates page. If you want to view these packages on the Select Updates page, ensure that you deselect this option.

(Only if you have selected yum as the patching tool) In the Advanced Options section, select one of the following patch application modes:

- **Most suitable architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, for the suitable RPM architectures that are installed on the Linux hosts that you are patching.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update packagename
```

- **Specific architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, on only those Linux hosts that have the RPM architecture of the selected package.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update packagename.arch
```

- **Specific version and architecture**, if you want yum to install only the specific version of the package selected on the Select Updates page, or update the existing version of the package to this specific version, on only those Linux hosts that have the RPM architecture of the selected package.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update epoch:packagename-ver-rel.arch
```

Click **Next**.

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



Note:

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

Click **Next**.

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

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

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

Click **Next**.

6. On the Credentials page, enter the credentials to be used for the updates.
Click **Next**.
7. On the Pre/Post script page, enter the scripts that need to be executed before/after the patching process, if any.
Click **Next**.
8. On the Schedule page, enter the details of the patching schedule that must be used.
Click **Next**.
9. On the Review page, review the update parameters.
Click **Finish**. A deployment procedure is submitted to update the selected packages. Follow all the steps of the procedure until it completes successfully.

Applying Ad Hoc or Emergency Patches on Linux Hosts

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

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Deployment Procedure Manager page, in the Procedure Library tab, select **Patch Linux Hosts**, then click **Launch**.
3. On the Package Repository page, in the LINUX Distribution section, select the tool that you want to use to update the RPMs. YellowDog Updater modified (yum), and up2date are two commonly used tools to patch Linux hosts.

Note:

If the Linux host to be patched is running on Oracle Linux 6 (OL6), then you must use the yum tool for patching. The up2date patching tool is not supported for this Linux version. If you do not use the yum tool in this scenario, the patching process fails on the *Configure Host For Patching* step with the following error:

You are not selecting 'yum' as the tool to update the RPMs in this system. 'yum' is the only supported tool for updating RPMs in Oracle Linux 6 operating system

(Only if you have selected yum as the patching tool) For the tool operation mode, ensure that you select **Package update and new package installation**. Since this method of patching Linux hosts without using a Linux patching group is meant for emergencies and is not based on a compliance report, you can only use it to install new packages, and not update existing packages.

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

In the Package Repository section, select the RPM repositories that you want to use.

In the Check GPG Signatures section, select **Check GPG signatures** to ensure that yum or up2date performs a GPG signature check on the packages obtained from the specified repositories. Sometimes, yum or up2date may require a public GPG key to verify the packages obtained from the repositories. This key may not be previously imported into the RPM database. To ensure that this key is imported, select **Import GPG key**, then specify the GPG Key URL.

In the Advanced Options section, by default, the **Hide obsolete updates** option is selected. Selecting this option hides the obsolete packages on the Select Updates page. If you want to view these packages on the Select Updates page, ensure that you deselect this option.

(Only if you have selected yum as the patching tool) In the Advanced Options section, select one of the following patch application modes:

- **Most suitable architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, for the suitable RPM architectures that are installed on the Linux hosts that you are patching.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update packagename
```

- **Specific architecture**, if you want yum to install the latest version of the selected package, or update the existing version of the package to the latest version, on only those Linux hosts that have the RPM architecture of the selected package.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update packagename.arch
```

- **Specific version and architecture**, if you want yum to install only the specific version of the package selected on the Select Updates page, or update the existing version of the package to this specific version, on only those Linux hosts that have the RPM architecture of the selected package.

If you select this option, Cloud Control runs the following yum command:

```
yum install|update epoch:packagename-ver-rel.arch
```

Click **Next**.

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



Note:

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

Click **Next**.

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

By default, every discovered Linux host is displayed on this page, and can be selected. However, if you want only those hosts that have an older version of at least

one of the packages (that you selected for the update operation in the previous step) to be displayed on this page, run the following command:

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

Click **Next**.

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

Click **Next**.

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

Click **Next**.

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

Click **Next**.

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

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

Managing Linux Configuration Files

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

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

Overview of Linux Configuration Files

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

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

HTTP server configuration file to this channel, then deploy the file from this channel to the target hosts.

Prerequisites for Managing Configuration Files

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

Creating a Linux Configuration File Channel

To create a configuration file channel, follow these steps:

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

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

Uploading Linux Configuration Files to a Particular Channel

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

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

Prerequisites for Uploading Linux Configuration Files

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

Uploading Linux Configuration Files

To upload configuration files, follow these steps:

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

6. For uploading from remote machine, click **Upload from Agent Machine**. Click **Select Target** and select the remote machine.

Before browsing for the files on this machine, set preferred credential for this machine.

7. After selecting the files, click **OK**.

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

Importing Linux Configuration Files from One Channel to Another

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

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

Prerequisites for Importing Linux Configuration Files

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

Importing Linux Configuration Files

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

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

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

Deploying Linux Configuration Files From a Particular Channel

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

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

Prerequisites for Deploying Linux Configuration Files

Before deploying configuration files, meet the following prerequisites:

- Ensure that the privileged patching user has write permission on the target machine location where each configuration file will be staged, and has SUDO privileges too.

- Ensure that there is at least one channel with some files uploaded.

Deploying Linux Configuration Files

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

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

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

Deleting a Linux Configuration File Channel

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

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

Prerequisites for Deleting a Linux Configuration File Channel

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

Deleting Linux Configuration File Channels

To delete a configuration file channel, follow these steps:

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

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

Oracle Grid Infrastructure and Oracle RAC Configuration Support

This section describes the configurations that OPlan supports for patching GI and RAC databases of versions 11.2.0.2 or higher, on Linux X64, Solaris X64, Solaris SPARC and AIX platforms. Enterprise Manager integrates with OPlan to generate the procedure dynamically. If you use OPlan, then the commands that run as *root* will use the

script available in the target Oracle Home. The commands required to run as *root* depend on the version and the mode of patching.

The following table lists the details:

Table 45-2 Oracle Grid Infrastructure and Oracle RAC Configuration Support

Version	Mode	Command
11.2	In-Place	<pre><CRS_HOME>/crs/install/rootcrs.pl -unlock <CRS_HOME>/rdbms/install/rootadd_rdbms.sh <CRS_HOME>/crs/install/rootcrs.pl -patch</pre>
11.2	Out Of Place	<pre>mkdir -p <PARENT_FOLDER_OF_CLONE_CRG_HOME>; cp -pR <CRS_HOME> <CLONE_CRG_HOME> /usr/bin/perl <CLONE_CRG_HOME>/crs/install/ rootcrs.pl -unlock -destcrshome=<CLONE_CRG_HOME> <CLONE_CRG_HOME>/root.sh mkdir -p <PARENT_FOLDER_OF_CLONE_DB_HOME>; cp -pR <DB_HOME> <CLONE_DB_HOME> /usr/bin/ perl <CLONE_DB_HOME>/root.sh /usr/bin/perl <CLONE_CRG_HOME>/crs/install/ rootcrs.pl -patch -destcrshome=<CLONE_CRG_HOME> /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -patch -destcrshome=<CRS_HOME></pre>
12.1	In-Place	<pre><CRS_HOME>/bin/crsctl stop crs /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -prepatch (-nonrolling) <CRS_HOME>/rdbms/install/rootadd_rdbms.sh /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -postpatch (-nonrolling)</pre>

Table 45-2 (Cont.) Oracle Grid Infrastructure and Oracle RAC Configuration Support

Version	Mode	Command
12.1	Out of Place	<pre> mkdir -p <PARENT_FOLDER_OF_CLONE_CRS_HOME>; cp -pR <CRS_HOME> <CLONE_CRS_HOME> /usr/bin/perl <CLONE_CRS_HOME>/crs/install/ rootcrs.pl -prepatch -dstcrshome=<CLONE_CRS_HOME> <CLONE_CRS_HOME>/root.sh mkdir -p <PARENT_FOLDER_OF_CLONE_DB_HOME>; cp -pR <DB_HOME> <CLONE_DB_HOME> <CLONE_DB_HOME>/root.sh /usr/bin/perl <CLONE_CRS_HOME> /crs/install/rootcrs.pl -postpatch -dstcrshome=<CLONE_CRS_HOME> /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -prepatch -dstcrshome=<CRS_HOME> /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -postpatch -dstcrshome=<CRS_HOME> /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -prepatch -rollback -dstcrshome=<CRS_HOME> /usr/bin/perl <CRS_HOME>/crs/install/ rootcrs.pl -postpatch -rollback -dstcrshome=<CRS_HOME> </pre>

Additional Linux Patching Tasks You Can Perform

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

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

- [Deleting a Package Channel](#)

Viewing Linux Patching Compliance History

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

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

Prerequisites for Viewing Linux Patching Compliance History

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

Viewing Linux Patching Compliance History

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

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

Note:

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

Patching Non-Compliant Linux Packages

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

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

Prerequisites for Patching Non-Compliant Linux Packages

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

Patching Non-Compliant Linux Packages

To patch non-compliant packages, follow these steps:

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

Rolling Back Linux Patch Update Sessions or Deinstalling Packages

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

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

Note:

- Rolling back upgrades is supported to a certain extent. When performing an upgrade such as from OEL 5.2 to OEL 5.3, many RPMs that are dependent on others are upgraded. When you apply RPMs, this dependency can be followed. However, when rolling back patch update sessions, this dependency must be followed in reverse order. This reverse operation is not supported by yum or up2date. Hence, you can use the roll-back feature to rollback a patch update session, but not to completely rollback a major upgrade such as from OEL 5.2 to OEL 5.3.
- Rolling back upgrades is not supported on hosts running on Oracle Linux 6.

Prerequisites for Rolling Back Linux Patch Update Sessions or Deinstalling Packages

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

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

Rolling Back Linux Patch Update Sessions or Deinstalling Packages

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

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

Registering a Custom Package Channel

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

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

Prerequisites for Registering a Custom Package Channel

Before registering a custom channel, meet the following prerequisites:

- Ensure that the RPM Repository is under `/var/www/html` and is accessible through HTTP protocol.
- Ensure that Apache is installed, and listening on port 80. To verify this, you can try connecting to the URL: `http://host`.
For example: `http://h1.example.com`. If this works, then it is confirmed that Apache is installed and listening on port 80.
- Ensure that metadata files are created by running `yum-arch` and `createrepo` commands.
- Ensure that a Management Agent is installed on the RPM repository host, and ensure that Management Agent is communicating with the OMS.
- Ensure that the Enterprise Manager User logs in with Super User privileges for registering a custom channel.

Registering a Custom Package Channel

To register a custom RPM Repository, follow these steps:

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

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

Cloning a Package Channel

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

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

Prerequisites for Cloning a Package Channel

Before cloning a channel, meet the following prerequisites:

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

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

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

Cloning a Package Channel

To clone a channel, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the source channel you want to clone, and click **Create Like**.

4. Enter the credentials to use for the source channel. The credentials must have both read and write access.
5. Enter a unique target channel name.
6. Click **Browse** to select the target host name.
7. Enter the directory location of the target channel. This directory should be under `/var/www/html`.
8. Enter the credentials to use for the target channel. This credential should have both read and write access.
9. Click **OK**.

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

Copying Packages from One Channel to Another

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

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

Prerequisites for Copying Packages from One Channel to Another

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

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

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

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

Copying Packages from One Channel to Another

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

1. In Cloud Control, from the **Setup** menu, select **Provisioning and Patching**, then select **Linux Patching**.
2. On the Patching Setup page, in the Linux Patching Setup tab, click **Manage RPM Repository**.
3. On the Manage RPM Repository page, select the source channel, and click **Copy Packages**.
4. Select the target channel.

5. From the source channel section, select and copy the packages to the target channel section.
6. Enter credentials for the source and target channels. These credentials should have read/write access to the machines.
7. Click **OK**.

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

Adding Custom Packages to a Channel

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

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

Prerequisites for Adding Custom Packages to a Channel

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

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

Adding Custom Packages to a Channel

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

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

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

Deleting a Package Channel

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

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

Prerequisites for Deleting a Package Channel

Before deleting a channel, meet the following prerequisites:

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

Deleting a Package Channel

To delete a channel, follow these steps:

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

Oracle Linux Home

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

Oracle Linux Home includes:

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

This chapter covers the following:

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

Oracle Linux Home Target

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

1. From Enterprise Manager Cloud Control home page, click **Enterprise** select **Cloud**, and then click **Oracle Linux Home**.

The **Oracle Linux Home** page appears.

2. A new system repository target "Oracle Linux Home" is created which is associated with all the Oracle Linux Hosts (Physical or VM's). Clicking on that will redirect to the "Oracle Linux Home" page.

Oracle Linux Home page


Oracle Linux Home Page has the following sections:

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

General Region

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

Figure 46-1 General Region




The screenshot shows a dashboard titled "General" with a settings gear icon. Below the title, it displays "Total Linux Hosts 7" with a green up arrow and "6" (representing up hosts) and a yellow gear icon and "1" (representing down hosts). Below this is a table with the following data:

OS Version	Total	Up	Down
Oracle Linux 6.7	1	-	1
Oracle Linux 6.9	1	1	-
Oracle Linux 7.3	4	4	-
Oracle Linux 7.4	1	1	-

Overview of Incidents and Problems

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

Figure 46-2 Overview of Incidents and Problems



The screenshot shows a dashboard titled "Overview of Incidents and Problems" with a settings gear icon. It is divided into two main sections: "Incidents" and "Problems".

Incidents

Updated in the last 24 hours: 1

Updated in last 7 days: 1

Breakdown of incidents updated in the last 7 days

Category	Severity	Incident	Warning	Problem
Availability	-	-	1	-
Performance	-	-	-	-
Security	-	-	-	-
Others	-	-	-	-


Problems

Total Open: 0

Host Flux

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


Figure 46-3 Host Flux



CPU


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

Figure 46-4 CPU Graph



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


Figure 46-5 CPU Utilization Table



Memory


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

Figure 46-6 Memory Graph



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

Figure 46-7 Memory Utilization Table



Oracle Linux Patching


The Oracle Linux Patching screen contains:

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


Patchable Linux Groups

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

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




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



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

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



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


Total Hosts : 1

100.0%
■ Non-Compliant (1)


Patchable Linux Hosts

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


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



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



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



Ksplice for Oracle Linux

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



Note:

The minimum required ksplice version for Enterprise Manager is 1.0.29.

This section covers the following topics:

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

Ksplice Metrics



Note:

To access the Ksplice Metrics, from the **Host** menu on a host's home page, select **Configuration**, and then click **Latest**. The Oracle Linux Host needs the latest Ksplice software to be installed and configured with Oracle Premium Support before these metrics can report data. For more information, see [Oracle Ksplice User's Guide](#).

The following attributes are covered under Ksplice:

1. Ksplice
 - Ksplice Version
 - This attribute reports the version of the Ksplice software installed on the target host.
 - Ksplice Status
 - This reports if the host is configured to get updates from the Ksplice Server or if Ksplice offline.
 - Base Kernel Version
 - This attribute queries the stock Kernel running in the system.
 - This version does not represent the patched version, it shows only the one that booted the system.

- Effective Kernel Version
 - This attribute reports the Effective Kernel which means the Kernel version after the live patching including security fixes and others.
 - It also reports the date when the last applied patch was built.
 - Kernel Status
 - This attribute reports if the kernel of the host is up-to-date or out-of-date. A system is updated if it has all the available Ksplice patches installed.
 - In case Ksplice is offline, it is based on the `uptrack-updates-`uname -r`` package installed on the system.
 - Kernel Patches Installed
 - This reports the count of Ksplice packages installed on the system.
 - User Space Status
 - This attribute reports if the host's User Space Ksplice aware packages are up to date or if it is out of date.
 - In case Ksplice is offline, it is based on the local repository configured on the system.
 - User Space Packages Installed
 - This attribute reports the count of Ksplice User Space packages installed on the system.
2. Kernel Installed Patches
 - This attribute reports the installed Ksplice patches in the system.
 3. Kernel Available Patches
 - This attribute lists the available Ksplice patches for the system, in short it list the patches that have not yet been installed.
 - This information is derived based on the Ksplice configuration.
 - In case a Ksplice host is configured with a Ksplice server, it gets that information from the server.
 - In case the Ksplice is offline, it reflects the data based on the `uptrack-updates-`uname -r`` package installed on the system.
 4. User Space Installed Packages
 - This attribute reports the Ksplice User Space packages installed on the system.

Ksplice Patching

Ksplice Patching region on the Oracle Linux Home Page uses the metrics collected in Ksplice metrics to collate the Ksplice status over all the monitored Oracle Linux Hosts. For information on metrics, see [Ksplice Metrics](#). It contains the following two sub-regions:

- Ksplice Status Region

This region displays the total number of

- Ksplice Patching Configured Linux Hosts: If you click this link, it opens a pop-up window that lists the Oracle Linux Hosts which are configured with the Ksplice software.
- Ksplice Patching Non-Configured Linux Hosts: If you click this link, it opens a pop-up window that lists the Oracle Linux Hosts which are not configured with Ksplice software.

The Ksplice Status region also contains the following two pie charts:

- Kernel Status
- User Space Status


Each pie chart shows the status of all the hosts that is, how many hosts are compliant, non compliant, or non-compliant. If you click on a particular host, it opens another page with appropriate hosts list.

- Ksplice Summary Region

This region displays the table of hosts that lists the following:

- Ksplice Client (Online/Offline)
- Kernel Status (Compliant/Non-Compliant/Compliance unknown)
- User Space Status (Compliant/Non-Compliant/Compliance unknown)
- Effective Kernel Version

Figure 46-8 Ksplice Patching Status



Ksplice Linux Hosts Page

Ksplice Linux Hosts page contains the following tables:

- Ksplice Enabled Hosts with Ksplice Software
- Ksplice software Version
- Ksplice Client (Online/Offline)
- Kernel Status (Compliant/Non-Compliant/Compliance unknown-in case of unconfigured/offline systems)
- Number of Kernel Installed Patches
- User Space Status (Compliant/Non-Compliant/Compliance unknown-in case of unconfigured/offline systems)
- Number of User Space Installed Patches
- Base Kernel Version
- Effective Kernel Version

Host Ksplice Patches/Packages details

If you click the host name in the Ksplice detail table, a new page opens. This page lists the installed Ksplice patches on that host. If this host is a Ksplice Online host, it also lists the updates that are available on that host. These updates can be added or removed from this page.

If the host is a Ksplice Offline page, this page shows all the Ksplice Kernel patches available in the local repository or user space installed packages if there is any available in the local offline repository.

If Ksplice Enhanced Client Software is installed on the host, then it displays the list of installed/available user space patches. Otherwise, it shows the message “Install/Upgrade/Configure Ksplice Enhanced Client Software”.

The **Refresh** button in the screen helps in pulling the latest data to the dashboard. As this dashboard works in asynchronous mode, so in case there are some database changes or patches updated, you should use this **Refresh** button to sync the system. When you click this button, a dialog box appears asking for a confirmation.

For Install Updates, and Remove Updates you have to select/enter host root/sudo privileged credentials.

Note:


The Kernel Space Compliance status is based on in-memory kernel patches. It shows “Complaint”, if all the patches are installed, otherwise it shows “Non-Complaint”. When the patches are removed from the Ksplice Online System, it renders the system non-compliant. User Space Compliance Status is based on on-disk user space packages. When user installs the user space packages, in-memory process is updated on the on-disk packages as well and hence, the status shows “Compliant”. When user removes user space packages, it only removes the in-memory updates. When user space updates are on disk but not in memory, it continues to show “Compliant”. Though it does not remove the on-disk packages, the table still shows the user space updates to be “Compliant”.

Oracle Linux Home

The Oracle Linux Main menu contains the following options:

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


Figure 46-9 Oracle Linux Home



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

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

Figure 46-10 Target Navigation Tree of Oracle Linux Home



Database Fleet Maintenance

This chapter describes the procedure to standardize database environments by automatically patching and upgrading a large number of databases with minimal downtime.

It covers the following sections:

- [About Database Fleet Maintenance](#)
- [Getting Started with Fleet Maintenance](#)
- [Update De-mystified](#)
- [Upgrading Oracle Homes](#)
- [Cleaning Up Oracle Homes](#)
- [Database Fleet Maintenance - RESTful APIs](#)

About Database Fleet Maintenance

Database Fleet Maintenance allows administrators to maintain groups or pools of Oracle Homes and associated databases by applying database updates that include:

- Interim one-off patches including quarterly security patch updates (SPUs/CPUs), Patchset Updates (PSUs), Release Updates (RUs) and Release Update Revisions (RURs).

Any new database patches and updates are made available in the form of images. An image (also referred to as gold image or software end-state) represents the software binary that is patched to the required level. Each newly created image from a new patch for a specific database version is a new version. For example, an image for Oracle Database version 18.3 contains gold images (such as Ver. 1, Ver. 2, and so on) at different patch levels.

Database fleet maintenance can be used for:

- Single Instance Oracle Homes and associated databases
- RAC Oracle Homes and associated databases
- Grid Infrastructure Homes and associated Grid Instances
- Oracle Restart Homes (Grid Infrastructure for Standalone Server) and associated databases
- Standby databases (Single Instance and RAC)
- RAC One Node

Fleet maintenance involves the following steps:

- Identifying configuration pollution using the Software Standardization Advisor.
- Creating reference environments. Reference environment is a pre-requisite for Gold Image creation. During each patching cycle, reference environment is patched with desired set of patches to be rolled to the target DBs.

- Creating gold images using reference environments and versioning them for additional changes.
- Subscribing databases and clusters to a gold image on which the appropriate patches have been applied.
- Deploying the image and switch associated targets from the old Oracle Home to the new Oracle Home.

This maintenance activity can be performed as and when a new version of the gold image is available.


 Note:

- Oracle provides latest updates and best practices about Database patching and upgrade with Fleet Maintenance on the My Oracle Support Web site at <https://support.oracle.com>, My Oracle Support Note ID 2435251.1.
- From the 12.2 release, classic patching using Lifecycle Management Graphical User Interface is deprecated. You must use the Database Fleet Maintenance feature to patch database and GI environment; where version ≥ 12.2 .

Getting Started with Fleet Maintenance

The following figure lists the steps required to perform database fleet maintenance. It covers various configurations such as RAC and Single Instance databases with or without Data Guard, Oracle Restart, RAC One-Node, and Grid Infrastructure Homes.

Figure 47-1 Database Fleet Maintenance Steps



Meeting the Prerequisites

Before you start the fleet management operations, you must ensure that the following prerequisites are met:

- The Enterprise Manager Software Library has been setup and configured.
- The database and cluster targets to be patched have been discovered in Enterprise Manager.
- A reference environment (Database and Oracle Home) representing the target state has been discovered in Enterprise Manager. This reference environment is required to create the gold image.
- Space is available on the hosts where the patching is to be carried out. More details are available in [Staging and Deploying the Software](#).

The following shows table that target types that are supported and the values you need to provide for the target_type and purpose fields in emcli commands:

Entity	Target_type	Purpose (for deploy operation)	Purpose (for Update operation)	Purpose (for Roll- back operation)	Purpose (for Cleanup software operation)
Single Instance database	oracle_data-base	DE- PLOY_DB_SOFT- WARE	UP- DATE_DB	ROLL- BACK_DB	CLEAN- UP_SOFT- WARE
RAC database	rac_data-base	DE- PLOY_RAC_SO FTWARE	UP- DATE_RAC DB	ROLL- BACK_RAC DB	CLEAN- UP_SOFT- WARE
Oracle Restart	has	DEPLOY_SI- HA_SOFTWARE	UP- DATE_SIHA	ROLL- BACK_SIHA	CLEAN- UP_SOFT- WARE
Grid Infrastruc- ture	cluster	DE- PLOY_GI_SOFT- WARE	UPDATE_GI	ROLL- BACK_GI	CLEAN- UP_SOFT- WARE
PDB	oracle_pdb	DE- PLOY_DB_SOFT- WARE/ DE- PLOY_RAC_SO FTWARE/+DE- PLOY_CDB/ ATTACH_CDB	UP- DATE_PDB	ROLL- BACK_PDB	-NA-
CDB/RAC CDB	oracle_data- base/ rac_data- base	DE- PLOY_DB_SOFT- WARE/ DE- PLOY_RAC_SO FTWARE/+DE- PLOY_CDB/ ATTACH_CDB	UP- DATE_DB/ UP- DATE_RAC DB	ROLL- BACK_DB/ ROLL- BACK_RAC DB	CLEAN- UP_SOFT- WARE

**Note:**

- All values in the table above are case sensitive.
- A PDB target may also be patched using an existing Container Database at a higher patch level, in this case ATTACH_CDB operation must be performed and deploy-software and deploy-CDB operations must be skipped.

Discovering Configuration Pollution

Patching is one of the most challenging phases in the product lifecycle. With multiple patch cycles and various database configurations, different patch plans are required to satisfy the patching needs of the users. To address this, administrators must first identify their specific database configuration needs and select the standard configurations that are to be enforced in their environment.

With the Enterprise Manager Oracle Database plug-in, administrators can use the Software Standardization Advisor that scans the database configuration across all targets and generates a report (on-demand) of the existing environment. The report identifies the number of unique configurations (Platform + Version + Database-Type + Patches) and lists the Oracle Homes and databases for each configuration.


For example, in a 12.1.0.2 Linux-64 2-node RAC environment, three distinct groups of patches may have been applied. Each of these groups may have multiple Oracle Homes and databases associated with them. To meet this requirement, the administrator can define a standard configuration that can help simplify the environment.

The Software Standardization Advisor also recommends a list of standardized configurations and lists all the Oracle Homes on which the configurations should be applied. Oracle recommends that administrators run this report quarterly to understand the state of their database configurations. This list becomes the basis of fleet maintenance.

To use the Software Standardization Advisor, follow these steps:

1. Login to the Enterprise Manager Console with administrative privileges.
2. From the **Targets** menu, select **Databases**.
3. From the **Administration** menu, select **Software Standardization Advisor**.
4. Choose the relevant tab, **Database** or **Grid Infrastructure** and click run/re-run to run the analysis.
5. Click **Generate Report** to generate a report as a spreadsheet.

A sample environment is shown below:



The current system configuration is as follows:

Product	Release	Platform	No of Oracle Homes	Current Patch Configuration	No of Databases	Database
OracleData-base12c	12.1.0.2.0	Linuxx86-64	2	19769480;20299023;415564.	1	SI: [salesCDB.localdomain]
OracleData-base12c	12.1.0.2.0	Linuxx86-64	3	19769480;20299023.	2	SI:[fm1.oracle.com;HRdb0000]
OracleData-base12c	12.1.0.2.0	Linuxx86-64	3	No patches applied	3	SI:[fm2.oracle.com;db01;FINdb000]

The recommended system configuration is as follows:


Product	Release	Platform	No of Oracle Homes	Current Patch Configuration	No of Databases	Database
OracleData-base12c	12.1.0.2.0	Linuxx86-64	8	19769480;20299023;415564.	6	SI: [HRdb0000;fm2.oracle.com;salesCDB.localdomain;fm1.oracle.com;FINdb000;db01]

Creating a Gold Image

The standard configuration described in the Discovering Configuration Pollution section is referred to as a *gold image*. It is the end of state software definition that contains information about the base software version plus the additional patches. The 'end-state definition' is the logical term; the physical software binary component that represents the end-state is called as a gold image. ("Gold" is the qualifier signifying the ideal standard or the level for the software configuration.)

The following figure shows how a gold image can be created.

Figure 47-2 Creating a Gold Image



You can create an image for a database or grid infrastructure home using the `emcli` command shown below. This image is stored in the Enterprise Manager Software Library and used for provisioning.

Prerequisites:

- The Enterprise Manager Software Library has been setup and configured.
- The database and cluster targets to be patched have been discovered in Enterprise Manager.
- Reference environment (Database and Oracle Home) representing the target state must have been discovered in Enterprise Manager. This reference environment is required to create the gold image.

Verb

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/user/input_rac"
```

where the `input_rac` file contains the variables explained below. For a complete description of input variables for `db_software_maintenance` family of verbs, see `db_software_maintenance` in the *Enterprise Manager Command Line Reference Guide*.

Description of the Input Variables

The input variables are described below:

- **IMAGE_NAME:** The name of the gold image. This name must be unique across images.
- **IMAGE_DESCRIPTION:** A description of the image.
- **REF_TARGET_NAME:** The Oracle home target that will be used to create this gold image. This is the Database or Grid Infrastructure Oracle Home from the existing environment on which the desired set of patches have been applied. For example, if we intend to patch our 18c Fleet of DBs with the latest Release Update (RU), we need to specify the Oracle Home target name for a reference DB Oracle

Home where 18c RU has already been applied. To find the reference target name, enter the following query on the Enterprise Manager repository:

```
SELECT distinct target_name FROM mgmt$target_properties WHERE tar-
get_name IN
SELECT target_name FROM mgmt_targets WHERE target_type='oracle_home'
AND host_name=<Host Name of this Oracle Home> AND property_name='IN-
STALL_LOCATION' AND property_value=<path of Oracle Home>
```

- **IMAGE_SWLIB_LOC**: Path in software library where the payload of the gold image will be stored.
- **REF_HOST_CREDENTIALS**: This is applicable for Database Gold Image only. For Grid Infrastructure or Oracle Restart Homes, this parameter should be listed as "REF_GI_CREDENTIALS". This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential> is the named credentials for the host on which the reference Oracle Home is located. This user must be the owner of Oracle home.
 - <Credential Owner> is the Enterprise Manager user who owns this Named Credential.

Note:

The **REF_HOST_CREDENTIALS** parameter is applicable for database gold images only. For Grid Infrastructure or Oracle Restart Homes, this parameter is listed as **REF_GI_CREDENTIALS**.

- **WORKING_DIRECTORY**: The temporary location on the host of reference Oracle home target.
- **STORAGE_TYPE_FOR_SWLIB**: The Software Library storage type. This can be OMS Shared or OMS Agent File system.
- **STORAGE_NAME_FOR_SWLIB**: The storage name for the Software Library. To retrieve this storage name, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**. In the **Software Library** page, from the **Actions** menu, select **Administration**. On this page, check the value in the Name column to retrieve the storage name for the Software Library.
- **VERSION_NAME**: An image can have multiple versions. A default version of the image is created and more versions can be added.

Sample Output

The following output is displayed:

```
'Create Gold Image Profile deployment procedure has been submitted successfully with
the instance name: 'CreateGoldImageProfile_SYSMAN_12_08_2015_05_09_AM' and execu-
tion_guid='25F92D9A00164A45E053D903C40A9B4B'
```

The operation is performed using a deployment procedure, and user needs to wait for its completion before performing next steps. Status of the gold image creation can be checked using the following command:

```
emcli get_instance_status -exec=25F92D9A00164A45E053D903C40'
```

Retrieving a List of Available Gold Images



Note:

To format the output of the emcli commands listed below, set the emctl property - oracle.sysman.dbprov.gis.emcli.verbs.tableLength to the width of the terminal. For example, emctl set property -name oracle.sysman.dbprov.gis.emcli.verbs.tableLength -value 160

Verb

When the image has been successfully created, you can get a list of images that are available with the following command:

```
emcli db_software_maintenance -getImages [-columnName=<column name(s)>]
```

Example

```
Image Id Image Name Description
Version
72D9D7C656A11AB2E053BC3A41 RACS_12.1_RAC_12.1.0.2.180
12.1.0.2.0 12CRACS 0A7A43
```

Description of the Input Variables

The input variables are:

- **columnName:** It is an optional comma-separated list parameter that can be used to limit the number of columns in the output. Supported values are image_id, image_name, image_version and image_description..

Sample Output

```
Image Id Image Name Description
Version Platform Name Creation Date Owner Modified By
FE55AD7AB28974EFE04313B2F00AD4A0 11.2.0.4.0 RAC DB Gold Image for 112040 RAC DB Homes
11.2.0.4.0 Linux x86-64 2015-12-28 13:22:42.0 TESTSUPERADMIN TESTSUPERADMIN
Verb getImages completed successfully
```

Listing the Versions of an Image

After you have retrieved a list of the available images, you can view a list of versions available for a specific image with the following command:

```
emcli db_software_maintenance -getVersions -image_id=<image_id>
[-columnName=<column names>]
```

Description of the Input Variables

The input variables are:

- **columnName:** It is an optional comma-separated list parameter that can be used to limit the number of columns in the output. Supported values are *image_id*, *image_name*, *image_version*, and *image_description*.

Sample Output

```
*****
*****
POSITION                                VERSION ID                                VERSION NAME
STATUS                                DATE CREATED
EXTERNAL ID                                HASHCODE
*****
1                                277DF28F2D30393BE053D903C40AC6    CUSTOMER 112044
VERSION ACTIVE    2015-12-22 05:54:45.0
ORACLE:DEFAULTSERVICE:EM:PROVI    C3448035451:B1400395227
SIONING:1:COMP:COMP_COMPONENT:S    10
UB_ORACLEDB:277DF28F2D2C393BE0
53D903C40AC610:0.1
*****
*****
2                                277E39B74D684C6BE053D903C40A59    CUSTOMER 112047
VERSION CURRENT    2015-12-22 06:14:40.0
ORACLE:DEFAULTSERVICE:EM:PROVI    C3448035451:B3881579444
IONING:1:COMP:COMP_COMPONENT:S    EB
UB_ORACLEDB:277E39B74D664C6BE0
53D903C40A59EB:0.1
*****
*****
TOTAL ROWS:2
```

Describing an Image

To view the content (patches/bugs) of an image, run the following emcli command. It shows the current version of the given image by default unless a specific *version_id* is provided.

Verb

```
emcli db_software_maintenance -describeImage -image_id=<Image Id> [-version_id=<version id>] [-bugs=true]
```

Sample Output

# PATCH#	PATCH DESCRIPTION
1 23054319	OCW Patch Set Update : 11.2.0.4.160719 (23054319)
2 19769489	Database Patch Set Update : 11.2.0.4.5 (19769489)
3 21352635	Database Patch Set Update : 11.2.0.4.8 (21352635)

Verifying if Image is Applicable

This step verifies if the image can be used to patch a specified database target. This is done by comparing the bug fixes available in the current Oracle home of the database target and the image.

Verb

Syntax: `emcli db_software_maintenance -checkApplicability -image_id=<image_id> -target_list=<target_list> -target_type=<target_type>`

Description of Input Variables

- `target_list`: List of database or grid infrastructure home or Oracle restart Home targets.
- `target_type`: Refer to the table “EMCLI command inputs based on Entity type”.

Examples:

RAC Database

```
emcli db_software_maintenance -checkApplicability -image_id=
FE55AD7AB28974EFE04313B2F00AD4A0 -target_list=RACDB1 -target_type=rac_da-
tabase
```

Oracle Restart

```
emcli db_software_maintenance -checkApplicability -im-
age_id=2EA08B6339234B7E053BC3A410A2826 -target_list=HAS1 -target_type=has
```

Grid Infrastructure

```
emcli db_software_maintenance -checkApplicability -im-
age_id=72F8E4FCDE017B99E053BC3A410AF226 -target_list=CLUSTER1 -tar-
get_type=cluster
```

Single Instance

```
emcli db_software_maintenance -checkApplicability -im-
age_id=6E4EA63DDC00535CE0531636B10ABA1A -target_list=DB1 -target_type=ora-
cle_database
```

Sample Output

```
Checking applicability for target [DB1]...
Image is applicable
```

This command can show one of the following results:

- **Applicable:** The image and database target contain the same set of bug fixes. The image can be applied on the specified target.
- **Applicable and Image has more bug fixes:** The image contains more bug fixes than those applied on the database. The list of extra bugs is displayed. The image can be applied on the specified target.
- **Not Applicable:** The database contains more bug fixes than those included in the image. The list of missing bugs is displayed. The administrator has to create a new version of the image that includes the missing bugs before the database can up- take the same.

Deleting an Image or Versions of an Image

Verb

To delete a version of an existing image, use the following command:


```
emcli db_software_maintenance -deleteVersion -version_id=<version_id> [-reportOnly=true|false]
```

Description of the Input Variables

The input variables are:

- Version_id: the ID of the version that needs to be deleted.
- reportOnly: This is an optional parameter and the default value is set to false. When it is specified as true, it lists the version details that will be deleted but the operation is not carried as the name suggests. This is used as a precautionary step before actual deletion of a version.

Verb

To delete an image, which will in turn delete all versions of an image and all metadata related to an image such as subscriptions and software library components, use the following command:

```
emcli db_software_maintenance -deleteImage -image_id=<image_id> [-reportOnly=true|false]
```

Description of the Input Variables


The input variables are:

- **Image_id:** the ID of the version that needs to be deleted.
- **reportOnly:** This is an optional parameter and the default value is set to false. When it is specified as true, it lists the version details that will be deleted but the operation is not carried as the name suggests. This is used as a precautionary step before actual deletion of a version.

Subscribing the Targets to the Selected Image

To achieve standardization, Oracle Homes must be subscribed to the image. A group of Grid Infrastructure Homes can subscribe to a image created with a reference Grid Infrastructure Home. A target should be subscribed to only one gold image at a given time.

Figure 47-3 Subscribing to an Image

**Verb**

Syntax: `emcli db_software_maintenance -subscribeTarget -target_name=<target name> -target_type=<target type> -image_id=<image id>`

where:

- target_name: The name of the target that needs to be patched.
- target_type: Refer to the table “EMCLI command inputs based on Entity type”
- image_id: The ID of the gold image to which this target is to be patched. You can get the image_id by running the emcli command. For information about the command, see [Retrieving a List of Available Gold Images](#).



Note:

If the target is a primary database, the standby database target is automatically subscribed to the same image.

Sample Output

```
Target 'RACDB1' subscribed successfully.
Verb subscribeTarget completed successfully
```




Note:

If the target is a primary database, standby database target is automatically subscribed to the same image.

Verifying the Subscription

You can view a list of all targets that have been subscribed to a specific image. In a large data center, there may be several images and multiple targets subscribing to each of these images. For example, the following figure shows three images and several targets subscribed to each of these images:

Figure 47-4 Verifying the Subscription



Verb

```
emcli db_software_maintenance -getImageSubscriptions -image_id=<image_id>
```

Sample Output

```
*****
*
```

```
TARGETNAME  TARGETTYPE  VERSIONNAME  DATESUBSCRIBED
CLUSTER1    CLUSTER      112044 Version  2015-12-22 06:59:08.0
```

```
*****
*
```

Total Rows:1

To verify the subscription at the target level, use the following command:

```
emcli db_software_maintenance -getTargetSubscriptions -target_name=<target_name> -target_type=<target_type>
```

Description of the Input Variables

- target_name: name of the target.
- target_type: refer to table “EMCLI command inputs based on Entity type”
- version_name: This is observed to be empty sometimes. A value is displayed when the existing Oracle Home of the target has same set of patches as any existing Gold Image version.

Sample Output

```
TARGET NAME TARGET TYPE IMAGE ID IMAGE NAME VERSION ID VERSION NAME PARENT NAME PARENT TYPE DATE SUBSCRIBED
=====
DB1 oracle_database 48B75BB7B5E0462 SIDB_11204 48B78E3967D56BE PSU Jan17 <null>
<null> Jan 2017-02-21
```

To verify the subscription of the container databases for PDB patching, use the following command:

```
emcli db_software_maintenance -getSubscriptionsForContainer -target_name=<CDB_TARGET_NAME> -target_type=<target_type> -image_id=<image_id>
```

Sample Output

```
emcli db_software_maintenance -getSubscriptionsForContainer
TARGET NAME TAREGT TYPE IMAGE ID
=====
conttarget oracle_database 73659339F2364E7CE053BC3A410A2016
```

Unsubscribing Targets from an Image

Verb

To unsubscribe a target from an image, use the following command:

```
emcli db_software_maintenance -unsubscribeTarget -target_name=<target_name> -target_type=<target_type> -image_id=<image_id>
```

where:

- target_name is the name of the RAC database target that needs to be patched.
- target_type – refer to the table “EMCLI command inputs based on Entity type”.

- `image_id` is the ID of the gold image to which this target is to be patched. You can get the `image_id` by running the `emcli` command. For information about the command, see [Retrieving a List of Available Gold Images](#).

Creating an Updated Version of the Image

After you have verified that all the targets have been subscribed to the selected image, a new version of the image must be created. This is the target version to which the databases are to be patched.

Verb

Syntax: `emcli db_software_maintenance -createSoftwareImage -input_file=data:<input_file_path>`

For a description of input file, see [Creating Gold Image](#).

Example: `emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/user/input_rac"`

For a RAC database, the `input_file` contents can be as follows:

```
IMAGE_ID=<ID of the existing gold image under which the version is to be created>
REF_TARGET_NAME=<Oracle home target name for 112044 Oracle home>
<IMAGE_SWLIB_LOC=Oracle Home Provisioning Profiles/11.2.0.4.4/linux_x64
REF_HOST_CREDENTIALS=REF_HOST_CREDS:SYSMAN
WORKING_DIRECTORY=/tmp
STORAGE_TYPE_FOR_SWLIB=0msShared
STORAGE_NAME_FOR_SWLIB=swlib
VERSION_NAME=PSU 112044 Version
```

Sample Output

```
Create Gold Image Profile deployment procedure has been submitted successfully with
the instance name: 'CreateGoldImageProfile_SYSMAN_12_08_2015_05_09_AM' and execu-
tion_guid='25F92D9A00164A45E053D903C40A9B4B'
```

Changing Version Status to Current

After an updated version of the image has been created (see [Creating a Gold Image](#), this new version must be marked as "Current". This indicates that this is the target version to which the databases are to be patched at the end of the fleet operation.

Verb:

`emcli db_software_maintenance -updateVersionStatus -version_id=<version_id> -status=CURRENT`

where `version_id` is the target version to which the databases are to be patched.

Sample Output

```
Version ID '269EA1638D1D7472E053D903C40ABF1B' updated successfully.
```

```
Verb updateVersionStatus completed successfully.
```

Staging and Deploying the Software

Fleet maintenance can be performed in a short period of time because the gold image (the version of the image that is marked as 'Current' to indicate the current gold standard) is copied to a parallel Oracle Home which is then turned into a newly patched Oracle Home. This command is used to create the parallel Oracle Home and copy the gold image there. In other words, this command allows you to stage and deploy the new Oracle Home on the subscribing Oracle Homes. The **Deploy** command automatically uses the CURRENT version of the subscribed image while creating the new Oracle Home.

Verb

```
emcli db_software_maintenance -performOperation -name=<User specified Operation name> -purpose=<Operation Type> -target_type=<type of target> -target_list=<list of targets> -normal_credential=<credential name> -privilege_credential<privileged credential name> -input_file="data:<location of input file>" [-standbyAutoDeploy=true|false] [-procedure_name_prefix="CustomUpdate"]
```

Description of the Input Variables

where:

- name: This is the unique name of the operation.
- purpose: Refer to the table EMCLI Command Inputs Based on Entity Type.
- target_type: The type of target on which this operation is being performed. Refer to the table EMCLI Command Inputs Based on Entity Type for valid values.
- target_list: This is a comma separated list of targets which need to be patched.
 - Targets of homogenous types are supported in a single fleet operation.
 - unique list of hosts based on this target list is displayed and start stage of Oracle home software on those hosts.
 - If targets running from same Oracle home are provided in this list, the stage and deploy operation will be triggered only once and not for all targets.
- normal_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- `start_schedule`: The date on which the stage and deploy is to be started if it is to be started in the future.
Format: "start_time:yyyy/mm/dd HH:mm"
Note: This is an optional parameter. If no date is provided, the fleet operation will start immediately.
- `standbyAutoDeploy`: This is an optional parameter with default value as True. If the target is a primary database, a new Oracle home with the same gold image version as Primary is deployed automatically on the Standby host. This parameter disables the automatic deployment of software on standby host when specified as false. If this value is false, standby staging/deploy can be performed independently using the above `emcli` command.
- `procedure_name_prefix`: This is an optional parameter. It is used to provide a custom name for the deployment procedure Instance name. The supported operations are DEPLOY, UPDATE, CLEANUP and ROLLBACK.
- `input_file` : This file will contain the following entries:
 - `NEW_ORACLE_HOME_LIST=<Absolute path to file system location for new Oracle Home>`
 - `workingDir=<Name of temp directory>`: This is the location in which the new Oracle Home will be installed on all the hosts. The Credential Owner must have read / write access to this location.
 - `dispatchLoc=<Dispatch location>`: The location in which the scripts will be staged on the host. These scripts can be executed by a **root** user.
 - `SKIP_PREREQUISITE_CHECKS=<true|false>`: The default value is false.
 - `SKIP_CVU_CHECK=<true|false>`: The default value is false.
 - `PREREQS_ONLY=<true|false>`: The default value is false. This can be used for identify errors during the prerequisite checks. The actual deployment will not happen when the value is set to "true".

Note: This verb submits a procedure which must be completed before you can proceed with the next steps.

Sample Output

```
Processing target "CLUSTER1"...
Checking if target is already running from the current version...
Check Passed.
Checking image subscription...
Check image subscription passed.
Operation 'Deploy 1120407 GI_Home' created successfully.
```

Space Requirements

- The temporary space in the working directory (`workingDir`) must be approximately twice the size of the Reference Home.

- The dispatch location (dispatchLoc) requires the same amount of space as the workingDir plus an additional 10 MB. If this parameter is not specified, the default working directory will be used.
- For the Oracle Home, space required is the same as Reference Home. It is recommended to have additional storage available for the growth of log files and so on.

 Note:

In case of primary-standby DataGuard setup, you must only provide primary target in `-target_list='PRIMARY_DB'` parameter.

Migrating the Listeners

If there are listeners running from the database home, use the following verb to migrate them to the parallel Oracle Home. You can use this command to migrate listeners running from Oracle database homes. Listeners running from Grid Infrastructure homes are automatically migrated during the GI update process.

Verb

```
emcli db_software_maintenance -performOperation -name="Update Listener" -
purpose=migrate_listener -target_type=oracle_database -target_list="DB1" -
normal_credential="NC_HOST_CREDS:SYSMAN" -privilege_credential="HOST
PRIV:SYSMAN" -start_schedule="start_time:yyyy/mm/dd HH:mm"
where:
```

- name: This is the unique name of the operation.
- purpose: MIGRATE_LISTENER
- target_type: The type of target on which this operation is being performed. This can be "rac_database" for RAC and "oracle_database" for single instance databases.
- target_list: This is a comma separated list of targets which are to be migrated.
- normal_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- start_schedule: The date on which the stage and deploy is to be started if it is to be started in the future.

Format: "start_time:yyyy/mm/dd HH:mm"

Note: This is an optional parameter. If no date is provided, the fleet operation will start immediately.

Sample Output

Processing target "CLUSTER1"...

Checking if target is already running from the current version...

Check Passed.

Checking image subscription...

Check image subscription passed.

You can monitor the patch operation status with the following command:

```
emcli get_instance_status -exec=25BE102DD7CB544EE053D903C40A784D -details -xml
```

Deploying Container Databases

Verb

In case of container databases, users have flexibility to update one or more PDBs at a time rather than the entire container being updated at the same time. This is made possible by creating a new container database or using an existing container on staged/patched Oracle Home and moving the PDBs to the new container on demand. A new CDB is created using the structure-only template. Deploy software must be performed before a DEPLOY_CDB operation.

```
emcli db_software_maintenance -performOperation -purpose="DEPLOY_CDB" -
target_list="<CDB Name that is being patched> -target_type="<target type>"
-name="Operation Name" -description="Operation description" -db_prefix|
db_name="<DB Name Prefix or DB Name>" -normal_credential="<credential
name>" -privilege_credential="<credential name>" -database_creden-
tial="<SYSDBA credential name>"
```

If a container database already exists with the required image, then the same can be attached for housing the PDBs when they are updated.

```
emcli db_software_maintenance -performOperation -purpose="ATTACH_CDB" -
target_list="<CDB Name that is being patched> -target_type="<target type>"
-name="Operation Name" -description="Operation description" -destina-
tionCDB="<Container database to which PDBs will be migrated>" -normal_cre-
dential="<credential name>" -privilege_credential="<credential name>" -da-
tabase_credential="<SYSDBA credential name>"
```

Updating the Database / Cluster

After the stage and deploy operations have been completed, the targets need to be migrated to the newly deployed Oracle homes.

Verb

```
emcli db_software_maintenance -performOperation -name=<User Specified Op-
eration> -purpose=<Operation Type> -target_type=<target type> -tar-
get_list=<list of targets> -normal_credential=<credential name> -privi-
lege_credential=<credential name> [-rolling=<true/false>] -input_file="da-
ta:<location of input file>"
```

where:

- name: This is the unique name of the operation.
- purpose: Refer to the EMCLI Command Inputs Based on Entity Type table.
- target_type: Refer to the EMCLI Command Inputs Based on Entity Type table.
- target_list: This is a comma separated list of targets which need to be patched.
 - Targets of homogenous types are supported in a single fleet operation.
- normal_credential: This must be entered in the format <Named Credential: Cre-
dential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home
will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Cre-
dential.
- privilege_credential: This must be entered in the format <Named Credential: Cre-
dential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home
will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Cre-
dential.

These credentials are used to run scripts as root.

- rolling: This is an optional flag with the default value as true. The update procedure works in "Rolling Patch" mode by default but you can override this if necessary. If the gold image has OJVM patches, you should mark this value as 'false' which will take all instances of RAC DB down for update. This option is applicable for DB patching only (not for Grid Infrastructure or Oracle Restart).
- input_file: This is an optional parameter. The parameter(s) that can be specified here are: workingDir: =<Name of temp directory>.

Sample Output

```
Processing target "CLUSTER1"...
```

```
Checking if target is already running from the current version...
```

```
Check Passed.
```

```
Checking image subscription...
```

```
Check image subscription passed.
```

You can monitor the patch operation status with the following command:

```
emcli get_instance_status -exec=25BE102DD7CB544EE053D903C40A784D -details -xml
```

Rolling Back Cluster / Database

This command is used to switch the database back to the previous Oracle home after the Update operation has been completed.

Verb

```
emcli db_software_maintenance -performOperation -name="Rollback RAC DB" -
purpose=ROLLBACK_RACDB -target_type=<target type> -target_list=<list of
targets> -normal_credential="NC_HOST_CREDS:SYSMAN" -privilege_creden-
tial="HOST_PRIV:SYSMAN" [-rolling=true/false] [-node_list="host1.exam-
ple.com"]
```

where:

- name: This is the unique name of the operation.
- purpose: There are standard purposes that can be performed by Fleet Operations which can be:
 - ROLLBACK_DB
 - ROLLBACK_RACDB
 - ROLLBACK_GI
- target_type: The type of target being provided in this operation which can "rac_database" or "oracle_database".
- target_list: This is a comma separated list of targets which need to be patched.
 - Targets of homogenous types are supported in a single fleet operation.
 - A unique list of hosts based on this target list is displayed and start stage of Oracle home software on those hosts.
 - If targets running from the same Oracle home are provided in this list, the stage and deploy operation will be started only once and not for all targets.
- normal_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- non_rolling: By default, rollback is performed on all nodes. If the non_rolling back flag is enabled, you can select the list of nodes (using the node_list command) that are to be rolled back

- **rolling:** By default rollback is performed in rolling fashion. This flag is used when the current Oracle home has patches that were applied in non-rolling mode (OJVM) and need to be rolled back.
- **node_list:** This is a comma separated list of hosts on which the instances need to be updated.

For example: If RACDB is running on a 4 node cluster host1, host2, host3, and host4 and you choose to update the instances in only 2 hosts at a time, the value of this parameter needs to be specified as `node_list="host1, host2"`.

Sample Output

```
Processing target "RACDB"...

Checking if target is already running from the current version of the image...

Check Passed.

Checking image subscription...

Check image subscription passed.

Operation Rollback RAC DB' created successfully.
```

Non-Rolling Applicable Patches

If the gold image contains a patch that is non-rolling applicable (E.g. OJVM patch pre-Jan 2017 PSUs), then every update or switch must be performed with the flag set to `"rolling=false"`. This starts the database in Upgrade mode before applying the data changes. Then the database is shut down and restarted in Normal mode.

Updating and Rolling Back Standby Databases

This section is applicable for all primary and physical standby configurations (RAC - RAC, RAC - SI, SI - SI, RAC One - RAC One). Both the primary and standby databases need to be switched to the new Oracle home in this use case. A single switch operation cannot handle both primary and standby databases.

The procedure to update and roll back standby databases varies depending on the patch_type (rolling applicable or not)

For rolling applicable patches like database PSUs, follow these steps:

- **Update the standby database**

```
emcli db_software_maintenance -performOperation -name=<User Specified Operation Name> -purpose=<Operation Type> -target_type=<Target Type> -target_list=<List of targets> -normal_credential=<credential name> -privilege_credential=<credential name> -dataguard_role=standby -input_file="data:<location of input file>"
```

Description of Input Parameters

- **input_file:** This is an optional parameter. The parameter(s) that can be specified here are:
 - * **workingDir=**<Name of temp directory>
 - * **disableDG=**<true|false>. The default value is false. This should be set to 'true' if the standby is managed by Data Guard

- Update the primary database

```
emcli db_software_maintenance -performOperation -name=<User Specified
Operation Name> -purpose=<Operation Type> -target_type=<Target Type> -
target_list=<List of targets> -normal_credential=<credential name> -
privilege_credential=<credential name> -dataguard_role=primary [-ig-
noreStandbyPrereq=true|false] -input_file="data:<location of input
file>"
```

Description of Input Parameters

- ignoreStandbyPrereq: The default value is false. This disables verification check if the Standby is on the same image version to which the Primary is being moved.
- input_file: This is an optional parameter. The parameter(s) that can be specified here are:
 - * workingDir=<Name of temp directory>
 - * enableDG=<true|false> The default value is false. This should be set to 'true' if the standby is managed by Data Guard.

For non-rolling applicable patches such as OJVM (pre-Jan 2017 PSUs), follow these steps:

- Update the standby database

```
emcli db_software_maintenance -performOperation -name=<User Specified
Operation Name> -purpose=<Operation Type> -target_type=<Target Type> -
target_list=<List of targets> -normal_credential=<credential name> -
privilege_credential=<credential name> -dataguard_role=standby -star-
tupAfterSwitch=false -input_file="data:<location of input file>"
```

Description of Input Parameters

- startupAfterSwitch: For non-rolling patches, it is necessary to bring up the Primary database first from the patched home. It will ensure that the Standby database is left in a shutdown state during the Update operation. The standby database will be started in a separate step after the Primary database is patched and started.
- input_file: This is an optional parameter. The parameter(s) that can be specified here are:
 - * workingDir=<Name of temp directory>
 - * disableDG=<true|false>. The default value is false. This should be set to 'true' if the standby is managed by Data Guard

- Update the primary database with dataguard_role=Primary

```
emcli db_software_maintenance -performOperation -name=<User Specified
Operation Name> -purpose=<Operation Type> -target_type=<Target Type> -
target_list=<List of targets> -normal_credential=<credential name> -
privilege_credential=<credential name> -dataguard_role=primary roll-
ing=false -input_file="data:<location of input file>"
```

Description of Input Parameters

- input_file: This is an optional parameter. The parameter(s) that can be specified here are:
 - * workingDir=<Name of temp directory>

- Start the standby database from a patched home

```
emcli db_software_maintenance -performOperation -name=<User Specified
Operation Name> -purpose=<Operation Type> -target_type=<Target Type> -
target_list=<List of targets> -normal_credential=credential name> -
privilege_credential=<credential name> -dataguard_role=standby -star-
tupDatabase=true -skipSwitchDatabase=true -input_file="data:<location
of input file>"
```

Description of Input Parameters

- startupDatabase: It should always be set to **true**. This enables the starting of the Standby database from the patched home.
- skipSwitchDatabase: It should be set to **true**. It skips the switching of database to patched home since it is already performed in earlier steps.
- input_file: This is an optional parameter. The parameter(s) that can be specified here are:
 - * workingDir=<Name of temp directory>
 - * enableDG=<true|false> The default value is false. This should be set to 'true' if the standby is managed by Data Guard.
 - * primary_dbhost_creds=<credential name of primary database host>. This is required because the Data Guard is enabled through the Primary Host. This can be the same as what is used in the Primary database update.

Note:

The rollback process is the same but the purpose value will be different. Refer to the EMCLI Command Inputs Based on Entity Type table for the correct value.

Node Wise RAC Database / Cluster Update

A cluster update always requires that all the RAC database instances running on that node be shut down during the switch process of the cluster instance. Hence a cluster update followed by a RAC database update results in an instance getting restarted twice.

There may be several situations where the administrator may require more control over the switch process. For example, the administrator may choose to perform a node-wise update of the cluster and RAC databases in order to avoid multiple restarts of the database instances. The administrator may also need to perform node specific pre-post steps.

This option provides this control by enabling the user to perform the following tasks for each node:

- Switch the cluster instance
- Leave the RAC database instances shutdown

For example, consider RAC databases RACDB_112 and RACDB_121 are running on this cluster. The instances RACDB_112_1 and RACDB_121_1 running on this specific node will continue to remain shut down after the cluster instance is switched.

```
emcli db_software_maintenance -performOperation -name="Update Cluster"
-purpose=UPDATE_GI -target_type=cluster -target_list= CLUSTER1 -nor-
mal_credential="NC_HOST_CREDS:SYSMAN" -privilege_creden-
tial="HOST_PRIV:SYSMAN" -rolling=true -node_list="host1.example.com" -
startupDatabase=false
```

- Switch all the RAC database instances on the same node.
- This step will switch the instances RACDB_112_1 and RACDB_121_1 to the new home and will restart the same.

For example, consider RAC databases RACDB_112 and RACDB_121 are running on this cluster. The instances RACDB_112_1 and RACDB_121_1 running on this specific node will continue to remain shut down after the cluster instance is switched.

Verb

```
emcli db_software_maintenance -performOperation -name="Update RAC DB" -
purpose=UPDATE_RACDB -target_type=rac_database -target_list= RACDB -nor-
mal_credential="NC_HOST_CREDS:SYSMAN" -privilege_creden-
tial="HOST_PRIV:SYSMAN" -rolling=true -node_list="host1.example.com"
```

where:

- name: This is the unique name of the operation.
- purpose: Refer to EMCLI Command Inputs Based on Entity Type.
- target_type: Refer to EMCLI Command Inputs Based on Entity Type.
- target_list: This is a comma separated list of targets which need to be patched.
 - Targets of homogenous types are supported in a single fleet operation.
 - A unique list of hosts based on this target list is displayed and start stage of Oracle Home software on those hosts.
 - If targets running from the same Oracle home are provided in this list, the stage and deploy operation will be launched only once and not for all targets.
- normal_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- node_list: This is a comma separated list of hosts on which the instances need to be updated.

For example: If RACDB is running on a 4 node cluster host1, host2, host3, and host4 and you choose to update the instances in only 2 hosts at a time, the value of this parameter needs to be specified as node_list="host1, host2"

Note: This is an optional parameter. If no date is provided, the fleet operation will start immediately.

Sample Output

```
Processing target "RACDB"...
```

```
Checking if target is already running from the current version...
```

```
Check Passed.
```

```
Checking image subscription...
```

```
Check image subscription passed.
```

PDB Update

Update PDB moves a PDB from its predecessor container to the newly available successor container. This can be executed only if the Attach CDB / Deploy CDB operation has been performed.

```
emcli db_software_maintenance -performOperation -purpose="UPDATE_PDB" -  
target_list="<PDB_TARGET_Name>" -target_type="oracle_pdb" -name="Operation  
Name" -description="Operation description" -normal_credential="<credential  
name>" -privilege_credential="<privilege credential name>" -database_cre-  
dential="<SYSDBA credentials of the database>"
```

Update De-mystified

The following section highlights how Oracle Homes can be switched to the parallel Oracle Homes with least amount of downtime.

Database Type	Steps Performed
Single Instance	• Target is put in Blackout mode. • Configuration files like spfile, password, listener.ora, sqlnet.ora, tnsnames.ora are copied over to the parallel Home. • 'oratab' file entries are updated • Instance is brought down and brought up with the parallel home as the new Oracle Home. • Target is brought of Blackout mode. • For pre-12c databases, SQLs are run using SQL*Plus and for 12c database, the datapatch utility is called to make data changes.

Database Type	Steps Performed
Grid Infrastructure	• All relevant targets are put in Blackout mode. • Backup CRS configuration files from the source home (\$SRC_HOME/crs/utl/crsconfig_dirs). • Configuration files like spfile, password, listener.ora, sqlnet.ora, tnsnames.ora are copied over to the parallel Home. • 'oratab' file entries are updated • HAS Instance, ASM Instance, Management database instance, Scan Listeners and Grid Listeners and all database instances are brought down and brought up with Parallel Home as new Oracle Homes. • Configuration scripts are run. • All targets are brought out of Blackout mode. • If the management database is used, for pre-12c database, SQLs are run using SQL*Plus and for 12c database, datapatch utility is called to make data changes.
RAC Rolling Node	For each node: • Target is put in Blackout mode. • Configuration files like spfile, password, listener.ora, sqlnet.ora, tnsnames.ora are copied over to the parallel Home. • 'oratab' file entries are updated • Instance is brought down and brought up with Parallel Home as new Oracle Home. • Target is brought out of Blackout mode. • For pre-12c databases, SQLs are run using SQL*Plus and for 12c database, the datapatch utility is called to make data changes.
RAC Non-Rolling Node	• All RAC nodes are put in Blackout mode. • Configuration files like spfile, password, listener.ora, sqlnet.ora, tnsnames.ora are copied over to the parallel Home. • 'oratab' file entries are updated • Instance is brought down and brought up with Parallel Home as new Oracle Home concurrently. • For pre-12c databases, SQLs are run using SQL*Plus and for 12c database, the datapatch utility is called to make data changes.

Upgrading Oracle Homes

There are two types of Upgrades supported – namely minor and major. A minor upgrade is moving a Oracle Home from one patchset to another (Ex: 12.1.0.1 to 12.1.0.2) whereas a major upgrade is moving a database from one major release to another (ex: 12.2.0.1 to 18.3 or 18.x to 19.x). All the steps such as Gold Image creation, image subscription and deploy image remain the same. Please note that we don't support node-wise update for Upgrades. The update command for “Database Upgrade” requires database credentials to be provided as part of the command.



Note:

Refer to the Oracle Database Upgrade Guide for the supported upgrade paths.

The typical EMCLI command format is given below:


```
emcli db_software_maintenance -performOperation -name="Update RAC DB" -
purpose=UPDATE_RACDB -target_type=<target_type> -target_list==<List of
targets> -normal_credential="NC_HOST_CREDS:SYSMAN" -privilege_creden-
tial="HOST_PRIV:SYSMAN" -database_credential=<SYSDBA credentials of the
database>
```

In case of the primary and standby databases, upgrade needs to be done through the rolling upgrade user interface in Enterprise Manager. The wizard can be launched from the target database home page of the primary or standby database. On the Database home page, select **Oracle Database, Provisioning**, then select **Upgrade Database**. Refer to <Data Guard Rolling Upgrade> section for more details.

Cleaning Up Oracle Homes

Since Oracle homes consume storage, they can be cleaned up to release space when they are no longer in use. To achieve this, set the following emctl property and run the emctl command, also shown below to delete unused homes of a given target. Oracle homes will be deleted only if there are no targets running from them.

```
emctl set property -name oracle.sysman.emInternalSDK.db.gis.lineagelength
-sysman_pwd $em_sysman_pwd -value <value>
```

where value is the number of Oracle Homes that must be retained. The default value is 3.

For example, if there are four versions of an Oracle Home (with the target running from the fourth version), only the earliest of them can be deleted if the value is set to 3. Similarly if the value is set to 1, only the current Oracle home from which the target is running will remain and the other three homes will be deleted.

```
emcli db_software_maintenance -performOperation -name=" User specified Op-
eration name " -purpose=<Operation Type> -target_list=<List of targets> -
target_type=<target type> -normal_credential=<credential_name> -privi-
lege_credential=<privilege credential_name> -workDir=/tmp [-reportOn-
ly=true|false>
```

Description of Input Variables

- name: This is the unique name of the operation.
- purpose: CLEANUP_SOFTWARE
- target_type: Refer to the EMCLI Command Inputs Based on Entity Type table.
- target_list: This is a comma separated list of targets which need to be patched.
- normal_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.

- <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- reportOnly: The default value is false. When set to 'true', it prints the eligible homes for deletion but the homes are not deleted.

Custom Pre and Post Scripts for Fleet Operations

In order to support automated maintenance activities for different Fleet Maintenance operations, the support for configuring Pre/Post scripts for different operations is implemented. Pre/post scripts need to be uploaded as EM Software Library Entity (Directive) prior to their usage with Fleet Maintenance operations. Pre/post script usage is supported with DEPLOY, UPDATE, CLEANUP and ROLLBACK operations.

Creating and Using Custom Pre/Post Scripts:

1. Develop your pre/post script. Review **Custom Script Considerations**.
2. Upload custom scripts to the Software Library. Review **Software Library Entity for uploading the script(s)**.
3. Capture the URN for SWLIB Entity. Review **How to get URN for SWLIB entity**.
4. Specify the Pre/Post script SWLIB URN in EMCLI input file. Review **Usage**.

Custom Script Considerations

Customer provided script will be invoked during Fleet Maintenance operations by passing a single argument absolute path to an input parameter file containing several parameters as key value pairs. This input parameter file (scriptInputFile.txt) is generated as part of the procedure activity and relevant DP variables are added to this file. These key value pair can be read from this file in the pre/post script and are intended to be used for maintenance operation being automated by the custom script. Input parameter file contains parameter values such as current oracle home, new oracle home etc. Sample input parameter file provided below:

CLEANUP_SOFTWARE operation - Sample

```
MAINTENANCE_PURPOSE=CLEANUP_SOFTWARE
CUSTOM_PRE_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Directives:none:74A730047930C5FDE053DF0FC40A3E69:0.1
ORACLE_HOME_TARGET_DETAILS=abc01def.example.com:/scratch/cuser/app/cuser/product/12.2.0/dbhome_1;
CUSTOM_POST_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Directives:none:74A730047935C5FDE053DF0FC40A3E69:0.1
DEPLOY_TYPE=SIDB
ORACLE_HOME_NORMAL_CRED_NAME=CUSER:SYSMAN
SOURCE_HOST=abc01def.example.com
ORACLE_HOME_ROOT_CRED_NAME=ORACLE_ROOT:SYSMAN
```

DEPLOY_DB_SOFTWARE operation - Sample

```
MAINTENANCE_PURPOSE=DEPLOY_DB_SOFTWARE
SOURCE_HOME_LOCATION=/scratch/cuser/app/cuser/product/12.2.0/dbhome_2
```

```
target_type=oracle_database
target_name=sidb122
NOT_WINDOWS=true
ORACLE_HOME_NORMAL_CREDSET_NAME=HostCredsNormal
DISPATCHER_LOC=/tmp
GOLD_IMAGE_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Compo-
nent:SUB_OracleDB:74AA263BE602CF3BE053DF0FC40A2ED6:0.1
ROOT_SCRIPT_LOC=/tmp/1311624666555972//
MASSDB-194177ce43d441b89a504584d146aafb
CUSTOM_PRE_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Di-
rectives:none:74A730047930C5FDE053DF0FC40A3E69:0.1
CUSTOM_POST_SCRIPTS_URN=oracle:defaultService:em:provisioning:
1:cmp:COMP_Directives:none:74A730047935C5FDE053DF0FC40A3E69:0.1
ORACLE_HOME_TARGET_DETAILS=den01nre.example.com:OraDB12Home2_4_den01nre.ex-
ample.com_3012;
WORK_DIR_LOC=/tmp/1311624666555972//1311626758529303/
DEPLOY_TYPE=SIDB
CREDSET_TARGET_TYPE=OracleHome
ORACLE_HOME_ROOT_CREDSET_NAME=HostCredsPriv
NEW_ORACLE_HOME_LIST=/scratch/cuser/app/cuser/product/12.2.0/dbhome_Dep2
```

Software Library Entity For Uploading The Script

Create a Software Library entity of type Directive (for uploading the custom scripts to EM).

1. Select **Enterprise > Provisioning and Patching > Software Library**.
2. Select the directory where you want to create the **Software Library** entity.
3. Click **Actions > Create Entity > Directives**.
4. Provide all the basic details on the **Directive** screen.
5. Click **Add** under **Command Line Arguments** on the **Configure** screen.
6. The **INPUT_FILE** to be added as mentioned in the list below. Click **OK**.
 - Argument Prefix="
 - Argument Suffix="
 - Property Name=INPUT_FILE
7. Click **Next**.
8. Click **Add** in **Specify Source** and select the script files to upload the custom script on the **Select File** screen. Choose the driver script as the main file.
9. Click **Next**.
10. Click **Save and Upload**.
11. An entity is created successfully in a Software Library.

How to Get URN for SWLIB Entity

To get URN for the SWLIB through EMCLI:

```
$OMS_HOME/bin/emcli list_swlib_entities -show_entity_rev_id -name="<Name
of SWLIB entity>"
```

To get URN for the SWLIB through EM UI:

1. Select **Enterprise > Provisioning and Patching > Software Library**.
2. Click **View > Columns > Internal ID**. Select this option.
3. Go to the uploaded entity and copy the **Internal ID** for the uploaded entity. The internal ID is the URN.

Pre/Post Script Usage

To execute the Custom Scripts, use `<-input_file>` option with fleet command and provide the URN for the pre/post script SWLIB entity created earlier.

```
emcli db_software_maintenance
-performOperation -name="Deploy 12.2 Home"
-purpose="DEPLOY_DB_SOFTWARE"
-target_type=oracle_database
-target_list="DB122"
-normal_credential="ORACLE:SYSMAN"
-privilege_credential="ORACLE_ROOT:SYSMAN"
-input_file="data:/scratch/input_db0H.prop"
```

Contents of input_db0H.prop file:

```
NEW_ORACLE_HOME_LIST=/scratch/cuser/app/cuser/product/12.1.0/dbhome_Dep1
CUSTOM_PRE_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Di-
rectives:none:74A730047930C5FDE053DF0FC40A3E69:0.1
CUSTOM_POST_SCRIPTS_URN=oracle:defaultService:em:provisioning:
1:cmp:COMP_Directives:none:74A730047935C5FDE053DF0FC40A3E69:0.1
```

Database Fleet Maintenance - RESTful APIs

Gold Image REST APIs

Gold Image is the end of state software definition that contains information about the base software version along with the additional patches. The 'end-state definition' is the logical term; the physical software binary component that represents the end-state is called as a gold image. For more information on gold image , see the standard configuration described in [Discovering Configuration Pollution](#) .

Get a List of Software Images

Features	Description
Request Method	GET
URI	<code>https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/websvcs/restful/emws/db/goldimage/softwareimages</code>
Request Headers	Authorization: Basic Accept: application/json

Features	Description
Response	<pre>{ "totalResults": 2, "items": [{ "name": "RAC121_a", "id": "4B7738536B6E7888E053057FB10ACF8C", "description": "RAC121", "version": "12.1.0.2.0", "platformName": "Linux x86-64", "dateCreated": "2017-03-24 15:09:21.0", "owner": "SYSMAN", "lastModifiedBy": "SYSMAN" }, { "name": "RACImage11204", "id": "4B620EC24DCE61FAE053057FB10AC7D0", "description": "RACImage11204", "version": "12.1.0.2.0", "platformName": "Linux x86-64", "dateCreated": "2017-03-23 13:16:51.0", "owner": "SYSMAN", "lastModifiedBy": "SYSMAN" }] }</pre>

Create Software Image

Features	Description
Request Method	POST
URI	<code>https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/websvcs/restful/emws/db/goldimage/softwareimages</code>
Request Headers	Authorization: Basic Content-Type: application/json
Payload	<pre>{ "imageName": "RAC121_a", "imageDescription": "RAC121", "refTargetName": "OraDB12Home1_1_example.exampledomain.com_190", "imageSwLibLocation": "DB Provisioning/12.1.0.2.0/goldimage", "refHostCredential": "NC_HOST_2017-03-18-103539:SYSMAN", "workingDirectory": "/tmp", "swLibStorageType": "OmsShared", "swLibStorageName": "swlib", "versionName": "PSUNo" }</pre>

Features	Description
Response	<pre>{ "messages": ["Create Gold Image operation has been submitted successfully with the instance name : 'CreateGoldImageProfile_SYS-MAN_03_27_2017_10_13_AM' and execution_guid=4BB1192C1A2F2AB3E053057FB10A792E", "You can track the status of operation using the following:", "EMCLI: emcli get_instance_status - exec=4BAFFB4FD4ED1B34E053057FB10A99BF", "Browser: https://blr123.example.com:11111/em/faces/core-jobs-procedureExecutionTracking?execution-GUID=4BB1192C1A2F2AB3E053057FB10A792E"] }</pre>

Description of the Input Variables in payload

- **imageName:** The name of the gold image. This name must be unique across images.
- **imageDescription:** A description of the image.
- **refTargetName:** The Oracle home target that will be used to create this gold image. This is the database or Grid Infrastructure Oracle Home from the existing environment on which the 11.2.0.4 PSU and all the one-off patches have been applied. To find the reference target name, enter the following query on the Enterprise Manager repository:


```
SELECT distinct target_name FROM mgmt$target_properties
WHERE target_name IN (SELECT target_name FROM mgmt_targets
WHERE target_type='oracle_home'
AND host_name=<Host Name of this Oracle Home>
AND property_name='INSTALL_LOCATION'
AND property_value=<path of Oracle Home>)
```
- **imageSwLibLocation:** Path in software library where the payload of the gold image will be stored.
- **refHostCredential:** This is applicable for Database Gold Image only. This must be entered in the format <Named Credential: Credential Owner>
- **workingDirectory:** The temporary location on the host of reference Oracle home target.
- **swLibStorageType:** The Software Library storage type. This can be OMS Shared or OMS Agent File system.
- **swLibStorageName:** The storage name for the Software Library. To retrieve this storage name, from the Enterprise menu, select **Provisioning and Patching**, then select **Software Library**. In the Software Library page, from the **Actions** menu, select **Administration**. On this page, check the value in the **Name** column to retrieve the storage name for the Software Library.

- **versionName:** An image can have multiple versions. A default version of the image is created and more versions can be added.

Search Software Images

The search can be filtered using various inputs as query parameters. The allowed query parameters are:

- Name
- Version
- Description
- Owner
- Target_type

Following is an example using query parameter “Name”:

Features	Description
Request Method	GET
URI	https://<OMS_CONSOLE_HOST>/em/websvcs/restful/emws/db/goldimage/softwareimages/search?name=%RAC%_a
Request Headers	Authorization: Basic Accept: application/json
Response	<pre>{ "totalResults": 1, "items": [{ "name": "RAC121_a", "id": "4B7738536B6E7888E053057FB10ACF8C", "description": "RAC121", "version": "12.1.0.2.0", "platformName": "Linux x86-64", "dateCreated": "2017-03-24 15:09:21.0", "owner": "SYSMAN", "lastModifiedBy": "SYSMAN" }] }</pre>

Delete a software image

Features	Description
Request Method	DELETE
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/websvcs/restful/emws/db/goldimage/softwareimages>
Request Headers	Authorization: Basic Content-Type: application/json

Features	Description
Payload	<pre>{ "image_id" : "4B7738536B6E7888E053057FB10ACF8C" }</pre>
Response	<pre>{ "messages" : "Image with id=4B7738536B6E7888E053057FB10ACF8C de- leted successfully" }</pre>

Subscribing Targets to the Selected Image

Features	Description
Request Method	POST
URI	<code>https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/websvcs/restful/emws/db/goldimage/imagesubscriptions</code>
Request Headers	Authorization: Basic Content-Type: application/json
Payload	<pre>{ "imageId" : "4B620EC24DCE61FAE053057FB10AC7D0", "targetName" : "Exemplerep_Database", "targetType" : "example_database" }</pre>
Response	<pre>{ "messages": ["Target 'Exemplerep_Database' subscribed successfully."] }</pre>

Description of the Input Variables in payload

- `imageId`: The GUID of the image.
- `targetType`: The target type of the target to be subscribed to this image.
- `targetName`: The name of the target.

Listing Subscriptions of an Image

Features	Description
Request Method	GET

Features	Description
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/websvcs/restful/emws/db/goldimage/imagesubscriptions?image_id=4B620EC24DCE61FAE053057FB10AC7D0
Request Headers	Authorization: Basic Accept: application/json
Response	[{ "imageId": "4B620EC24DCE61FAE053057FB10AC7D0", "targetGuid": "207B57A3FE300C86F81FE7D409F5DD1C", "dateSubscribed": "2017-03-24 16:01:25.0" }]

UnSubscribing Targets from the Selected Image

Features	Description
Request Method	DELETE
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/websvcs/restful/emws/db/goldimage/imagesubscriptions
Request Headers	Authorization: Basic Content-Type: application/json
Payload	{ "imageId" : "4B620EC24DCE61FAE053057FB10AC7D0", "targetName" : "Exemplerep_Database", "targetType" : "example_database" }
Response	{ "messages": ["Target 'Exemplerep_Database' unsubscribed successfully."] }

Description of the Input Variables in payload

- imageId: The GUID of the image.
- targetType: The target type of the target to be unsubscribed from this image.
- targetName: The name of the target.

Listing Target Subscriptions of an Image

Features	Description
Request Method	PUT
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/websvcs/restful/emws/db/goldimage/targetsubscriptions
Request Headers	Authorization: Basic Content-Type: application/json
Payload	<pre>{ "imageId" : "4B7738536B6E7888E053057FB10ACF8C ", "targetName" : "Exemplerep_Database", "targetType" : "example_database" }</pre>
Response	<pre>{ "totalResults": 1, "items": [{ "targetName": "Exemplerep_Da- tabase", "targetType": "example_data- base", "imageId": "4B7738536B6E7888E053057FB10ACF8C", "imageName": "RAC121_a", "subscriptionDate": "2017-03-27 10:06:00.0" }] }</pre>

Description of the Input Variables in payload

- imageId: The GUID of the image.
- targetType: The target type of the target.
- targetName: The name of the target.

Creating Version of Image

Features	Description
Request Method	POST
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/websvcs/restful/emws/db/goldimage/versions
Request Headers	Authorization: Basic Content-Type: application/json

Features	Description
Payload	<pre>{ "imageName": " 4B7738536B6E7888E053057FB10ACF8C", "refTargetName": "Or- aDB12Home1_1_slc06nag.example.com_190", "imageSwLibLocation": "DB Provision- ing/12.1.0.2.0/goldimage", "refHostCredential": "NC_HOST_2017-03-18-103539:SYSMAN", "workingDirectory": "/tmp", "swLibStorageType": "OmsShared", "swLibStorageName": "swlib", "versionName": "PSUNo" }</pre>
Response	<pre>{ "messages": ["Create Gold Image oper- ation has been submitted successfully with the instance name : 'CreateGoldImageProfile_SYS- MAN_03_27_2017_10_13_AM' and execu- tion_guid=4BB1192C1A2F2AB3E053057FB10A792 E", "You can track the status of opera- tion using the following:", "EMCLI: emcli get_instance_status - exec=4BAFFB4FD4ED1B34E053057FB10A99BF", "Browser: https://blr123.exam- ple.oracle.com:11111/em/faces/core-jobs- procedureExecutionTracking?execution- GUID=4BB1192C1A2F2AB3E053057FB10A792E"] }</pre>

Description of Input parameters in payload

- **imageId:** The GUID of the image.
- **refTargetName:** The Oracle home target that will be used to create this gold image. This is the database or Grid Infrastructure Oracle Home from the existing environment on which the 11.2.0.4 PSU and all the one-off patches have been applied. To find the reference target name, enter the following query on the Enterprise Manager repository:

```
SELECT distinct target_name FROM mgmt$target_properties
WHERE target_name IN (SELECT target_name FROM mgmt_targets
WHERE target_type='oracle_home'
AND host_name=<Host Name of this Oracle Home>
AND property_name='INSTALL_LOCATION'
AND property_value=<path of Oracle Home>)
```
- **imageSwLibLocation :** Path in software library where the payload of the gold image will be stored.
- **refHostCredential:** This is applicable for Database Gold Image only. This must be entered in the format<Named Credential: Credential Owner>

- **workingDirectory**: The temporary location on the host of reference Oracle home target.
- **swLibStorageType**: The Software Library storage type. This can be OMS Shared or OMS Agent File system.
- **swLibStorageName**: The storage name for the Software Library. To retrieve this storage name, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**. In the Software Library page, from the **Actions** menu, select **Administration**. On this page, check the value in the **Name** column to retrieve the storage name for the Software Library.
- **versionName**: An image can have multiple versions. A default version of the image is created and more versions can be added.

Listing Versions

Features	Description
Request Method	GET
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/websvcs/restful/emws/db/goldimage/versions?image_id=4B7738536B6E7888E053057FB10ACF8C
Request Headers	Authorization: Basic Content-Type: application/json
Response	<pre>{ "totalResults": 1, "items": [{ "name": "PSUNos", "versionId": "4BA-FA861A7ED6603E053057FB10A0588", "imageId": "4B7738536B6E7888E053057FB10ACF8C", "externalId": "oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_OracleDB:4BA-FA861A7E96603E053057FB10A0588:0.1", "status": "CURRENT", "position": "1", "hashCode": "C1622270664:B<NO_PATCHES>", "dateCreated": "2017-03-27 10:13:28.0", "provisioningStatus": "READY" }] }</pre>

Deleting Versions

Features	Description
Request Method	DELETE

Features	Description
URI	https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/websvcs/restful/emws/db/goldimage/versions
Request Headers	Authorization: Basic Content-Type: application/json
Payload	{ "versionId": "4BB099D42CC8671EE053057FB10A07DA" }
Response	{ "messages": ["The following version will be deleted. 4BB099D42CC8671EE053057FB10A07DA"] }

Deploy APIs

Feature	Description
Request Method	POST
URI	/em/websvcs/restful/emws/db/fleetmaintenance/performOperation/deploy
Request Headers	Authorization: basic <generated credential value> Content-Type: application/json
Body	{ "name": "<Operation Name>", "targetName" : "<target name>", "targetType" : "<Target type>", "normalCredString" : "KKHANUJA:SYSMAN", "privCredString" : "KKHANUJA_ROOT:SYSMAN", "standbyAutoDeploy" : "", "skipCVUChecks" : "", "skipPrereqs" : "", "newOracleHomeLoc" : "/scratch/kkhanuja/oh1710" }

Description of the Input Variables

- name: Name of the operation.
- targetName: Name of the target.
- targetType: Refer to CLI Command Inputs Based on Entity Type table.
- normalCredString: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.

- <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privCredString: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- standbyAutoDeploy: This is an optional parameter with default value as True. If the target is a primary database, a new Oracle home using the same gold image version as Primary is deployed automatically on the Standby host. This parameter disables the automatic deployment of software on standby host when specified as false. If this value is false, the Standby staging/deploy operation can be performed independently using the emcli command.
- skipCVUChecks: true|false. Default value is false
- skipPrereqs: true|false. Default value is false
- WorkingDirectory: <Name of temp directory>
- newOracleHomeLoc: <Location where Oracle Home will be created>

A sample of the response received is shown below:

Update Oracle home software operation for RAC Database t05nwe can be monitored using the following EMCLI command:

```
emcli get_instance_status -exec=4943D8E3626E420EE0530854F10A18A9 -details -xml
```

Browser: <https://blr123.example.com:11111/em/faces/core-jobs-procedureExecution-Tracking?executionGUID=4943D8E3626E420EE0530854F10A18A9>

Update / Rollback

Feature	Description
Request Method	POST
URI for Update	em/websvcs/restful/emws/db/fleetmaintenance/performOperation/update
URI for Rollback	em/websvcs/restful/emws/db/fleetmaintenance/performOperation/rollback
Request Headers	Authorization: basic <generated credential value> Content-Type: application/json

Feature	Description
Body	<pre>{ "name": "<Operation name>", "targetName": "<name of the target>", "targetType": "<target type>", "normalCredString": "<credential name>", "privCredString": "<privileged credential name>" "nodeList": "<list of nodes>" "dgRole": "primary standby" "startupAfterSwitch": "true false" "startupDBAfterSwitch": "true false" "skipSwitchDatabase": "true false" "ignoreStandbyPrereq": "true false" }</pre>

Description of the Input Variables

- name: Name of the operation
- targetName: Name of the target
- targetType: Refer to CLI Command Inputs Based on Entity Type table.
- normalCredString: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privCredString: This must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where new Oracle home will be deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

- nodeList: comma separated list of nodes
- dgRole: primary|standby
- startupDBAfterSwitch: The default value is true. It is used for node-wise patching. To leave a database instance down after a cluster instance is patched
- skipSwitchDatabase: It should be set to 'true'. It skips the switching of database to patched home since it is already performed in earlier steps.
- ignoreStandbyPrereq: The default value is false. This disables verification check if the Standby is on the same image version where the Primary is being moved to
- startupAfterSwitch: For non-rolling patches, it is necessary to bring up the Primary database first from the patched home. It will ensure that Standby is left in shutdown state during the Update operation. Standby will be started in a separate step after Primary is patched and started.

A sample of the response received is shown below:

Update Oracle home software operation for RAC Database t05nwe can be monitored using the following EMCLI command:

```
emcli get_instance_status -exec=4943D8E3626E420EE0530854F10A18A9 -details  
-xml
```

Browser: <https://blr123.example.com:11111/em/faces/core-jobs-procedureExecution-Tracking?executionGUID=4943D8E3626E420EE0530854F10A18A9>

Performing Engineered System Software Updates

This chapter explains how you can use Enterprise Manager Cloud Control 13c to apply Quarterly Full Stack Download Patch (QFSDP) on Exadata and Exalytics targets. In particular, this chapter covers the following:

- [Overview of Exadata System Software Update](#)
- [Configuring Options for Exadata Component Software Updates](#)
- [Updating Exadata Database Servers](#)
- [Updating Exadata Storage Servers](#)
- [Updating Exadata Infiniband Switches](#)
- [Rolling Backup Deployed Software Updates](#)
- [Patching Oracle Identity Management Targets](#)
- [Overview of Exalytics System Software Update](#)
- [Configuring the Options for Oracle Exalytics Updates](#)
- [Updating Oracle Exalytics Compute Nodes](#)
- [Updating Oracle Exalytics Business Intelligence Instance](#)

Overview of Exadata System Software Update

Using Quarterly Full Stack Download Patches (QFSDP), you can update each component of an Exadata target, individually, independent of the other Exadata target components. Also, you can apply a different QFSDP on each Exadata target component. This approach minimizes downtime, as you can update a few Exadata component targets at a time, keeping the other components up and running. You need not update all the Exadata target components at once.

To update Exadata target components, you must use the Exadata Target Maintenance page. To access the Exadata Target Maintenance page, from the **Targets** menu, select **Exadata**. Click the Exadata target that you want to update. On the Exadata target home page, from the **Database Machine** menu, select **Software Update**.




Table 48-1 Exadata Update Information

Parameter Name	Description
Recommended QFSDP	The patch recommendation feature for Exadata can function in two different modes: Online Mode and Offline Mode. In the Online Mode, the Download System Patch & Evaluate Patch Recommendations Job downloads the QFSDPs from the support site, and evaluates the Patch Recommendations. In the Offline Mode, you must upload the QFSDP to Software Library using EMCLI, following which, the Download System Patch & Evaluate Patch Recommendations Job checks for patches in the Software Library, and evaluate the Patch Recommendations for discovered Engineered System Targets. The Recommended QFSDP section provides a link to the readme of the patch, and the timestamp of the actual patch download.
Last QFSDP Check	The QFSDP check happens at an interval of 24 hours. The time when the last check happened is recorded in this section.
Targets	All the targets configured on the Exadata system and their status is recorded in this section.
Targets Needing Updates	The targets that are part of the system but are not at the latest patch level are the ones that need to be patched. The information about the number of targets that must be patched to bring them upto the recommended patch level is available here.
Updates in Progress	When you deploy or rollback a patch, the deployment procedure is submitted and the updates are said to be in progress. The number of targets that are involved in the update process is mentioned here.
Components	The Software Update page enables you to update the following components of the Exadata targets: - Oracle Database Server - Oracle Exadata Storage Server - Oracle Infiniband Switch For any of the selected targets, you can view the following details: - Last QFSDP Applied: This is the latest patch applied on the selected component. - Update Status: The different status of the targets. For details, see Table 48-2.

Table 48-2 Update Status

Status	Description
	This symbol implies that the selected target(s) of the components require updates to bring them to the recommended patch level.
	This symbol implies that the patch is currently being deployed, and the procedure is in progress. If it is successful, then a tick mark will appear, if not you will see a cross mark against the respective targets.
	This symbol implies that the selected target of the component is up-to-date with the latest patch.
	This symbol implies that patching failed, and that the patch could not be applied on the designated targets.

For a selected component, in addition to details like the target name, status, version, update status, and when the last analyze and deploy happened, you have the following details:

- Selected Patch Details: Details like the bug fixed and knowledge articles related the update are displayed here.
- Operation Status: Information of the current operation on the targets is available in this section. For more granular details, select the Operation Tab to view the:
 - Summary: All the information that you have entered is available in the summary section. After the Analyze step completes, you can view the patching steps. Click **Download** to download and view the Oplan-based steps that are generated.
 - Problems to Resolve: Any issue encountered during analysis and deploy phase is recorded in this section.
 - Information Messages: This section contains the information text.
 - Log Files and Diagnostics: When the Analyze, Deploy, or Rollback phase is in progress, a list of static links are displayed. Click on the link to fetch the on-demand files and view the details of the log files and trace files that contain information about the patching steps performed on the targets.

After the patching process is complete, the log files are automatically fetched and stored in the OMS Repository. Click the link to download the zip file. The zip contains log and trace files that contain information about the patching steps are performed on the targets.

Configuring Options for Exadata Component Software Updates

Before patching an Oracle Exadata target component, you can choose to set up the options for patching, such that the options are reused for subsequent patching of targets within the same component present in the same Exadata system. Setting up the patch options enables you to proceed with the analysis phase of a subsequent Exadata target component patching operation without having to re-enter the options.

To set up the update options for a particular Exadata component, follow these steps:

1. From the **Targets** menu, select **Exadata**.
2. Click the name of an Exadata target.
3. On the Exadata target home page, from the **Database Machine** menu, select **Software Update**.
4. In the Component Navigation section, select the required Exadata component.
5. Click **Settings**.
6. Specify the following options:

- **Operation Mode:** *(only for Oracle Exadata Storage Server and Oracle Infini-band Switch targets)* Select the mode of patching that you want to use. You can choose to patch the component targets in rolling mode or parallel mode. Note that for an Infiniband Switch target, only the rolling mode of patching is supported.

Out of the component targets selected for patching, if you want to patch a single target at a time, select **Rolling**. This is the default option, and it involves very little downtime, see [Pausing the Patching Operation While Patching Targets in Rolling Mode](#).

However, if you want to simultaneously patch all the component targets selected for patching, select **Parallel**. This option involves downtime, as all the component targets are shut down for a significant period. However, this process consumes less time, as all the component targets are patched simultaneously.

- **Where to Stage:** If you want the wizard to stage the patches from Software Library to a temporary location before the patch is applied on the component target, select **Yes** for **Stage Patches**, then specify the location where you want the patches to be staged. However, if you have already manually staged the patches that you want to apply, select **No** for **Stage Patches**, then specify the location where you staged the patches. If the stage location is a shared location, select **Shared Location**.

To manually stage QFSDP, download the patch, navigate to the location (parent directory) where you want to stage the patch, create a subdirectory with the same name as the patch zip file, then extract the contents of the patch zip file into this subdirectory.

For example, if you downloaded patch 699099.zip, and the stage location, which is the parent directory, is /u01/app/example/em/stagepatch, then, in this parent directory, create a subdirectory titled 699099 and extract the contents of the zip file. Specify /u01/app/example /em/stagepatch as the stage location.

- **Credential Information:** Provide the required credentials. You can choose to use preferred credentials, or override them with different credentials.
- **Customization:** For patching Oracle Database Server targets, you can select the following options:
 - **Restart target after deployment:** Select this option to restart the Oracle Database Server target after applying the Quarterly Full Stack Download Patch (QFSDP).
 - **Run script before deployment:** Select this option to run a script on the Oracle Database Server target before deploying the QFSDP. Specify the location where the script is hosted.

- **Run script after deployment:** Select this option to run a script on the Oracle Database Server target after deploying the QFSDP. Specify the location where the script is hosted.
 - **Notifications:** Specify whether or not you want to enable email notifications when the patching operation is scheduled, starts, requires action, is suspended, succeeds, and fails.

To enable email notifications, select **Receive notification emails when the patching process**, then select the required options. If a warning message, mentioning that the sender or the receiver email address is not set up, is displayed, perform the action mentioned in the warning.
7. Click **Apply** to save the settings. Once saved, the setting will be available for the next set of patching operation.

Updating Exadata Database Servers

To patch the Exadata Database Server or Compute Node Server components of an Exadata target, follow these steps:

1. From the **Targets** menu, select **Exadata**.
2. Click the name of the Exadata target that you want to patch.
3. On the Exadata target home page, from the **Database Machine** menu, select **Software Update**.
4. In the Component Navigation section, select **Oracle Database Server**.
5. Select the targets that you want to patch, then click **Select QFSDP**.

The Select Quarterly Full Stack Download Patch window displays a list of QFSDPs that can be applied on the Oracle Database Server target. The recommended (latest) patches are marked with a tick. Select the QFSDP that you want to apply, then click **Select**.

6. Click **Analyze**.

Specify a schedule for the analysis. If you haven't specified any schedule, the Deployment Procedure will run immediately.

Verify the provided options. If you haven't already specified the options (as described in [Configuring Options for Exadata Component Software Updates](#)), provide the required options.

Click **Submit**.

7. After the analysis is complete, click **Deploy**. Deploy uses the options selected for Analyze. Once deployed, the Last QFSP Applied is applied to the update, and the current revision will show the latest version.

Updating Exadata Storage Servers

To patch the Oracle Exadata Storage Server components of an Exadata target, follow these steps:

1. From the **Targets** menu, select **Exadata**.
2. Click the name of the Exadata target that you want to patch.

3. On the Exadata target home page, from the **Database Machine** menu, select **Software Update**.
4. In the Component Navigation section, select **Oracle Exadata Storage Server**.
5. Select the targets that you want to patch, then click **Select QFSDP**.

The Select Quarterly Full Stack Download Patch window displays a list of QFSDPs that can be applied on the Oracle Database Server target. The recommended (latest) patches are marked with a tick. Select the QFSDP that you want to apply, then click **Select**.

Figure 48-1 Select QFSDP

Select Quarterly Full Stack Download Patch X

[Search Patch](#)

▲ **View Software Update Knowledge**

888828.1 [Exadata Database Machine and Exadata Storage Server Supported Versions](#)

1270094.1 [Exadata Critical Issues](#)

Patch Number	Description	Update Version	Product/Family	Updated	Size
21206075 ✓	QUARTERLY FULL STACK DOWNLOAD PATCH FOR EXADATA (JUL 2015 - 12.1.2.1)	12.1.2.1...	Oracle Database	10+ hours ago	6.1 GB

✓ Recommended Patch

[Cancel](#) [Select](#)

[Read Me File](#)

☆ **Patch 21206075: QUARTERLY FULL STACK DOWNLOAD PATCH FOR EXADATA (JUL 2015 - 12.1.2.1)**

Last Updated Jan 20, 2015 11:37 PM (8+ months ago) Size 6.1 GB

Product	Oracle Database Family	Download Access	Software
Release	Oracle 12.1.0.2.0	Classification	Security
Platform	Linux x86-64	Patch Tag	

The Software Update Knowledge section contains information about software patches and releases for Oracle Exadata Database Machine running Exadata Storage Server Software and a list fixes and workarounds that are deemed critical for Exadata Storage Server.

For information on installing the Quarterly Full Stack Download Patch for Oracle Exadata, click **ReadMe File**.

6. Click **Analyze**.

Specify a schedule for the analysis. If you haven't specified any schedule, the Deployment Procedure will run immediately.

Verify the provided options. If you haven't already specified the options (as described in [Configuring Options for Exadata Component Software Updates](#)), provide the required options.

Click **Submit**.

**Note:**

Target-level patch monitoring has been enabled for Oracle Exadata Storage Server components of an Exadata target. When there are multiple targets selected for patching, monitoring the patching status of individual targets is possible, and does not depend on the deployment procedure completion or failure.

For example, assuming host 1, host 2, and host 3 are being patched, and the deployment procedure fails because patching host 3 wasn't successful; target-level patch monitoring still allows you to monitor the patching status of host 1 and host 2 independently.

7. After the analysis is complete, click **Deploy**.

Updating Exadata Infiniband Switches

To patch the Oracle Infiniband Switch components of an Exadata target, follow these steps:

**Note:**

The Exadata Infiniband Switches supports only component-based update. You cannot update individual targets.

1. From the **Targets** menu, select **Exadata**.
2. Click the name of the Exadata target that you want to patch.
3. On the Exadata target home page, from the **Database Machine** menu, select **Software Update**.
4. In the Component Navigation section, select **Oracle Infiniband Switch**.
5. Select the targets that you want to patch, then click **Select QFSDP**.

The Select Quarterly Full Stack Download Patch window displays a list of QFSDPs that can be applied on the Oracle Database Server target. The recommended (latest) patches are marked with a tick. Select the QFSDP that you want to apply, then click **Select**.

6. Click **Analyze**.

Specify a schedule for the analysis. If you haven't specified any schedule, the Deployment Procedure will run immediately.

Verify the provided options. If you haven't already specified the options (as described in [Configuring Options for Exadata Component Software Updates](#)), provide the required options.

Click **Submit**.

7. After the analysis is complete, click **Deploy**.

Rolling Backup Deployed Software Updates

To rollback a Quarterly Full Stack Download Patch (QFSDP) that is deployed on an Oracle Exadata component, follow these steps:



Note:

You can not select the patch you want to rollback, instead, you can only rollback the latest QFSDP applied using Cloud Control.

1. From the **Targets** menu, select **Exadata**.
2. Click the name of the Exadata target on which the Quarterly Full Stack Download Patch (QFSDP) that you want to rollback is deployed.
3. On the Exadata target home page, from the **Database Machine** menu, select **Software Upgrade**.
4. In the Component Navigation section, select the Exadata target component on which the QFSDP that you want to rollback is deployed.
5. *(Applicable for Oracle Exadata Storage Server and Oracle Infiniband Switch only)* From the **Rollback** menu, select **Prerequisite Only** to analyze if the patch is suitable for a roll back. If the analysis is successful, you can proceed to the next step.
6. From the **Rollback** menu, select **Rollback** to rollback to the last deployed patch.
Specify a schedule for the rollback operation. Provide the required options, then click **Rollback**.

Patching Oracle Identity Management Targets

Using patch plans, you can patch those Oracle Access Management Server and Oracle Identity Management Server targets that were provisioned using Identity Management Lifecycle tools. Other Oracle Identity Management targets are not supported for patching.

To patch Oracle Access Management Server and Oracle Identity Management Server targets (that were provisioned using Identity Management Lifecycle tools) using patch plans, you must have an Identity Management Pack Plus license. Also, the 12.1.0.6 Enterprise Manager for Oracle Fusion Middleware plug-in (or higher) must be deployed on the OMS, and on all the Management Agents running on the hosts on which the Oracle Access Management Server and Oracle Identity Management Server targets are deployed.

To patch Oracle Access Management Server and Oracle Identity Management Server targets using a patch plan, follow these steps:

1. Identify the patches that you want to apply, using patch recommendations, or by searching for the patches.

For information on how to search for patches, see [Searching for Patches on My Oracle Support](#) and [Searching for Patches in Oracle Software Library](#).

2. Create, analyze, and deploy a patch plan containing the required patches, as described in [Creating, Analyzing, Preparing, and Deploying Patch Plans](#).

Overview of Exalytics System Software Update

Enterprise Manager Cloud Control 13c enables you to apply Patch Set Updates (PSU) on Exalytics targets. Using PSUs, you can patch an Exalytics target component individually, independent of the other Exalytics target components. Also, you can apply a different Patch Set Update (PSU) on each Exalytics target component. This approach minimizes downtime, as you can patch a few Exalytics target components at a time, while keeping the other components up and running.

For example, you can apply a particular Patch Set Updates on two compute nodes of an Exalytics target at a particular point in time, and apply a different PSUs on two other compute nodes of the Exalytics target one hour later. You need not patch all the compute nodes of an Exalytics target at once.

To patch Exalytics target components, you must use the Exalytics Patching page. To access the Exalytics Patching page, from the **Targets** menu, select **Exalytics**. Click the Exalytics target that you want to patch. On the Exalytics target home page, from the **Exalytics System** menu, select **Patching**.

Configuring the Options for Oracle Exalytics Updates

Before patching an Oracle Exalytics target component, you can choose to set up the options for patching, such that the options are reused for subsequent patching operations. Setting up the patch options beforehand enables you to proceed with the analysis phase of a subsequent Oracle Exalytics patching operation without having to re-enter the options.

To set up the options for an Oracle Exalytics patching operation, follow these steps:

1. From the **Targets** menu, select **Exalytics**.
2. Click the name of the Oracle Exalytics target for which you want to set up the patch options.
3. On the Oracle Exalytics target home page, from the **Exalytics System** menu, select **Software Update**.
4. In the Component Navigation section, select either **Compute Node** or **BI Instance**, then click **Settings**.
 - **Operation Mode:** Select the mode of patching that you want to use. You can choose to patch the component targets in rolling mode or parallel mode.

Out of the component targets selected for patching, if you want to patch a single target at a time, select **Rolling**. This is the default option, and it involves very little downtime. While patching your component targets in rolling mode, you can choose to pause the execution of the patching deployment procedure after each node is patched. For information on how to do so, see [Pausing the Patching Operation While Patching Targets in Rolling Mode](#).

However, if you want to simultaneously patch all the component targets selected for patching, select **Parallel**. This option involves downtime, as all the component targets are shut down for a significant period. However, this process consumes less time, as all the component targets are patched simultaneously.

- **Stage Patches:** If you want the wizard to stage the patches from Software Library to a temporary location before the patch is applied on the component target, select **Yes** for **Stage Patches**, then specify the location where you want the patches to be staged. However, if you have already manually staged the patches that you want to apply, select **No** for **Stage Patches**, then specify the location where you staged the patches. If the stage location is a shared location, select **Shared Location**.

To manually stage the patches, download the patch, navigate to the location (parent directory) where you want to stage the patch, create a subdirectory with the same name as the patch zip file, then extract the contents of the patch zip file into this subdirectory.

- **Host Credential** (only for Compute Node): Provide the required credentials. You can choose to use preferred credentials, or override them with different credentials.
- **Oracle Home Credentials** (*only for BI Instance*): Provide the required credentials. You can choose to use preferred credentials, or override them with different credentials.
- **WebLogic Domain Credentials** (*only for BI Instance*): Provide the WebLogic Domain Credentials. You can choose to use preferred credentials, or override them with different credential.
- **Customization**(only for Compute Node): For patching Compute Nodes, by default, the wizard uses a static, Oracle-supplied deployment procedure to apply patches. Click **Create Like and Edit** to customize the default deployment procedure.

For patching Compute Node targets, you can select the following options:

- **Restart target after deployment:** Select this option to restart the Oracle Database Server target after applying the Quarterly Full Stack Download Patch (QFSDP).
- **Run script before deployment:** Select this option to run a script on the Oracle Database Server target before deploying the QFSDP. Specify the location where the script is hosted.
- **Run script after deployment:** Select this option to run a script on the Oracle Database Server target after deploying the QFSDP. Specify the location where the script is hosted.
- **Notifications:** Specify whether or not you want to enable email notifications when the patching operation is scheduled, starts, requires action, is suspended, succeeds, and fails.

To enable email notifications, select **Receive notification emails when the patching process**, then select the required options. If a warning message, mentioning that the sender or the receiver email address is not set up, is displayed, perform the action mentioned in the warning.

Updating Oracle Exalytics Compute Nodes

To patch the compute nodes of an Oracle Exalytics target, follow these steps:

 Note:

When you deploy a Patch Set Update (PSU) on an Oracle Exalytics compute node, the Oracle Integrated Lights Out Manager (ILOM) firmware installed on the compute node is also patched by default. You do not need to patch the ILOM firmware separately.

1. From the **Targets** menu, select **Exalytics**.
2. Click the name of the Oracle Exalytics target that you want to patch.
3. On the Oracle Exalytics target home page, from the **Exalytics System** menu, select **Software Update**.
4. In the Component Navigation section, select **Compute Node**.
5. Select the compute nodes that you want to patch, then select **Add Patch**.

The Select Patch window displays a list of Patch Set Updates that can be applied on the compute nodes. The recommended (latest) patches are marked with a tick. Select the Patch Set Update (PSU) that you want to apply, then click **Select**.

6. Click **Analyze**.

Specify a schedule for the analysis. Verify the provided options. If you haven't already specified the options (as described in [Configuring the Options for Oracle Exalytics Updates](#)), provide the required options .

Click **Submit**.

7. After the analysis is complete, click **Deploy**.

Updating Oracle Exalytics Business Intelligence Instance

To patch the BI Instance of an Oracle Exalytics target, follow these steps:

 Note:

When you deploy a Patch Set Update (PSU) on an Oracle Exalytics compute node, the Oracle Integrated Lights Out Manager (ILOM) firmware installed on the compute node is also patched by default. You do not need to patch the ILOM firmware separately.

1. From the **Targets** menu, select **Exalytics**.
2. Click the name of the Oracle Exalytics target that you want to patch.
3. On the Oracle Exalytics target home page, from the **Exalytics System** menu, select **Software Update**.
4. In the Component Navigation section, select **BI Instance**.
5. Select the BI Instance that you want to patch, then select **Select PSU**.

The Select Patch Set Updates window displays a list of Patch Set Updates that can be applied on the BI Instances. Select the Patch Set Update (PSU) that you want to apply, then click **Select**.

6. Click **Analyze**.

Specify a schedule for the analysis. Verify the provided options. If you haven't already specified the options (as described in [Configuring the Options for Oracle Exalytics Updates](#)), provide the required options.

Click **Submit**.

7. After the analysis is complete, click **Deploy**.

Part XI

Configuration, Compliance, and Change Management

This part contains the following chapters:

- [Managing Configuration Information](#)
- [Managing Compliance](#)
- [Managing Database Schema Changes](#)
- [Additional Setup for Real-time Monitoring](#)

Managing Configuration Information

This chapter explains how Oracle Enterprise Manager Cloud Control (Cloud Control) simplifies the monitoring and management of the deployments in your enterprise.

This chapter covers the following:

- [Overview of Configuration Management](#)
- [Overview of Configuration Searches](#)
- [Overview of Configuration Browser](#)
- [Overview of Configuration History](#)
- [Overview of Comparisons and Templates](#)
- [Overview of Configuration Extensions and Collections](#)
- [Overview of Parsers](#)
- [Overview of Relationships](#)
- [Overview of Configuration Topology Viewer](#)

Overview of Configuration Management

Cloud Control collects configuration information for all managed targets across the enterprise. Collected configuration information is periodically sent to the Management Repository over HTTP or HTTPS, allowing you to access up-to-date configuration information for your entire enterprise through Cloud Control.

Cloud Control enables you to view, save, track, compare, search, and customize collected configuration information for all managed targets known to Enterprise Manager. Additionally, the Configuration Topology Viewer provides a visual layout of a target's relationships with other targets; for example, you can determine a system's structure by viewing the members of a system and their interrelationships.

[Table 49-1](#) provides a snippet of configuration information collected for a small sampling of target types as an example.

Table 49-1 Collected Configurations for Various Targets

Target Type	Collected Configuration Information
Host ¹	• Hardware (includes memory, CPU, I/O device, and network information) • Operating system (includes installed patches and patch sets) • Oracle software (includes installed products and their components, patch sets, and interim patches applied using OPatch) • Other software (Operating system includes installed packages and patch sets) • Ksplice kernel patches and select userspace libraries

Table 49-1 (Cont.) Collected Configurations for Various Targets

Target Type	Collected Configuration Information
Database ²	- Database and instance properties - Initialization and System Global Area parameters - Tablespace, datafile, and control file information - Redo logs, rollback segments, and high availability information - Licensing information
Middleware such as WebLogic Server	- Node Manager, machine, Web service, and Web service port configurations - Resource Adapter, including outbound - Web and EJB modules - Server information - JDBC Datasource and Multi Datasource - Resource usage - Virtual hosts - Startup Shutdown classes - Jolt Connection Pool - Work Manager - JMS Topic, Queue and Connection Factory - Network channels
Elastic Cloud Infrastructure	- Switch details - Storage appliance details - Compute node details (including associated "Host" target GUID) - Switch ports configuration - Network topology (switch port - device association metric)
VM Server Pool	- Server Pool configuration details (total disk space and memory available, for example) - VM Guest member details
Client	- VM Server member details - Hardware - Operating system (includes properties, file systems, patches) - Software registered with the operating system - Network data (includes latency and bandwidth to the Web server) - Client-specific data that describes configuration for the browser used to access the client configuration collection applet - Other client-oriented data items
Non-Oracle Systems	- Hardware details including vendor, architecture, CPU, and I/O device information. - Operating system details including name, version, software and package lists, kernel parameters, and file system information. - OS Registered software including product name, vendor, location, and installation time.

¹ The default collection period for host configuration information is 24 hours.

² The default collection period for database configuration information is 12 hours.

Use Cloud Control to manage enterprise configurations:

- Search collected configuration data
- Compare configurations

- View latest and saved configurations as well as inventory and usage details
- Monitor configuration history for changes
- Build configuration extensions and introduce custom target types
- Collect and analyze external client configurations
- Perform root cause analysis and impact analysis

Overview of Configuration Searches

Use configuration search to search configuration data across the enterprise. Cloud Control ships with a set of configuration searches provided by Oracle, which you can use as a starting point to explore the volume of configuration data collected. As you work with a provided search, you can tailor the search criteria to refine or broaden the results, and save the altered search under a new name.

Perform powerful searches across the enterprise using sophisticated combinations of search filters, options, and relationships.

Enhance the search filtering criteria by adding your own SQL query statements. Save interesting search results by printing a report or exporting to a file.

To access the search capability, from the **Enterprise** menu, select **Configuration**, then select **Search**. The Configuration Search library page enables you to perform the following tasks:

- [Managing Configuration Searches](#)
- [Creating a Configuration Search](#)

Managing Configuration Searches

To manage configuration searches, access the Configuration Search Library from the **Enterprise** menu. Select **Configuration**, and then select **Search**.

On the Configuration Search Library page, you can perform the following tasks for existing configuration searches:

- [Searching for a Configuration Search](#)
- [Running a Configuration Search](#)
- [Editing a Configuration Search](#)
- [Deleting a Configuration Search](#)
- [Importing or Exporting a Configuration Search](#)

Searching for a Configuration Search

The Configuration Search Library page displays a table containing saved configuration searches.

To search for an existing configuration search, scroll down the table or use the Search option by doing the following:

1. Expand **Search**.
2. Specify the name of the configuration search.

3. Specify the owner of the configuration search.

 Note:

The search name and owner fields recognize containment, so you can specify a text string as a partial name to find all searches where the name or owner contains the string.

4. Select the target type.
5. Select **Latest** or **History** search type.
6. Select the mode that was used to create the search, such as **All**, **Modeler**, or **SQL**.
7. Specify if the configuration search is system defined or not.
8. Click **Go**. The search results are displayed in the table.

The View menu enables you to change how the content gets displayed in the table. You can select the columns that you want to view, sort the table content in ascending or descending order, and reorder the columns.

Running a Configuration Search

To run a configuration search, do the following:

1. On the Configuration Search library page, select a configuration search from the table, and then click **Run** to execute the search.
2. The Edit/Run Search page displays the search parameters applied in the search execution, the number of selected configuration items, and the search results.
3. You can edit the configuration items by clicking on the edit icon or the link next to it. The Apply Configuration Items dialog box appears. You can search for configuration items in the left panel, and then refine the configuration items using the right panel. Once you have selected the configuration items that you want to apply to the configuration search, click **Apply**. Else, click **Reset** to obtain the previous configuration items, or **Cancel**.
4. You can also export, print, and detach the configuration search page.

Editing a Configuration Search

You can only edit configuration searches that have an opened lock icon displayed next to the search name. Oracle saved configuration searches are usually locked searches which you cannot edit.

To edit an existing configuration search, do the following:

1. On the Configuration Search Library page, select a configuration search in the table and click **Edit**.
2. The Run/Edit Search page displays the search parameters and the results of the search.
3. Change the search criteria to achieve the desired results.

4. You can edit the configuration items by clicking on the edit icon or the link next to it. The Apply Configuration Items dialog box appears. You can search for configuration items in the left panel, and then refine the configuration items using the right panel. Once you have selected the configuration items that you want to apply to the configuration search, click **Apply**. Else, click **Reset** to obtain the previous configuration items, or **Cancel**.
5. Click **Save** to overwrite the existing search. Click **Save As** to save the edited search under a new name. If you are working with an Oracle-provided search, use Save As.
6. You can also export, print, or detach the configuration search.

Deleting a Configuration Search

To delete an existing search, on the Configuration Search Library page, select the configuration search from the table, and click **Delete**.

In the dialog that opens, confirm the operation. You must be either the owner or the administrator of the configuration search in order to delete a search. A search in use cannot be deleted. Oracle recommends that you not delete a configuration search provided by Oracle.

Importing or Exporting a Configuration Search

The Configuration Search Library page enables you to search for and import or export existing configuration searches.

To import a configuration search, click **Actions**, and then select **Import**. In the Import box that appears, browse for the configuration search file that you want to import, and click **OK**.

To export a configuration search, select a configuration search from the table, click **Actions**, and then select **Export**. A dialog box appears with the options to **Open** or **Save** the configuration search as an XML file.

Creating a Configuration Search

The Configuration Search Library page enables you to create a configuration search using any of the three usecases:

- [Creating a New Configuration Search](#)
- [Creating a Configuration Search from an Existing Configuration Search](#)
- [Creating a Configuration Search Using SQL](#)

Creating a New Configuration Search

To create a new configuration search, do the following:

1. From the **Enterprise** menu, select **Configuration**, and then select **Search**.
2. On the Configuration Search Library page, select **Create...**, and then select Configuration Search.
3. On the Configuration Search page, in the New Search section, select the target type of the configuration search. The table containing the targets gets refreshed to display the target type that you have selected.

4. To apply configuration items to the configuration search, do the following:
 - a. Click the Configuration Items **Add** link.
 - b. In the Apply Configuration Items dialog box, in the left panel, in the search field, specify the name of the configuration item. You can view a list of the configuration items in this panel in a flat or hierarchical view.
 - c. Select a configuration item from the left panel. You can refine the selected configuration item using the search criteria options in the right panel. Ensure that you deselect the configuration items in the right panel that you do not want to see in the search results.
 - d. Click **Advanced Search Options** if you want to further refine your search. This search method enables you to view the configuration items in groups, and add configuration items to each group. This search also provides OR and AND operators, unlike the Simple search which just gives you the option of OR operator.

You can further refine each group, by clicking the refine icon next to the group name. This enables you to select any of the following conditions:

None - Displays results based on specified property values.

Exists - Displays targets that contain the configuration item identified by the specified property values. For example, display database instances that contain patch 13343438.

Does not exist - Display targets that do not contain the configuration item identified by the specified property values. For example, display database instances that do not contain patch 13343438.

The first selection option returns not only matching entities but also actual property values. The rest return only the matching entities.

- e. Click on **Related Target Types** link to associate the target type with other targets. For example, you may want to know the Management Agent that is monitoring a host you have selected.
- f. Click **Apply** to add the configuration items to the configuration search. Click **Reset** if you want to revert to the previously saved configuration items.

The number of configuration items selected will now be displayed next to Configuration Items. The configuration items get added to New Search section and are displayed in the table against the targets.

5. As you add criteria, click **Search** to see the results. Continue to revise the search by adding and removing filters until the results are satisfactory.

Notice in the search results table that the column names are a concatenation of the filters you specify and elect to display. So, for example, If you filter on hardware vendor name for target type host, the column name in the search results table reads Host Hardware Vendor Name.

6. Click **Advanced** if you want to specify more specific search criteria such as target name, member of, and the host that the target is on. You can also choose to **Reset** all the changes that you have made.
7. Click **Save As**. In the Create Configuration Search dialog box, specify a name for the configuration search, and click **OK**.

Creating a Configuration Search from an Existing Configuration Search

To create a configuration search from an existing configuration search, do the following:

1. From the **Enterprise** menu, select **Configuration**, and then select **Search**.
2. On the Configuration Search Library page, search for and select the configuration search that you want to copy from the table.
3. Click **Create Like...**
4. In the Copy Configuration Search dialog box, specify a name for the new configuration search that you are creating.
5. Click **OK**.
6. Select the new row in the table and click **Edit**. Make the desired changes to the search parameters and then, save the configuration search.

Creating a Configuration Search Using SQL

Sometimes, despite all the filtering criteria, search results still fall short. To refine the search even further, you can create a configuration search using SQL by following these steps:

1. From the **Enterprise** menu, select **Configuration**, and then select **Search**.
2. On the Configuration Search Library page, click **Create** and then select **Search Using SQL**.
3. On the Search Using SQL page, you can create a SQL Query statement and then click **Search** to run the search. You can also edit the SQL Query statements for an existing configuration search by selecting the search from the table on the Configuration Search Library page and then clicking **Search Using SQL**.



Note:

You use views in this case. You cannot access the underlying tables. Your SQL edits apply only to the current search execution. If you want to preserve the edited statement, you can choose to **Export** as an XML file or **Print** the SQL statement.

4. Click **Save As**. In the Create Configuration Search dialog box, specify a name for the configuration search, and click **OK**.

Overview of Configuration Browser

Use the Configuration Browser to view configuration data in the context of a single managed entity. Configuration data can include:

- Configuration items and properties
- System configuration data as well as all system members and their configuration data

- System and target relationships (immediate, member of, uses, used by, and so forth)
- Configuration extension collection data

The browser window consists of left and right panes. The left pane is a tree hierarchy. The right pane consists of tabs that display information in tables. As you navigate in the tree, your selection dictates the contents in the right pane. Depending on the selection, tabs appear containing data such as properties and values, relationships, a hierarchical structure of a system and its members, and file contents in both a parsed and raw text format.

You can take any of several actions as you view a configuration in the browser. These actions are available from the **Actions** menu above the tabs. The tree hierarchy in the left pane also has context menus available.

This section covers the following topics:

- [Viewing Configuration Data](#)
- [Working with Saved Configurations](#)
- [Working with Inventory and Usage Details](#)

Viewing Configuration Data

The Configuration Browser enables you to view a target's latest or saved configuration data. While viewing configuration data, you can access configuration features such as compare and history.

1. From the **Targets** menu, select **All Targets**.
2. In the table of returned targets, right-click in the row of the desired target.
3. In the popup menu, select **Configuration**, then select **Last Collected** or **Saved**. In the case of saved configurations, select in the table of saved configurations the one you want to browse, then click **View**. The browser opens to display the (latest or saved) configuration data for the selected target.

Note that these same selections (**Last Collected** and **Saved**) are available in the **Configuration** menu on a target's home page that appears in the top-left corner and typically takes the name of the target type, for example, Host or Web Cache.

4. The browser display differs depending on the target type
 - For standard targets, the tree hierarchy on the left shows the target node at the top, beneath which appear configuration item categories and nested configuration items. Select the target node and the tabs on the right show target properties and various relationships (immediate, member of, uses, used by). Immediate relationships indicate direction: source and destination. Thus, for example, a source target type of database has an immediate relationship (hosted by) with a destination target type of host.

As you traverse the tree on the left, the tab on the right becomes the tree selection and displays the properties and values for the selection in table rows. So, for example, if the target type is host and you select Hardware in the tree on the left, the tab on the right becomes Hardware, and the table row displays values for Host Name, Domain, Vendor Name, and so forth. As the table view changes, look to the lower-right corner to see the number of rows the table contains. For multirow tables, use the search filter to drill down to specific properties and values. Add additional search filters as needed.

- When target type is a system, the tree hierarchy on the left shows the following:
 - The root target at the top level
 - A nested node one level down for each configuration item associated with the root target
 - A folder at the same level as the nested node for each member type
 - A node for each member within the member type beneath the member folder

Select the root target and the tabs on the right show target properties, a system topology table, and various relationships (immediate, member of, uses, used by). Select a configuration item in the tree on the left, and the tab on the right displays the item's properties and values. Note that this applies only to configuration items associated with the root target. Select a member target on the left and the tab on the right displays the member target properties. Note, however, that configuration data for the target does not display.

To see the member target's configuration data, you have to right click on the member, and then select **Latest Configuration**. The browser display then becomes the same as for a standard target. There is a bread crumb above the tree hierarchy on the left that enables you to return to the system view. If you subsequently save the member configuration, the link to the configuration data changes to **Saved Configuration**.

- Select a configuration extension file in the tree on the left; separate tabs for a parsed view and a raw text view of the file appear in the tables on the right.
5. To view the configuration details of all the members of the target, click **Configuration Report**. This exports all the configuration details into a zip file which gets downloaded. Extract the XLS file to view all the configuration details of the members.
 6. (Optional) If you want to save this configuration snapshot, select **Save Latest** in the **Actions** drop-down menu above the tabs. In the dialog that opens, enter a description by which to distinguish the configuration, then click **Submit Job**. Click **OK** to exit the dialog. The save action is also available on the right-click menu while selecting a target tree node. Saving a configuration saves all the configuration and relationship data for the selected target. It also saves the relationship and configuration data for all member targets.
 7. Other options in the **Actions** menu include:
 - **Go to Homepage**—returns to selected target home page.
 - **Export**—opens a dialog where you can browse to a file location and save the configuration as a CSV file.
 - **Topology**—opens the Configuration Topology Viewer showing the viewed target's relationships.
 - **Compare**—displays the comparison workflow page, where the viewed target's configuration is preselected as the one against which to compare other configurations.
 - **Search**—displays the configuration search page where the viewed target is the search object.
 - **History**—displays the history page for the viewed target's configuration.

- **Refresh**—triggers a collection of the viewed target's configuration data and subsequent refresh of the browser's tree hierarchy. Applicable only when viewing a latest configuration (last collected). Note that a manual refresh on a composite target applies only to the target itself, not to its members.

Working with Saved Configurations

Saved configurations are snapshots in time of collected data preserved for future reference. You may simply want to view the saved data, or you may want to use it as the basis of a comparison.

You can save standard as well as composite configurations. Saving a configuration saves all configuration item and relationship data for the selected target and for all member targets.

Note that there are various ways to save a configuration:

- While viewing a table of all targets, right-click a target and select **Configuration**, then select **Save**.
- While viewing a target's last collected configuration in the Configuration Browser, select **Save Latest** from the **Actions** drop-down menu.

A save, particularly one that involves systems or groups, can take several minutes. So, for performance reasons, a save action submits a job that occurs asynchronously. To check job status, do the following:

1. From the **Enterprise** menu, select **Job**, then select **Activity**.
2. Click **Advanced Search** and set the following criteria:
 - Set **Job Type** to ECM Save (or Save Latest).
 - Set **Target Type** to Targetless.
3. Click **Go**.
4. Drill down in the search results for save details.

To view a saved configuration:

1. From the **Enterprise** menu, select **Configuration**, then select **Saved**.
2. In the table of saved configurations, select the configuration you want to browse, then click **View**.
3. Navigate the tree hierarchy to expose the following categories of data:
 - Managed entities, configuration items, their properties, and relationships
 - System structures
 - Configuration extension collections

You can also view a saved configuration in the Configuration Browser: right-click a target tree node and select **Configuration**, then select **Saved**.

To compare a saved configuration:

1. From the **Enterprise** menu, select **Configuration**, then select **Saved**.
2. In the table of saved configurations, select the configuration you want to compare against, then click **Compare**.

3. The selected configuration becomes the first configuration in the comparison workflow. Continue the process of setting up the comparison.

To import a previously exported configuration:

1. From the **Enterprise** menu, select **Configuration**, then select **Saved**.
2. Click the **Import** button.
3. In the dialog that opens, browse to the location of the exported configuration data and click **Import**.

Upon refreshing, the imported configuration appears in the table of saved configurations.

Working with Inventory and Usage Details

In the Inventory and Usage Details page, you can:

- View inventory summaries for deployments such as hosts, database installations, and fusion middleware installations on an enterprise basis or for specific targets.
- View inventory summary information in the context of different dimensions. For example, for host inventory summary, you can view by platform, vendor, or OS version.
- Drill down multiple levels of inventory details.
- See trends in inventory counts charted across a time line. Chart bars are color-coded to match the view selection.
- Switch to a pie chart to break down the inventory data for the rollup option by color-coded percentages.
- For Hosts (OS Patches) and Databases (Patches Applied), click a patch indicator to link to patch details.
- Repeatedly revise selections to refresh chart and details based on new selections.
- Export deployment and details tables to CSV files.

To view inventory and usage details:

1. From the **Enterprise** menu, select **Configuration**, then select **Inventory and Usage Details**.

Alternatively, you can click **See Details** in the Inventory and Usage region of the Grid Summary page.

2. Select the entity you want to examine and choose a rollup option. For example, show all deployed hosts rolled up by platform. Note that the page refreshes automatically upon selection.
3. For patch updates, click **Yes** to view patch details.
4. Select the radio button to specify how to display the inventory chart.
 - The trend chart shows inventory counts across a time line. Use the magnifier icon to zoom the view. You can adjust the date range by sliding the horizontal scroll bar under the chart.
 - The pie chart breaks down the inventory data for the selected rollup option by percentages in an appealing color-coded visual.

5. Click **Table View** to convert the trend chart to table format. Close the table to return to the chart view.
6. Select one or more rows in the deployments table and click the **View Details** button to refresh the chart and details table based on the selected rows.
7. In any given row in the top table there is a count bar next to the count that represents a percentage of the maximum count. For example, if the maximum number of hosts by platform is four, the bar for hosts represented on two platforms would be half as long. Click the bar to refresh the details table and chart for the row.

Note that you can export either the master (deployments) table or the details table. In either case, click the **Export** button to open a dialog where you can browse to a file location and save the table as a CSV file.

**Note:**

In a shared Oracle Home deployment the inventory location is `/scratch/myCentralInventory`. This inventory location inside the Oracle Home is used by all the other hosts to execute the *attachHome* script.

The *oraInst.loc* file inside the Oracle Home points to this inventory location. Ensure that this inventory location is locally available in all destination hosts. Else, this location is created if required. This way there is coherence as all the hosts using the same Oracle Home also use the same inventory location.

Overview of Configuration History

Configuration history is a log of changes to a managed entity (target) recorded over a period of one year. The recorded history includes changes both to configurations and to relationships. Relationships are the associations that exist among managed entities.

Configuration history is a powerful tool for monitoring change activity across the enterprise. Consider these use cases:

- You have noticed that an Oracle RAC system has been underperforming for the past month. As an administrator it would be useful to know what changes have occurred during that time. Have members been added or removed? Have there been configuration changes to the system itself?
- The daytime administrator notices that detected changes are the result of a patch having been applied and adds an annotation to that effect. The overnight administrator is alerted to the changes and sees the annotation upon follow-up.
- A hardware memory change to production hosts has been detected. The administrator wants to keep the IT group posted on any future changes in this area. The administrator schedules a recurring job to check history specifically for changes to hardware memory on production hosts and to notify the IT group of such changes.

While viewing a configuration history, you can:

- Track changes to targets over time by specifying and refining search criteria.
- View change history and manipulate how the information is presented.

- Annotate change records with comments that become part of the change history. Annotations have a timestamp and an owner.
- Schedule a history search to capture future changes based on the same criteria.
- View the status of scheduled history jobs.
- Notify others of future change detection.
- Save change history details to a file.

This section covers the following topics:

- [Accessing Configuration History](#)
- [Working with Configuration History](#)
- [Viewing History Job Activity](#)

Accessing Configuration History

Use any of the following methods to access configuration history:

- From the **Enterprise** menu, select **Configuration**, then select **History**. Proceed with a configuration history search.
- Perform a search of all targets. Right-click in a row of returned targets and select **Configuration**, then select **History** in the popup menu. View the results for the selected target, identified by type and name in the respective search criteria fields; change the filtering criteria to see a different result. Select a specific configuration item, for example, or change the date range.
- On a target home page, select **Configuration**, then select **History** in the target type-specific menu (top left corner). View the results for the target, identified by type and name in the respective search criteria fields; change the filtering criteria to see a different result. Select a specific configuration item, for example, or change the date range.

Working with Configuration History

The Configuration History Perform the following tasks within configuration history:

- Drill down within configuration change history
- Enter annotations and comments
- Schedule a recurring history search and send the results
- Save a change history to a file
- Save configuration history
- Create a search using SQL

Searching History

To search configuration history, follow these steps:

1. From the Enterprise menu, select Configuration, and then select History.
2. In the New Search section, select the target type. The Include Member Target Changes check box is active only if you select a composite target type (system or group).

3. Select and specify the search criteria for the target name.
4. Click the Add link to add configuration items to the search.

In the Apply Configuration Items dialog box, you can search for configuration items in the left panel, and then refine the configuration items using the right panel. Once you have selected the configuration items that you want to apply to the configuration search, click **Apply**. Else, click **Reset** to obtain the previous configuration items, or **Cancel**.
5. Limit the scope of the search to a specific type of change, such as Change, Deleted Item, or New Item. All types of change is selected by default.
6. Specify the number of days for which you want the changes to be discovered. The default time is the last 7 days.
7. For a more refined search, click Advance, and do the following:
 - Add Relationship Items by clicking the Add link. Select a relationship type in the dialog box and click OK. This link is only enabled if a specific target type has been selected and if the Include Member Target Changes checkbox is not selected.
 - Choose whether to show all criteria-based history records or group them by timestamp and target. The default is Grouped.
 - You can refine the time and date range of the changes discovered by specifying the Before and After time ranges.
 - Specify search criteria for Annotation, On Host, and Member Of options.
8. Click **Search** to trigger the operation. A progress indicator verifies ongoing search activity. Results appear in the table at the bottom.

**Note:**

One search strategy to consider is perform a gross-level search to see the volume of changes, then go back and refine the search by adding filters.

Working with History Search Results

Each row represents a target satisfying the search criteria in which a change was detected, where a change constitutes something that was added, deleted, or modified. Numbers in parentheses on the tabs reflect the number of respective configuration and relationship changes detected. A search on target name for relationships returns matches on all source targets, destination targets, and targets that contain the target name.

- Click **See Real-Time Observations** to search actions monitored by compliance rules. Observations are the actions that users have taken on a host or target that were configured to be monitored through real-time monitoring rules.
- Click **Export** to save the search results to a CSV file such as a spreadsheet. The value in each column represents a comma-separated value.
- Click the number in the History Records column to display the changes detected for the selected target.

In the change details table, select a table row and do any of the following:

- Click **Details** to see the change details in a pop-up window, including old and new values, and the specifics of any annotations. The **Change** link in the Type of Change column pops up the same window.
- Click **See Real-Time Observations** to search actions monitored by compliance rules. Observations are the actions that users have taken on a host or target that were configured to be monitored through real-time monitoring rules.
- Click **Add Annotation** to enter a comment about the change.
- Click **Export** to save the search results to a CSV file such as a spreadsheet. The value in each column represents a comma-separated value.

Annotating Configuration Changes

To annotate configuration changes:

1. Select the change row in the results table. To add the same annotation to multiple lines, use the multiselect feature (Ctrl+click or Shift+click).
2. Click the **Add Annotation** button.
3. In the window that pops up, type your comment and click **OK**. Your comment appears in the designated column. Your login name and a timestamp are associated with your comment and available in the pop-up window that opens when you view the change details.

Note that you can also remove an annotation, provided you are the one who entered the comment (or have super administrator privileges). Select the row that contains the annotation and click the **Remove Annotation** button. Confirm the removal in the pop-up message that opens.

Scheduling a History Search and Creating a Notification List

You can schedule the change history search to run as a background job (click the **Schedule and Notify** button). The search can be run once-only or on a recurring basis. Run the search immediately or at some later date. You also can supply e-mail addresses to which to send a link to the search results.

Use a scheduled history search as a tracking mechanism to generate alerts when changes occur.

1. Specify the job schedule:
 - If not now, when. Click **Later** to activate the calendar widget where you can select a date and time.
 - How often. Select report frequency in the drop-down list. Default is once-only.
 - Wait how long. If the job fails to run as scheduled, cancel within a specified time frame.
 - Keep going. Maintain the job schedule for the specified period.
2. Enter the e-mail addresses of those to be directed to the change history search results. Use a comma to separate addresses.
3. Click **OK** to schedule the job.

Saving History to a File

You can capture the snapshot of change history you have culled for further review and to share with a wider audience, by saving the change details to a CSV file. Click **Export** and follow instructions in the export dialog.

Saving Configuration History

You can save your configuration history search, by clicking on the **Save As** button on the Configuration History page. In the Configuration History - Oracle Enterprise Manager dialog box that appears, specify a name for the saved configuration history.

Click **OK**.

Your configuration history search gets stored in the Configuration Search Library. To navigate to the Configuration Search Library, from the **Enterprise** menu, select **Configuration**, and then click **Search**.

You can select your saved configuration history and perform tasks such as edit, delete, and run.

Creating a Search Using SQL

Sometimes, despite all the filtering criteria, search results still fall short. To refine the search even further, on the Configuration History page, click **Search Using SQL**, and then select Configuration Changes or Relationship Changes.

On the Search Using SQL page, you can edit the SQL Query statement to extend search expressions and rerun the search. Note that you use views in this case; you cannot access the underlying tables.

Click **Search**.

You can also save the SQL search by clicking **Save As**. In the Create Configuration Search dialog box that appears, specify a name for the SQL search, and click **OK**.

Your SQL search gets stored in the Configuration Search Library. To navigate to the Configuration Search Library, from the **Enterprise** menu, select **Configuration**, and then click **Search**.

You can select your saved configuration history and perform tasks such as edit, delete, and run.

Viewing History Job Activity

View a list of all current and past history searches. Use search criteria to filter the list of history jobs (click the **History Job Activity** button). For example, show all scheduled history searches started over the past 24 hours; or, show all successful history searches involving hosts started over the past 31 days. The jobs engine purges history jobs older than 31 days.

The history jobs you can view beyond your own depend on your role and access level granted.

Select a table row and click **View Result** to go to the Jobs page that reports the history search. From there you can drill down to the changes the history search detected.

The job name is a hyperlink that takes you to the same place. Use the bread crumb on the Jobs page to navigate back to the list.

If you are the job owner or otherwise have the proper access level, you can perform list maintenance by deleting history jobs that no longer have relevance.

Overview of Comparisons and Templates

This section describes the template creation process and the use of rules in the process. It also provides information on setting up comparisons and managing comparison templates.

This section covers the following topics:

- [About Comparison Templates](#)
- [Working with Comparison Templates](#)
- [Specifying Rules](#)
- [About Rules Expression and Syntax](#)
- [Understanding Rules by Example](#)
- [About Comparisons](#)
- [Considerations Before Creating a Comparison](#)
- [Working with Comparison Results](#)

About Comparison Templates

A comparison template is an exemplar for fine-tuning a comparison of like configurations. A template is associated with a specific target type, which determines the configuration item types, items, and properties to be compared. Oracle provides a set of default templates to support certain target types. A template enables you to establish specific settings to take into account when comparing configurations of the given target type; for example, which property differences to ignore, and which property differences trigger an alert. You also can use constraints to establish acceptable values for specific properties. A configuration being compared that does not comply with the constraint constitutes a difference.

A template can invoke rules, or expressions, to be evaluated in determining when there is a match for comparison purposes, and when to disregard differences detected in a comparison.

Templates can be used as is, or as a guideline. So, for example, you might decide that an existing comparison template, with just a few tweaks, can meet your requirements. Perhaps the template ignores property differences that you are concerned about. In this case, use the create-like feature to make the adjustments to an existing template and save it under another name.

For systems, you design a system template that references member templates, based on the target types that make up the system. Create the member templates before you create the system template.

Working with Comparison Templates

This section describes how to create, edit, and otherwise manage comparison templates.

Creating or Editing a Comparison Template

Use these instructions when creating a new template or editing an existing template; this includes create-like.

1. From the **Enterprise** menu, select **Configuration**, then select **Comparison & Drift Management**. Click the Templates tab.
2. To search for a template, click **Search**. You can search for multiple target types from the target type list. Select the target types that you want to search for. You can also specify the Template Name, the name of the Owner, and specify if the template is a default template or an Oracle provided template. Click **Search**.
3. Each template has a lock beside the template name. A closed lock represents an Oracle provided template. These templates cannot be edited. The templates that have an open lock are user defined comparison templates, and can be edited.
4. For a new template, click **Create** and provide a name and target type. To base a template on an existing one, select the template row, click **Create Like**, and provide a name. In either case, the action creates a new template row.
5. Select the appropriate template row in the table and click the **Edit** button. The **Template Details** page appears.

The compared configurations' target type drives the hierarchy of configuration item types and configuration items on the left. The settings in play for the respective properties on the right derive from the selected template, unless you are creating a new template from scratch, in which case there are no settings.

A system comparison takes an overall template and a template for each system member. Thus there is an additional tab for **Member Settings**. Edit the tab as follows:

- Optionally select the member template to use for each system member type.
 - For any given member type, you can elect to compare configurations by checking the check box.
 - For member types that you are comparing, select a target property to use as a matching key. The default is target name, but typically you would want to use a distinguishable property to align comparison entities, such as department or location.
6. To create or edit the comparison template for each member, select the **Member Settings** tab.

You can choose to view the mapping display in a tree or table format. If you set the Mapping Display to **Tree**, the View mapping and Comparison results will display the system members in a hierarchical tree format. If you set the Mapping Display to **Table**, the View mapping and Comparison results will display the system members in a table format. You can edit the following:

- Optionally select the member template to use for each system member type.

- For any given member type, you can elect not to compare configurations by clearing the check box.

 Note:

When you clear a check box for a system member, the children instances of the system member will automatically be ignored during the comparison that uses this template.

- For member types that you are comparing, select a target property to use as a matching key. The default is target name, but typically you would want to use a distinguishable property to align comparison entities, such as department or location.
7. In the **Template Settings** tab, select a configuration item type or item in the left pane to expose its properties in the right pane. A key icon denotes a property that is defined as a key column in the configuration item type's metadata.

 Tip:

Notice the Compare check box column on the **Template Settings** tab. This is a powerful feature that enables you to streamline the comparison by selecting only those items you want to compare. When you select the check box, the comparison engine includes the corresponding configuration item type and all of its descendants.

Contrast this with the ability to compare individual columns and rows on the **Property Settings** tab, in which the settings are stored as part of comparison results, giving you the option to view the compared properties on the results page.

So, for example, in comparing host configurations, you may decide that any differences in CPU properties are immaterial. Simply expand the Hardware configuration item type and deselect the CPUs check box to exclude all properties associated with the item.

8. Click the **Property Settings** tab and check boxes for property differences to be compared and alerted if different. They are mutually exclusive. When you compare differences in a property value in this fashion, you are doing so unconditionally for all differences detected in the property value for the configuration item type.

Use a value constraint rule to filter the property value. In this case, the comparison engine compares the property value in the configurations being compared (the second through n configurations) to the constrained value. A property value that satisfies the constraint constitutes a difference. For example, test for a version equal to or greater than 6. Any instance in the compared configurations of a version property value under 6 constitutes a difference. Clearly, you would not set a value constraint if you checked compare differences. Specify a rule expression to set a value constraint. See [Specifying Rules](#) for details.

9. Repeat the preceding steps to set additional property settings on other configuration items.
10. Optionally, select an item in the left pane and click the **Rules for Matching Configuration Items** tab. For a given property, specify a rule expression to be evaluat-

ed to determine when a match exists between configuration instances. In other words, if the expression resolves to true, compare the instances. See [Specifying Rules](#) for details.

Match rules are column-based; they apply an AND logical operator. If you specify rules for multiple properties, they must all resolve to true to constitute a match.

11. Optionally, select an item in the left pane and click the **Rules for Including or Excluding Configuration Items** tab. For a given property, specify a rule expression to be evaluated:

- Compare all - Compares everything in the configuration item.
- Excludes items that match these rules - Compares everything except the properties listed.
- Includes items that match these rules - Only compares the properties listed, that is, ignores everything else.

The rules for including or excluding configuration items are row-based; they apply an AND logical operator within a subset of rules and an OR logical operator between rule subsets. So, if you specify two rules for property A and two rules for property B, either both rules set on property A OR both rules set on property B must resolve to true to constitute a match.

Managing Comparison Templates

In addition to creating and editing comparison templates, you manage them by doing the following:

- View a template's settings and composition; this is read-only
- Delete a template (requires the proper permissions)
- Share templates by exporting them in XML file format and importing them into other Cloud Control systems

Viewing a Comparison Template

You can view templates provided by Oracle and other users' templates to which you have access. Viewing a template is read-only: you see its makeup, but you cannot change anything, even temporarily.

1. Select a template in the Comparison Templates page and click the **View** button.
2. Expand items in the tree on the left and peruse the settings and rules on the various tabs.

Deleting a Comparison Template

Deleting a template is subject to the following constraints:

- You cannot delete an Oracle provided template. An Oracle provided template is represented by a closed lock beside the template name.
- You cannot delete a comparison template unless you have the proper permissions.
- You cannot delete a default comparison template.
- You cannot delete a comparison template currently in use.

To delete a template, select it in the Comparison Templates page, click **Delete**, and confirm the operation.

Exporting a Comparison Template

Use the export feature to save a template as an external file that can be imported into another Cloud Control system.

1. Select a template in the Comparison Templates page, select the **Actions** menu and click **Export**.

A platform-specific file dialog opens. For example, if you are using Firefox, the dialog notes that you have chosen to open the named template, which it identifies as an XML file. The dialog asks what you want Firefox to do, open the file in an XML editor or save the file.

2. Select the save radio button and click **OK**.
3. Browse to the desired location in the file system and save the file, changing the name if applicable. You cannot change the name of a template provided by Oracle on export.

Importing a Comparison Template

Any comparison template import must comply with the comparison template .xsd. So, for all intents and purposes, the import should be a previously exported template to ensure compliance.

1. In the Comparison Templates page, select the **Actions** menu, and click **Import**.
2. Browse to the template file location and click **Import**.

The imported template appears as a new row in the template table.

An exported template is associated with its owner. A template whose owner is not the same as the login ID of the person importing the template retains its original ownership. If you want to be the owner of the imported template, you have to edit the owner attribute in the template XML file prior to import, changing the value to your login ID. Or, you can simply remove the attribute, in which case the default owner will be set to the ID of the person initiating the import operation.

The Template Manager disallows import of a template provided by Oracle of the same name. Similarly, you could change the name attribute in the template XML file prior to import to allow the import to occur.

Specifying Rules

Specify rules in the context of creating or editing a comparison template (see [Creating or Editing a Comparison Template](#)).

Rules enable you to parse configuration data in order to fine-tune comparisons. In terms of the comparison, a rule applies the expression to the value of the selected item in the configuration instance that is being compared to the benchmark configuration. Matching rules are intended to devise a comparison key that aligns the instances being compared. Ignore rules are intended to establish a basis for disregarding any differences detected between instances being compared.

To access the rules functionality, select **Configuration** from the **Enterprise** menu, then select **Comparison & Drift Management**. Click the **Templates** tab on the left. Select a template to edit and click **Edit**.

Creating a Value Constraint Rule

Specify value constraint rules as follows:

1. Select a configuration item in the left pane.
2. Click the **Property Settings** tab in the right pane and select the property on which you want to set a value constraint.

When the Property Settings tab is selected, keys are displayed in the column to the left of the Property Name.

3. Click the **Edit Rule** button in the toolbar. In the dialog that opens:
 - a. Select an operator from the drop-down list.
 - b. Type an operands expression, then click **OK**. An operand is a value that you want to either include or exclude from the constraint. For example, if you want to exclude Patch ID 12,34,56,78, you would enter operand as '12', '34', '56', '78'.

To clear a rule, select the table row and click the **Remove Rule** button in the toolbar.

See [About Rules Expression and Syntax](#) for details on the formation of a rules expression.

Creating a Matching Rule

Specify matching rules as follows:

1. Select a configuration item in the left pane.
2. Click the **Rules for Matching Configuration Items** tab in the right pane, then click **New**.
3. Select a property in the drop-down list that appears under **Property Name**.
4. To create the rule, select the table row and click the **Edit Rule** button in the toolbar. In the dialog that opens:
 - a. Select an operator from the drop-down list.
 - b. Type an operands expression, then click **OK**.
 - c. To specify additional rules, click **New** and repeat Steps a and b

To clear a rule, select the table row and click the **Remove Rule** button in the toolbar.

For more information on the formation of a rule expressions, see [About Rules Expression and Syntax](#).

You can enter additional rules for the same or for a different configuration item. When there are multiple rules, they resolve in the order specified. Matching rules take an AND logical operator, which means all conditions must resolve to true to constitute a match.

Creating a Rule for Including and Excluding Configuration Items

Specify ignore rules as follows:

1. Select a configuration item in the left pane.
2. Click the **Rules for Including or Excluding Configuration Items** tab in the right pane.
3. Choose one of the following options: Compare all, Exclude those that satisfy rules, Include only those that satisfy rules. Click **New**.
4. Select a property in the drop-down list that appears under **Property Name**.
5. To create the rule, select the table row and click the **Edit Rules** button in the tool-bar. In the dialog that opens:
 - a. Select an operator from the drop-down list.
 - b. Type an operands expression, then click **OK**.
 - c. To specify additional rules, click **New** and repeat Steps a and b.

To clear a rule, select the table row and click the **Remove Rule** button in the tool-bar.

See [About Rules Expression and Syntax](#) for details on the formation of a rules expression.

You can enter additional rules for the same or for a different configuration item. When there are multiple rules, they resolve in the order specified. Including and Excluding rules take an AND logical operator for rules within a subset, and an OR logical operator between subsets. So, for two subsets, each with multiple rules, all rules in the first subset OR all rules in the second subset must resolve to true to constitute a match.

6. Select **New Or** to indicate the end of one rule subset and the beginning of another.

About Rules Expression and Syntax

A rule consists of an operator and operands. Taken together, they form an expression that resolves to a value that is then compared to the value of the selected item. A true condition satisfies the rule.

Operands can be literals (string literals are enclosed in single quotes), legal numbers, or dates of the form YYYY-MM-DD HH24:MI:SS.FF. Operands that directly reference the value of a configuration item must be of the same date type as that value. Operands in square brackets in the syntax are optional.

Operator	Operands
is equal to*	An optional literal value to match; string values are case-sensitive; if unspecified, expression evaluates value of the property to which the rule applies Note that a matching rule compares the values of the configuration items in the respective configuration to one another, not to a third specified value, so the operator does not take an operand in this case. [match-literal]

Operator	Operands
is case-insensitive equal to*	An optional case-insensitive string literal; if unspecified, expression evaluates value of the property to which the rule applies Note that a matching rule compares the values of the configuration items in the respective configuration to one another, not to a third specified value, so the operator does not take an operand in this case. <code>['match-literal']</code>
is greater than or equal†	A literal value to match; required <code>match-literal</code>
is greater than †	A literal value to match; required <code>match-literal</code>
is less than or equal to†	A literal value to match; required <code>match-literal</code>
is less than†	A literal value to match; required <code>match-literal</code>
is one of†	A comma-separated list of literal values, at least one of which must be specified, but only one of which need match <code>match-literal-1[,match-literal-n,...]</code>
is between†	A range specified as start and end literal values; both must be specified; range is inclusive <code>start-range-literal , end-range-literal</code>
contains†	A string literal on which to perform pattern matching; required <code>[FALSE TRUE,] 'pattern-literal'</code>
replace‡	FALSE (default) means string must comply with Oracle LIKE operator syntax; TRUE means string must comply with Posix regular expression syntax A string literal to match and replace with a second string literal <code>[FALSE TRUE,]'pattern-literal'[, 'replacement-literal'][,position-integer][,occurrence-integer]</code> FALSE (default) means string must comply with Oracle LIKE operator syntax; TRUE means string must comply with Posix regular expression syntax TRUE enables optional positional integer argument to indicate where within the column value to extract the string, and optional occurrence integer argument to indicate the position count to replace Mandatory pattern literal represents the string value to match If the replacement string literal is unspecified, replace the matched string literal with nothing

Operator	Operands
substring‡	Extract specified segment of string value [FALSE TRUE,]position-integer[, length-integer][, 'pattern-literal'[, occurrence-integer]] FALSE (default) means string must comply with Oracle LIKE operator syntax; TRUE means string must comply with Posix regular expression syntax Mandatory positional integer argument indicates where to begin string extraction: • If 0 or 1, returns all characters • If positive integer, starts extraction from beginning • If negative integer, starts extraction counting backwards from end Optional length integer argument indicates character count starting at position integer pattern literal represents the value to match; optional if the first argument is FALSE; required if TRUE occurrence integer argument indicates character count to match; valid only if pattern literal is specified

Notations are as follows:

- *—Enabled for value constraints, matching rules, and ignore rules
- †—Enabled for value constraints and ignore rules only
- ‡—Enabled for matching rules only

Understanding Rules by Example

These rule examples assume that you are in the process of creating or editing a template and are at the point where you have selected the configuration item in the tree on the left.

Matching Rule Examples

Suppose, when comparing the hardware of host configurations, you want, for matching purposes, to ignore case in respective vendor names. Here's a simple rule to make the comparison case-insensitive.

1. In the **Rules for Matching Configuration Items** tab, click **New**.

Note: For this example, ensure you are using a Host target type template.
2. Select **Vendor Name** in the drop-down list under Property Name.
3. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to is-case-insensitive-equal-to. As this operator takes no operands for a matching rule, you are done.
 - Click **OK**.

Suppose you want to compare WebLogic Servers, aligning on server name, where the names are different: ManagedServer1 and ManagedServer2, for example. To ensure the comparison occurs, you need to fashion a match on server name.

1. In the Template Settings tab, highlight **Server Information**.
2. In the **Rules for Matching Configuration Items** tab, click **New**. In the Property Name drop-down list, select **Machine Name**.
3. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to substring.
 - Set **Operands** to 1, 13.
 - Click **OK**.

Effectively, the rule says use the first 13 characters of the name (ManagedServer), thus excluding the qualifying integer.

4. Another way to achieve the same result:
 - Set **Operator** to replace.
 - Set **Operands** to true, '(*)(\d*)', '\1'.
 - Click **OK**.

This example uses a regular expression (TRUE) to resolve all characters prior to the qualifying integer.

For a more advanced example, consider a database instance comparison that requires a match on Datafiles file names within a Tablespace, where file names are of the form:

```
/u01/jblack_abc2d/oracle/dbs/dabc2/mgmt_ad4j.dbf
```

1. In the Template Settings tab, highlight the **Control files** configuration item.
Note: For this example, ensure you are using a Database Instance target type template.
2. In the **Rules for Matching Configuration Items** tab, click **New**.
3. In the Property Name drop-down list, select **File Name**.
4. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to replace.
 - Set **Operands** to true, '(/u01/)(.*)(oracle.*\dabc[0-9]+.*)(.*)', '\2\4'.
 - Click **OK**.

Effectively, the rule says use a regular expression (TRUE) to construct a matching key from the value between /u01/ and oracle, combined with what remains of the original filename after dabc2 /, or jblack_abc2d/mgmt_ad4j.dbf.

Ignore Rule Examples

Generally, you use ignore rules to ignore differences in collections that are row-oriented, as opposed to column-oriented. Configuration extension snapshots, for example, are row-oriented data collections.

Say, for example, you wanted to ignore in configuration extension parsed data, any row where the property `Attribute` identifies an internal ID or checksum.

1. In the **Rules for Including or Excluding Configuration Items** tab, click **New**.
2. Select **Attribute** in the drop-down list.
3. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to is one of.
 - Set **Operands** to 'id', 'checksum'.
 - Click **OK**.

The rule ensures that the comparison ignores any row in the collection data that contains either of the specified values.

Now consider an ignore rule that demonstrates how the comparison engine applies the logical operators AND and OR against the same configuration item type. In this example the objective is to ignore rows in configuration extension parsed data when any of three rule sets satisfies the following conditions:

```
Data Source = 'sqlnet.ora' AND Attribute = 'ADR_BASE'
OR
Data Source = 'tnsnames.ora' AND Attribute = 'HOST'
OR
Data Source = 'resources.xml' AND Attribute = 'authMechanismPreference'
```

Notice that the comparison engine applies the AND operator to rules within a set and the OR operator between rule sets. Rules for ignoring instances support inheritance; thus, in this case, the Data Source property is available in rules creation, as demonstrated in the example.

1. In the **Rules for Including or Excluding Configuration Items** tab, click **New**.
2. Select **Data Source** in the drop-down list.
3. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to is equal to.
 - Set **Operands** to 'sqlnet.ora'.
 - Click **OK**.
4. Click **New** and select **Attribute** in the drop-down list.
5. Select the table row and click the **Edit Rules** button in the toolbar to open the rule dialog.
 - Set **Operator** to is equal to.
 - Set **Operands** to 'ADR_BASE'.
 - Click **OK**.
6. Click **New Or** to insert a logical OR operator to signal the end of the first rule set.
7. Add two new rules where **Data Source** is equal to 'tnsnames.ora' and **Attribute** is equal to 'HOST'.
8. Click **New Or** to insert a logical OR operator to signal the end of the second rule set.

9. Add two new rules where **Data Source** is equal to 'resources.xml' and **Attribute** is equal to 'authMechanismPreference'.

The comparison ignores any row in the collection data that satisfies any of the three rule sets.

About Comparisons

Enterprise Configuration Management deals with the collection, storage, and monitoring of configuration data tied to managed entities within the enterprise. A host, for example, has configuration item types related to its hardware and software components—number of CPUs, memory, IO devices, OS platform and version, installed software products, and so forth.

Changes to configuration data invariably happen, typically because of common events like patches and upgrades. At some point a change to one component can affect the overall system in a negative way. Detecting the root cause becomes paramount.

Enterprise Manager provides the following types of comparisons:

- Configuration drift
Enables you to compare configurations of a target with configurations of another target of the same type.
- Configuration consistency
Reflects the changes of target members within a system. For example, you would use configuration consistency to ensure that the configuration parameters for all databases within a cluster database are the same

The comparisons can be done on the current configuration or on configurations previously saved (for example, just before applying a patch or doing an upgrade).

Comparisons allow you to:

- Ignore certain attributes during a comparison. Define comparison templates to disregard unnecessary attributes.
- Notify key personnel when differences are detected
- Design and share comparison templates with other administrators
- Compare complete target systems; match target system members automatically or manually
- Compare configuration file data as raw file content or in a parsed format

Note that after you define and save comparisons, a change in configuration *automatically* starts a recomparison. If you have setup notifications, you will be notified of the changes.

Considerations Before Creating a Comparison

Comparisons are an important factor in managing the enterprise. Setting up a comparison involves the following steps:

- Select the first configuration in the comparison (the one to compare against)
- Select additional configurations (the one or more configurations to compare to the first configuration)

- Select a comparison template to fine-tune the attributes being compared (or no template)
- For system comparisons, map members as needed. It's a way to selectively indicate how members of respective systems should match up in a comparison.
- Review your work
- Set up email notification by setting up mail servers and then creating incident rules
- Review BI Publisher Reports (From the **Enterprise** menu, select **Reports**, then select **BI Publisher Reports**.)

A follow-on step would be to review the results and drill down to differences details.

Steps in Setting Up a Drift or Consistency Comparison

When creating a drift or consistency comparison:

1. Determine what configuration items you want to compare.
2. Create a template or use an existing template.
3. Create a definition of either a [drift comparison](#) or a [consistency comparison](#).
4. Perform the comparison.
5. Set up email notification by setting up mail servers and then creating incident rules. See [Creating Notifications for Comparisons](#).
6. View the results from the Comparison & Drift Management Drift Results tab or the Comparison & Drift Management Consistency Results tab. You can also view the results as a BI Report. To view BI Reports: from the **Enterprise** menu, select **Reports**, then select **BI Publisher Reports**.

About One-Time Comparisons

One-time comparisons are used to immediately view the differences between target configurations. One-time comparisons differ from drift and consistency comparisons in that the comparison is only evaluated once, even if you choose to save the results for later viewing.

If you are thinking of creating a drift or consistency comparison, consider creating a one-time comparison to verify that the initial drift or consistency evaluation will not result in immediate differences. Using the results, you can change the errant configuration values, or fine-tune your comparison template to ensure that only configuration items of interest are compared.

To perform a one-time comparison:

1. From the **Enterprise** menu, select **Configuration**, then select **Comparison & Drift Management**.
2. In the One-Time Comparison section of the Dashboard page, click **Create Comparison**.
3. Select Basic or Advanced comparison.
 - The Basic One-Time Comparison is a simple comparison between two or more targets. The fields are:
 - Reference Target (Current)

Target against which the comparison is being made.

- Comparison Template

Template or pattern to be used for the comparison. This template can contain property settings, rules for matching configuration items, and rules for including and excluding configuration items.

- The Advanced One-Time Comparison provides more options than the Basic One-Time Comparison. For example, you can use saved configurations and perform a consistency comparison within a system.

The fields are:

- Reference Target (Current)

Target against which the comparison is being made. Use this option when comparing targets. Note that the list of the targets is restricted to 2,000. Use the Search option to restrict the list of targets.

- Reference Target (Saved)

Configuration that was created at an earlier time and which will be used as the base for the comparison. Use this option when you want to compare targets against your gold configuration.

- Consistency Target Type (use for systems)

When comparing systems, provide the type of system, for example, Cluster Database or Database System.

Use this option to test the consistency of targets among systems. Note that all cluster members have the same saved configuration. For example, you can determine if all the configuration database parameters are the same within a cluster database.

- Comparison Template

Select a template that has gone through rigorous testing or use this one-time comparison to fine tune the comparison template to compare only what you need.

An example of an existing template is one that you created or a template provided by Oracle. If you don't supply a template, there will be a one-to-one comparison between the fields in the reference target and the compared targets.

4. Click **Add** to select targets or **Add Saved** to select saved configurations for the comparison. Remember that you can only compare targets of the same target type.

Note that the more targets you add, the longer it will take for the compare operation to complete.

After you have added the targets, and you decide to minimize the number of targets, on the Target menu, highlight the targets to eliminate and click Remove.

5. Click **OK**. The comparison begins immediately and results are displayed on the Comparison Results page.

Depending on the options you chose, it may take a while for the results to display. Click the Enterprise Manager Refresh button until the comparison is completed and the In Progress icon disappears. Click the i icon located by Comparison Results name for a listing of the options chosen for the comparison.

About Configuration Drift

Configuration drift ensures consistency (uniformity) across a large number of targets. For example, ensures that the configuration of the personnel database is the same across all the personnel databases throughout the company. This is particularly useful in this day and age of company acquisitions.

To perform a Drift Management Comparison:

1. From the **Enterprise** menu, select **Configuration**, then select **Comparison & Drift Management**.
2. In the Drift Management section of the Dashboard page, click **Create Definition**.
3. On the Compare Configurations dialog box, select the Target Type and Template. For example, select Database Instance and the template of choice, in this case Database Instance Template. Then, click **OK**.

On the Drift Definition page, provide the following information:

- **Definition Name** - Make the name meaningful. The information from the Compare Configurations popup is replicated.
- **Source Configuration**

Select either Latest Configuration or Saved Configuration. Either one is the gold configuration against which the other targets will be compared. When using the Latest Configuration, choose the source target. When using the Saved Configuration, choose the appropriate configuration.

- **Advanced**

Expand this section to provide additional factors on the comparison.

- Choose the severity of the drift. Options are: Minor Warning, Warning, and Critical.
- Target Property Filter

You can specify specific target properties that determine which targets this definition can work against. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this definition, for instance for Linux OS, it will only be applicable to targets on Linux Operating Systems.

- Description

Describe this drift definition and provide as much detail as possible. This information provides information for others who will be using this definition in the future.

- Rationale

Explain the reason for this comparison, for example: This content will detect configuration drifts for the targets.

- Keywords

Enables you to categorize this drift definition for quick reference.

After you have provided the information, select one of these options:

- Save and Test Targets

Once a target is associated with a comparison, comparisons are automatically triggered for that target when appropriate. For example, when the target's configuration changes, a comparison will be triggered, and if the automatic comparison results in a difference from the reference target, notifications will be sent, and the differences will be available in any reports for the drift comparison.

Before associating targets with a comparison, you may want to run a test comparison. Configuration differences detected during a test comparison do not cause notifications to be sent, and will not appear in reports. The results from the test comparison can be used to verify that the comparison template is only comparing items of interest. Also, the test results can give a preliminary indication of how different a target is from the reference target. The **Save and Test Targets** option lets you run these test comparisons.

To test targets, in the Test Targets page, click **Add** to add a compared target. Select the target, and click **Run Test**. This will trigger a test comparison, and the number of differences generated will be placed in the Differences column of the table. To view the results of the test comparison, click on the difference count. Once you are satisfied with the comparison results, select the target and click **Associate** to permanently associate the target with the drift comparison, enabling automatic re-comparisons, notifications, and visibility of the target in reports.

- **Save and Associate Targets**

Saves the input, triggers automatic comparisons that use this comparison definition, and directs you to the Target Association for Drift Definition page.

Click **Add** and select the targets you want to associate with this definition. Click **OK**. You will be prompted on whether you really want to save the associations.

- **Save and Return**

Saves the input for future use. You can view the saved comparison definition in the Definition Library tab of the Comparison & Drift Management dashboard.

You are returned to the page from where the configuration drift was launched. For example, if you started from Drift Definition Library page, you will return to that page. If you started from the Overview tab, you will return to the Overview tab. If you started from the Drift Results page, you will return to the Drift Results page.

- **Cancel**

Aborts the operation. None of the input will be saved.

About Configuration Consistency

Configuration consistency reflects the changes of target members within a system or group. "For example, you would use configuration consistency to ensure that all of the database instances within a cluster database have the same configuration.

To perform a Consistency Management Comparison:

1. From the **Enterprise** menu, select **Configuration**, then select **Comparison & Drift Management**.
2. Locate the Consistency Management section of the Dashboard Overview page. Click **Create Definition**.

Note: You can create comparison definitions from several locations. You can create comparison definitions from:

- Overview tab

- Definition Library tab
 - Comparison Templates tab, after selecting a template
3. On the popup, select the Target Type and Template. For example, select Cluster Database and the template of choice. The template can be one that you have defined, or as in this case, Cluster Database Template. Then click **OK**.
 4. On the Consistency Definition Details page, change the Definition Name as needed. The Compare Template and Applicable To information are replicated.
 5. Click **Advanced** to provide information which is used for compliance.
 - Compliance Standard Name is the same name as used for the Consistency Definition Name. This is the name to search for when using compliance standards.
 - Rule Name is the same name as the compliance standard name.
 - Compliance Rule State is automatically considered Production.
 - Choose the severity of the drift. Options are: Minor Warning, Warning, and Critical. When a consistency is in a Critical state, it needs to be addressed in a timely manner. For example, if the space on a database is getting very low, it needs to be addressed before it crashes.
 - Target Property Filter

You can specify specific target properties that determine which targets this definition can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this definition, for instance for Linux OS, it will only be applicable to targets on Linux Operating Systems.
 - Description

Describe this consistency definition and provide as much detail as possible. This information provides information for others who will be using this definition in the future.
 - Rationale

Explain the reason for this comparison, for example: This content will detect configuration consistency for the systems.
 - Keywords

Enables you to categorize this consistency definition for quick reference.
 6. For consistency comparisons, Oracle chooses one target of each member target type as the reference target. All other members will be compared against the reference target of the same target type.

Note: Since all members of the same type should be the same, it should not matter which target is selected as the reference target. However, if you would prefer to choose specific targets as reference targets, click the **Edit** icon in the Reference Targets when associating targets, or click the **Reference Targets** button when creating a one-time comparison. This allows you to choose your own reference targets for each member target type.

After you have provided the information, select one of these options:

 - Save and Test Targets

Saves the comparison definition which compares the members against reference members within the system. You can then run test comparisons to verify that the comparison template only compares attributes of interest to you.

To review the results, click the number in the Differences column. The Comparison Results page appears.

- Save and Associate Targets

Saves the input and directs you to the Target Association for Consistency Definition page. Click **Add** and select the systems you want to associate with this definition. Click **OK**.

Once the targets are associated with a comparison, comparisons are automatically triggered for the targets when appropriate. For example, when a target's configuration changes, a comparison will be triggered, and if the automatic comparison results in a difference from the reference target, notifications will be sent, and the differences will be seen in any reports for the consistency comparison.

- Save and Return

Saves the input for future use. You will be returned to the Comparison Dashboard.

- Cancel

Aborts the operation. None of the input will be saved.

About the Definition Library

The Definition Library is the repository for all the drift and consistency definitions created using the Comparison & Drift Management Dashboard.

To access the Definition Library, select **Comparison & Drift Management** from the Enterprise menu. On the Dashboard page, click the **Definition Library**.

From this page, you can:

- Create a new definition
- Edit an existing definition (as needed)
- Run tests against a definition
- Associate targets or groups to a definition (perform the association *after* you have verified test results from the Test Association page)
- Delete a definition

Setting Up a Comparison Template

To set up a comparison template:

1. From the Enterprise menu, select **Configuration**, then select **Comparison & Drift Management**.
2. On the left of the Dashboard page, click the Templates icon.
3. On the Comparison Templates page, click **Create** to create a new template or **Create Like** to create a template similar to one that already exists. Using the Create Like option enables you to use existing templates and make minor changes as needed.

Provide data for the Name and Target Type. When providing a name, make it obvious so folks can determine the reason for the template. For historical purposes, provide a description that explains the reasons and specifics of this template. When you provide the target type, configuration items for that target are automatically added to the template definition.

4. After you create (or create like) the template, edit the template to delete or modify configuration items.

When the Save Only Differences option is checked, only differences will be saved when this template is used in drift, consistency, and one time comparisons. If the box is not checked, then all information will be stored in the comparisons in which this template is a part.

Note: When creating a system template, for example, a cluster database, the template page provides more information when there is a target with members like a Cluster Database.

5. Compare Configurations: This option provides the same functionality as if you were on the Dashboard page and clicked Create Comparison for One-Time Comparison, Create Definition for Drift Management, or Create Definition for Consistency Management.
6. After you associate targets to the comparison definition, the comparison will rerun automatically whenever there is a change to the configuration for a target, when the system members change, or when the template changes.
7. If you choose Table for the Mapping Display, and do not override the default flat map when defining the comparison, the members of the systems will be matched for comparison without regard to their level in the system hierarchy. Thus, when viewing the results, the original system hierarchy cannot be retrieved.

Chose the Table option only if you are not concerned about how the members of the system are related to each other.

Creating Notifications for Comparisons

To be notified of a comparison change, you need to setup and enable notifications. There are two parts to setting up notification: setting up the mail servers and creating an incident rule.

Follow these steps to setup notifications:

1. From the Setup menu, select **Notifications**, then select **Mail Servers**.
2. On the Mail Servers page, provide the Server Identity information then add the outgoing mail (SMTP) server.

After you have set up the Mail Servers, set up the Incident Rules as follows:

1. From the Setup menu, select **Incidents** then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click **Create Rule Set**.

Provide a name and description and apply the rule set to Targets. This should apply to All Targets.

3. Create a rule as follows:
 - a. Select - **Incoming events and updates to events**.
 - b. Select - **Specific events of type Compliance Standard Rule Violation**.

- c. Select either **Compliance Standard** or **Compliance Standard Rule**. If selecting a Compliance Standard, select a standard of type Drift or Consistency. If selecting a Compliance Standard rule, select a rule of type Drift or Consistency.
 - d. Add a rule - If you selected a compliance standard in the previous step, add a compliance standard of type **Configuration Drift** or **Configuration Consistency**. If you selected a compliance rule, select either a **Configuration Drift Rule** or a **Configuration Consistency Rule**.
 - e. Add Actions - **Basic Notifications** (E-mail to). Enter the e-mail addresses of those who are to be notified when the comparison detects a difference. Use a comma to separate addresses. Remember that the properties for which differences are alerted were specifically selected in the comparison template.
 - f. Click **Next** and specify rule name and description. Click **Continue**.
4. Click **Save**.

Working with Comparison Results

This section covers comparison results from the following perspectives:

- [About Consistency Management \(System\) Comparison Results](#)
- [About Drift \(Target\) Comparison Results](#)
- [Synchronizing Configuration Extension Files](#)

About Consistency Management (System) Comparison Results

Consistency results display when you click the **Consistency Results** tab on the Comparison & Drift Management Dashboard page. The view summarizes the results of all the consistency comparisons. Red signifies the number of inconsistent targets within the system, whereas green signifies the number of consistent targets within the system. To view the differences for a particular consistency definition, click the number of differences associated with the consistency definition. The Comparison Results page appears.

- Select **Only Differences** in the Show drop-down list to eliminate the "noise" of the same results.

The icons that appear in the view are mostly intuitive: equal–same, not equal–different.

The table displays a hierarchy of system and member target types where:

- The Target Type column displays the system and member tree hierarchy.
- The Result column shows comparison results based on the mappings established as part of comparison setup. A boxed 1 (left only) or 2 (right only) means there was nothing to compare to the first or second member target, respectively. Note that if the parent target configurations are the same, but one or the other parent has child members marked as left only or right only, the parents are marked as different.
- To resolve unmatched members, rerun the comparison, this time ensuring in the mapping step that the left and right member pairs appear in the mapped members table. Select an appropriate system comparison template with target matching rules defined, such that these members are mapped, or map the pairs manually.

- When the Member column displays both an equal and a not equal icon, it indicates equality at the parent level, but a difference in some member.
- To view a summary of all the differences found when comparing the system target and any member targets, click **Export**, located at the top of the table that displays the system members. An XLS report will be downloaded.

About Drift (Target) Comparison Results

Drift results display when you click the **Drift Results** tab on the Comparison & Drift Management Dashboard page. The view summarizes the results of all the drift comparisons. Red signifies how many targets have drifted from the gold standard, whereas green signifies how many targets are similar to the gold standard. To view the differences for a particular drift definition, click the number of differences associated with the drift definition. The Comparison Results page appears.

Simple Target (Non-System Target) Results

When a simple target (non-system target) comparison is completed, the left pane displays a hierarchy of configuration items for the target being compared, and, if applicable, configuration extensions. Refine the scope comparison results as follows:

- Select **Only Differences** in the Show drop-down list to eliminate the "noise" of the same results.
- Select **Left Only** to display items that are only present on the target displayed on the left and NOT present on the target displayed on the right.
- Select **Right Only** to display items that are only present on the target displayed on the right and NOT present on the target displayed on the left.

The icons that appear in the view are mostly intuitive: equal–same, not equal–different. The key icon denotes the key properties of the configuration item type. An indication of Out of Range means that the property value failed a value constraint set on the property. A boxed 1 (left only) or 2 (right only) means that the comparison did not find a matching item to compare to the first or second configuration, respectively.

System Drift Comparison Results

When a system drift comparison is completed, the system results page displays the system and its members along with its comparison results. Drill down from the system results to the simple target results to view additional configuration comparison result details.

Synchronizing Configuration Extension Files

Use this feature to perform on-demand file synchronization when a comparison of file-based configurations returns differences. Often, this involves configuration extensions that users create. See [Overview of Configuration Extensions and Collections](#), for information on configuration extensions.

Note:

This feature is available only for file-based configuration extensions. Differences resulting from comparisons of command-based or SQL query-based configuration extensions cannot be synchronized.

1. From the Enterprise menu, select **Configuration**, then select **Comparison & Drift Management**. Click the **Drift Results** tab on the left. On the Drift Results page, click the Drift Definition of interest.

On the Comparison Results page, locate the Configuration Items (Differences) region located on the left. Select all the Configuration Extension items of interest in the Configuration Tree. You can select multiple files, indicating that you'll be updating all of them in the same direction.

Click the **Synchronize** icon located to the right of the Configuration Item. This icon is present *only* for configuration extension (CE) nodes, and only for CE nodes that are eligible for synchronization.

Note: File synchronization is also available from the results of one-time comparisons.

2. The Synchronize File page displays the files selected on the Comparison Results page. If there are files that cannot be synchronized, such as those that have no differences, they are not submitted for synchronization.
3. Optionally, use the Preview feature to view the effect of the update on a file-by-file basis. Click the eyeglasses icon to view the file before and after the update in raw format.
4. Complete the Credentials and Setting sections as follows:
 - Specify the login credentials as necessary. You must have login access to the target destination and write permission on the directory or directories to be updated.
 - Select the appropriate radio button to indicate a destination directory. In either case (original or alternate), you must have write permission on the directory.
 - Select the appropriate radio button for how to proceed on conflict. The comparison is performed using data from the repository. A conflict arises when the file to be updated has changed on the target and is different from the data used for the comparison. Indicate what you want to do in this case—proceed or stop.
 - Note that irrespective of the selection for destination directory (original or alternate), the conflict check is always performed against files in the original directory.
 - Indicate the desired backup options (both are selected by default when the update target is the original directory):
 - Mark the appropriate check box if you want to save a snapshot of the configuration to be updated prior to synchronizing (give it a descriptive name so you can easily retrieve the file from saved configurations; defaults to a generic name—CCS Synchronization Saved Snapshot—which applies even if you blank the field).
 - Mark the appropriate check box if you want to make a backup copy of the configuration file before it's updated. Browse to a directory on which you have write permission.

These are not mutually exclusive options. With the former, you are saving time-stamped collection data in the OMS repository; whereas, with the latter, you are storing a copy of a file in a file system.

- If desired, perform an on-demand collection refresh of the destination target's configuration data immediately after file synchronization. This way, if you rerun the comparison or view the configuration in the Configuration Browser, the ef-

facts of the update will be visible. You can also run a manual refresh at any time, or wait for the next scheduled collection.

The check box is selected by default when the original destination directory is the update target. The check box is disabled if you specified an alternate directory, as there would be nothing to refresh in this case.

- When satisfied with the results, click **Synchronize**.

On the Synchronize Files popup, click the link to track the synchronization job. When the job completes, you can rerun the comparison to verify the update, assuming you requested a refresh. You can also open the configuration extension in the Configuration Browser and confirm the update there.

Not All Configuration Files Can Be Synchronized

You may notice in the comparison results differences view that some files, though different, cannot be selected for synchronization (their check boxes are disabled). There are several possible reasons, including:

- The destination file is non-writable.
- There is no source file.
- Files that do not have differences.
- During the configuration extension definition, the file was associated with a parser that does not support a process called reverse transform, which is, effectively, the ability to return the parsed form of a file to a syntax tree structure that can then be rendered back into a physical representation. Not all parsers support reverse transform.

Note: It is on the File Synchronization page where the files are marked as eligible or ineligible for synchronization. It is on this page where you can determine whether the selections are valid.

Comparison and Drift Management BI Publisher Reports

Business Intelligence Publisher (BI Publisher) is the primary reporting system that provides a single, Web-based platform for authoring, managing, and delivering interactive reports and all types of highly formatted documents.

Using BI Publishes, Enterprise Manager provides the following Comparison and Drift Management reports:

- Drift Report for Systems - Drift results report for system targets. It includes Drift definition summary, various roll-ups and Drift Comparison results. Examples of system targets are databases and fusion applications.
- Drift Report - Drift results report. It includes Drift definition summary, various roll-ups and Drift Comparison results. Use this report to view simple targets, for example, host.
- Patch Differences for Fusion Instances (Scheduled)

Note: This is a Bursting/Scheduler only report. Preconditions include:

- Comparison template for Fusion Instance and its corresponding Oracle Home target type is available.

- Oracle Home comparison template contains the 'Patches Installed In Oracle Home' configuration item, and has the 'Patches ID' and 'Patch Language' properties enabled.
- Comparison was performed using the proper template.
- Patch Differences for Fusion Instances
Preconditions include:
 - Comparison template for Fusion Instance and its corresponding Oracle Home target type is available.
 - Oracle Home comparison template contains the 'Patches Installed In Oracle Home' configuration item, and has the 'Patches ID' and 'Patch Language' configuration properties enabled.
 - Drift definition was created using the proper Fusion Instance template.

To access the Comparison and Drift Management BI Publisher reports:

1. From the **Enterprise** menu, select **Reports**, then select **BI Publisher Enterprise Reports**.
2. In the list of Enterprise Manager Reports, select **Comparison and Drift Management**.
3. In the login screen, provide your credentials to access the reports.

NOTE:

- To see information in these reports, you must have had to run a comparison
- Only differences are reported
- Information is grouped by compared targets

Overview of Configuration Extensions and Collections

Configuration extensions provide a way to identify files and other configuration data that Cloud Control does not already collect. These customized configurations can be collected on well-known target types or on a target type introduced as part of the configuration extension definition. A set of configuration extensions called blueprints is available for download from Oracle. They are called blueprints because they lay out precisely the files and data to collect for a given platform such as Apache Tomcat.

A typical life cycle of a configuration extension might be as follows:

- Create a configuration extension and deploy it to several targets.
- Assess its effectiveness over time.
- Modify and fine-tune the specification and redeploy, perhaps across a wider spectrum.
- Undeploy and delete the specification if no longer pertinent.

This section covers the following topics:

- [Working with Configuration Extensions](#)
- [About Configuration Extensions and Deployment](#)
- [Extending Configuration Data Collections](#)

- [Using Configuration Extensions as Blueprints](#)

Working with Configuration Extensions

This section describes how to create, edit, and otherwise manage configuration extensions. If you want to create a configuration extension that uses a custom target type, the suggested workflow is to first create the custom target type. Concurrent with that activity, you can also add a complementary new target to serve as a sample target instance.

Creating a Custom Target Type

If no existing target type satisfies your configuration extension requirements, you can create a custom target type.

Before creating a new target type, ensure that the administrator has installed the Software Library (from the **Setup** menu, select **Provisioning and Patching**, then select **Software Library**). This must be done once, after Cloud Control installation.

1. From the **Enterprise** menu, select **Configuration**, then select **Configuration Extensions**. On the Configuration Extensions page, select **Create Custom Target Type** from the **Actions** menu.
2. In the dialog that opens, provide a name for the custom target type, then click **OK**. As noted, it may take a while to complete the process.
3. When done, a message confirms target type creation and asks if you want to add a sample target instance. A sample target provides the basis for collecting configuration data. Click **Yes**.
4. A dialog opens associated with the custom target type you just added. Click the search icon to select a Management Agent to monitor the target you are adding, then click **Add Target**.
5. In the dialog that opens, provide target properties appropriate to the instance target type. In particular, the pertinent target property is the path to install home, as this is the likely location of configuration files relevant to the custom target type. Optionally, provide global properties such as cost center and lifecycle status. Click **OK**.

The target is now available as a sample target when you create a configuration extension for the custom target type.

It's not imperative that you add a new target instance during custom target type creation. You can do so subsequently by selecting **Add New Custom Target** from the **Actions** menu and following Steps 4 and 5 in the process above, this time selecting a custom target type from the drop-down list.

Creating or Editing a Configuration Extension

Use the following instructions to create, create like, or edit a configuration extension.

Given appropriate privileges, you can edit a configuration extension and save the edited version, in which case, the version number increases. You might also edit and save as a draft, or edit a draft for publishing. Note that when you edit a configuration extension, you cannot change the target type, as this would cause the underlying metadata to be incompatible with existing deployments of the configuration extension, see [About Configuration Extensions and Privileges](#).

**Note:**

When you edit a deployed configuration extension, it is automatically redeployed upon saving. This does not apply to saving as draft.

1. In the Configuration Extensions library, click the **Create** button; or, select an existing specification in the library and click **Create Like** or **Edit**.
2. On the Create Configuration Extension page, enter a name for the configuration extension and an optional description. The create like action requires minimally that you rename the specification.
3. Select a target type from the drop-down menu.
4. Optionally, set up a sample target. A sample target resides on the host from which you intend to collect configuration data. If you do not set up a sample target, you cannot browse the file system or use the preview feature in entering your specifications.

Click the search icon. A dialog opens containing known instances of the target type. Use the filtering criteria as necessary to locate the instance you want, then click **Select**.

5. For instructions on how to complete the Files & Commands tab, see [Using the Files & Commands Tab](#)
6. For instructions on how to complete the SQL tab, see [Using the SQL Tab](#).
7. After you complete the specification definition and have mapped credentials to the target type, use the preview feature to validate your entries, in particular, to ensure the parsed view is what you expect.
8. Save the new or edited specification. Remember that configuration extensions are in the public domain. Use the save-as-draft feature to keep the specification private while you test and refine it, see [About Configuration Extensions and Versioning](#).

If you are editing a draft, the buttons change as follows:

- **Publish** implies that you are making the draft public.
- **Save** implies that you are creating the next version of the draft.

When done, you can begin collecting configuration data by deploying the configuration extension to target instances, see [About Configuration Extensions and Deployment](#).

Using the Files & Commands Tab

Create file and command specifications as follows:

1. Click the search icon to browse to a default base directory location. This is where the configuration files reside, or where the commands you specify are to execute.

Click the **Use Property** button to open a dialog where you can select a target property to include as part of the directory path. These properties serve as variables, denoted by curly braces, to be substituted with actual values at runtime. You can type additional text in the box to supplement your selection. So, for example, you might select OracleHome and append a directory—`{OracleHome}/config`—to collect files on the target located in the config subdirectory under the Oracle Home

path. Note that the target type definition determines available target properties. User-defined properties do not appear in the list, as they are not available at the Management Agent.

2. Click **Advanced Settings** to specify the following:

- An alternate base directory for the sample target.
- The encoding to use in collecting the data at the Management Agent. Configuration data is stored in UTF-8 format in the repository. Oracle Default means use UTF-8 for XML files and the locale encoding of the target for all other file types; Target Locale means store all file types including XML in the locale encoding of the target; otherwise, select an encoding from the drop-down list. Selecting directly from the list automatically selects the accompanying radio button.
- Whether to use the Management Agent credentials (file and command specification only) or some other predefined credential set to access data on the target. If the customized credential set does not appear in the drop-down list, click **Create** to identify the credential set to use. Note that you must then specify the credentials that map to the credential set name you create. If you don't know a mapped name, you can specify a credential set when you open the Remote File Browser to add files as described in Step 3. See [Setting Up Credentials When Creating a Configuration Extension](#) for more information.

3. Click **Add** and select file or command as the specification type.

For a **file specification**, enter a file name in the space provided or browse the base directory to select a file on the target. Use of wildcards (* and **) is allowed, where ** indicates 0 or more subdirectories. In using wildcards (and as a general caveat), ensure that collections do not result in too many (or too large) files, and that the files collected be configuration-related, that is, files under administrative control that change relatively rarely, so as not to overload Cloud Control.

For a **command specification**, enter command syntax in the space provided or browse the base directory to a script. You must assign a unique alias to the command. The alias you assign appears in the Configuration Browser as a link when viewing the configuration extension hierarchy. When you click the link, it opens the command specification in the tab on the right. The same caveats as mentioned for files apply to command output; that is, that their results are constrained in number and size, and to configuration-related data.

Select a parser to convert the configuration file or command output into a standard format for storing in the repository. There is no default. If you do not specify a parser, only the raw data format is stored and available for viewing. See [Managing Parsers](#) for more information.

Optionally, specify post-parser rules to align tree nodes. See [Setting Up Rules](#) for information on entering rules.

4. Repeat Step 3 to specify additional files or commands.

Return to [Creating or Editing a Configuration Extension](#) and resume with 7.

Using the SQL Tab

Create SQL query specifications as follows:

1. Select credentials to use to connect to the database. If the customized credential set does not appear in the drop-down list, click **Create** to identify the credential set to use. Note that you must then specify the credentials that map to the credential

set name you create (see [Setting Up Credentials When Creating a Configuration Extension](#)). Configuration extensions only support database credentials with NORMAL roles, not those with SYSDBA, SYSOPER, or other roles.

2. Specify a JDBC connection to an Oracle database from which to extract data via an SQL query. The connection string can be either a URL or an abstraction of database target properties. It cannot be a combination of the two; that is, partial URL and some target properties.

The URL must contain the name of the target database host, applicable port number, and the Oracle Service name (SID); for example, `mydatabase.example.com:1521:ORCL`.

If you want to use target properties, leave the field blank. At runtime the application will substitute values for these target properties— `{MachineName}{Port}{SID}`—to make the connection.

3. Click **Add** and type or paste a SQL query in the provided text box. Ensure that the query is sufficiently selective to return only pertinent configuration-related data of manageable size and scope.

You must assign a unique alias to the query. The alias you assign appears in the Configuration Browser as a link when viewing the configuration extension hierarchy. When you click the link, it opens the SQL query in the tab on the right.

Database Query Parser should be preselected in the drop-down list.

Optionally, specify post-parser rules to align tree nodes. See [Setting Up Rules](#) for information on entering rules.

4. Repeat Step 3 to specify additional SQL queries.

Return to [Creating or Editing a Configuration Extension](#) and resume with Step 7.

Setting Up Credentials When Creating a Configuration Extension

If you create a credential set while creating a configuration extension, you have to specify the credentials that make up the credential set. To do this, you have to return to the Configuration Extensions library and proceed as follows:

1. From the **Setup** menu (top right of the page next to the Help menu), select **Security**, then select **Monitoring Credentials**.
2. Select the applicable target type in the table and click **Manage Monitoring Credentials**.
3. Select the row with the credential set name you created during the configuration extension definition for the given target type and click **Set Credentials**.
4. Enter the username and password for the credential set and click **Save** (or **Test and Save** for database credentials).
5. Return to the Files & Commands tab ([Using the Files & Commands Tab](#)) or SQL tab ([Using the SQL Tab](#)) description.

Setting Up Rules

Use rules to differentiate nodes in the parsed representation that have the same name. This is particularly important in comparisons and change history when trying to match nodes in the parsed tree, or when expressing SQL queries to verify compliance. Rules resolve to an identifier that is appended in square brackets to node text in the

tree as a way of uniquely identifying the node. An operation such as a comparison will then use the combination of node text and bracketed identifier for evaluation purposes.

A rule consists of a condition and an expression, both of which must be valid XPath expressions. The condition resolves to a node that requires the identifier. The expression resolves to a string computation for the identifier. You can use a special case SKIP expression to bypass the node specified in the condition; this is a convenient way to eliminate "noise." In other words, for purposes of comparison, ignore the node the condition resolves to.

Some parsers have default parser rules already defined. They execute automatically on the parsed representation. You can elect to use a subset of default rules, edit them, or override them with custom rules that you define.

The number in the **Rules** column is significant. Initially, the number is zero (0). A whole number greater than zero indicates the number of custom rules defined. Zero also appears for a parser that has default parser rules. So the appearance of a whole number in the column stipulates an override of default parser rules, if any, with the custom rules the number represents.

Set up rules as follows:

1. Click the **Parser Rules** button. The Edit Parser Rules page displays.
2. To define a custom rule, click **Add**. In the table row that appears, enter a condition and an expression as valid XPath expressions.

You can define multiple rules; they are applied to the parsed content in the order specified. Click **Return** when you are done.

Select a table row to delete a custom rule.

3. To manipulate default rules, click **Add Default Rules**.

Rules appear in table rows, provided the parser you selected has default parser rules. Edit and delete default rules as appropriate to your purposes. Remember that you are working with a copy of these rules; the originals remain safely intact.

Note that if you delete all rules, you are merely removing the copies you imported. Default parser rules will still fire unless overridden by custom rules.

For examples of rules, see [Using Parsed Files and Rules](#).

Return to the Files & Commands tab ([Using the Files & Commands Tab](#)) or SQL tab ([Using the SQL Tab](#)) description.

Managing Configuration Extensions

In addition to creating and editing configuration extensions, you manage them by doing the following:

- View the selected specification (read-only)
- Synchronize the selected specification with facets in the Compliance Library for real-time facet monitoring
- Share configuration extensions by exporting them in XML file format and importing them from the local file system
- Delete the selected specification (requires the proper permissions)

Viewing a Configuration Extension

You can view a configuration extension in read-only mode to get an idea of the make-up of a specification. Perhaps, for example, to see if it is a likely candidate on which to base a new specification.

1. In the Configuration Extensions library, select the specification table row and click **View Details**.
2. Peruse the settings and rules on the various tabs.

Enabling Facet Synchronization

You can synchronize a configuration extension specification with real-time monitoring facets to monitor real-time changes to the configuration files and queries that make up the configuration extension. Real-time monitoring enables you to know such things as when files and database settings change, who made the change, whether observations were automatically reconciled, whether the actions observed were authorized, and so forth.

When you synchronize configuration extensions with real-time monitoring facets, future changes to configuration extensions automatically propagate to corresponding facets, which means configurations are not only collected, compared, tracked, and so forth, but also are monitored for authorized real-time changes. Note that to associate a configuration extension with a facet and to subsequently edit a configuration extension synchronized with a facet requires the additional role of EM_COMPLIANCE_DESIGNER.

1. In the Configuration Extensions library, select the specification table row, then select **Enable Facet Synchronization** from the **Actions** menu.
2. The **Facet Synchronization** column displays a **Use Facet** link in the configuration extension table row. Click the link to go to the **Real-time Monitoring Facets** tab in the Compliance Library where you can manage the synchronization of facets with the configuration extension.

Exporting a Configuration Extension

You can export a configuration extension as an XML file that can subsequently be imported into the same or another system.

1. In the Configuration Extensions library, select the specification table row, then select **Export** from the **Actions** menu.
2. Browse to a file system location where you want to save the specification as an XML file. The saved file takes the name of the configuration extension by default.

Importing a Configuration Extension

Given appropriate privileges, you can import a configuration extension that was previously exported as an XML file.

1. In the Configuration Extensions library, select the specification table row, then select **Import** from the **Actions** menu.
2. Browse to the file location. Select the file and click the **Import** button on the dialog. The imported specification appears in the Configuration Extensions library.

Deleting a Configuration Extension

You must be the owner or otherwise have sufficient privileges to delete a configuration extension. Note that there are dependencies that potentially impact deletion, including deployments, job schedules, existing collections, and so forth.

1. In the Configuration Extensions library, select the specification table row and click **Delete**.
2. The system validates permissions and otherwise checks for dependencies that might prevent the deletion, although some dependencies cannot be verified until a job submission involving the configuration extension.

About Configuration Extensions and Versioning

When you create a configuration extension, you have options to save or save as draft. A normal save action makes the specification publicly available to the general user community. A save as draft action keeps the specification private to you. How you use these actions when creating and editing specifications influences the mechanics of versioning. Consider the following scenarios:

- You create and save a configuration extension; this is public version 1. You subsequently edit public1 and save as a draft; this becomes draft1. Public1 is still generally available. You edit draft1 and publish; this becomes public2. Note that in parallel, someone else with the proper permissions can also edit public1 and save as a draft to create version 1 of draft2.
- You create and save a configuration extension as a draft; this is version1 of draft1. You edit and save again; this becomes version 2 of draft1. Repeat the edit-and-save operation; this becomes version 3 of draft1. Edit version 3 of draft1 and publish; this becomes public version 1.

About Configuration Extensions and Privileges

Working with configuration extensions requires privileges specific to the given operation you want to perform.

Operation	Required Privilege (Role)
Create new target type	EM_PLUGIN_OMS_ADMIN To create a new target type, ensure that the administrator has installed a software library (from the Setup menu, select Provisioning and Patching , then select Software Library). This must be done once, after Cloud Control installation.
Create new target instance	EM_PLUGIN_AGENT_ADMIN
Create or import configuration extension	"Manage configuration extensions owned by user" (or the more powerful "Manage configuration extensions owned by any user")
Associate configuration extension with an auto-synchronized real-time monitoring facet	EM_COMPLIANCE_DESIGNER

Operation	Required Privilege (Role)
Edit or delete configuration extension	Differs, depending on the specific activity within the realm of editing: • Configuration extension owner requires "Manage configuration extensions owned by user"; nonowner requires "Manage configuration extensions owned by any user" • Schedule redeployment jobs for already deployed targets requires "Create" privilege for Job System resource type • For configuration extensions associated with real-time monitoring facet, EM_COMPLIANCE_DESIGNER
Deploy or undeploy configuration extension on a target	"Manage target metrics" privilege on the target instance; "Create" privilege for Job System resource type (to schedule deployment/undeployment;) EM_PLUGIN_AGENT_ADMIN (to deploy a plug-in to a Management Agent)
Create a new credential set	Superuser
View configuration extension definition	None
View configuration extension collected data	Regular "target instance view" privilege

Note that editing an imported configuration extension may be restricted to edits that do not change the version, depending on options set during export. One such permissible edit would be to credential set information.

About Configuration Extensions and Deployment

Deployment of a configuration extension means to direct the specification to a target where a monitoring Management Agent will collect configuration data based on the specification's definition. A configuration extension can be deployed to multiple targets. You must have sufficient privileges to deploy and undeploy configuration extensions.

Manage deployments by performing the following actions:

- [Deploying and Undeploying Configuration Extensions](#)
- [Editing a Deployment of Configuration Extensions](#)
- [Viewing a Configuration Collection](#)

Deploying and Undeploying Configuration Extensions

Deployment of a configuration extension means to direct the specification to a target where a monitoring Management Agent will collect configuration data based on the specification's definition. A configuration extension can be deployed to multiple targets. You must have sufficient privileges to deploy and undeploy configuration extensions.

To deploy a configuration extension:

1. In the Configuration Extensions library, select the specification table row and click **Manage Deployments**.
2. On the Deployments page, click **Add**. In the dialog that opens, search for and select targets of the specified target type where you want to deploy the configuration extension.
3. When you close the dialog (click **Select**), a new column appears denoting a pending action of **Deploy** and the status becomes **Selected for Deployment**.

4. Proceed as follows:

- Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Deployment job in progress**.
- Click **OK** to schedule the deployment and return to the library.
- Click **Cancel** to void the request and return to the library.

5. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

Usually, if you update a deployed configuration extension (CE), redeployment occurs automatically. However, redeployment will not be initiated when certain CE attributes are modified, such as the sample target.

To undeploy a configuration extension:

1. On the Deployments page, select the deployment in the table.

2. Click **Remove**. A new column appears denoting a pending action of **Undeploy**; status remains **Deployed**.

3. Proceed as follows:

- Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Undeployment job in progress**.
- Click **OK** to schedule the undeployment and return to the library.
- Click **Cancel** to void the request and return to the library.

4. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

When viewing configuration extensions in the library, a green check mark in the Deployments column denotes currently deployed configuration extensions. The number in the column indicates how many targets the configuration extension has been deployed to. Click the number to navigate to the relevant deployments page.

Editing a Deployment of Configuration Extensions

To edit a deployment, follow these steps:

1. In the Configuration Extensions library, locate the appropriate table row and click the numerical link in the Deployments column. In addition, you can, after selecting the configuration extension table row, click the **Manage Deployments** button in the toolbar.
2. On the Deployments page, select the deployment in the table and click **Edit**.
3. The type of configuration extension, that is, file/command-based or SQL-based, determines the make-up of the dialog that opens. Specify a base directory to override the default base directory currently in effect, or change the JDBC URL, as appropriate. Click **OK**.
4. After closing the Edit dialog, proceed as follows:
 - Click **Apply** to confirm the action while remaining on the Deployments page. The action column disappears, and the status becomes **Redeployment job in progress**.
 - Click **Save** to initiate the redeployment and navigate back to the Configuration Extension library page

- Click **Cancel** to void the request and return to the library.
5. Click **Refresh Status** on the Deployments page to confirm a successful outcome.

Note that the edit applies to the deployment of the specification; it does not change the configuration extension definition.

Viewing a Configuration Collection

You must have sufficient privileges to view a configuration extension's collected data.

1. In the Configuration Extensions library, locate the appropriate table row and click the deployments link.
2. On the Deployments page, select the deployment in the table and click **View Configuration**.
3. In the Configuration Browser popup window, peruse details of the configuration extension by selecting nodes in the tree hierarchy on the left:
 - The root node represents the target instance being monitored. The right pane displays target properties and immediate relationships.
 - The next level down in the tree represents a template for the specification. The right pane displays specification details such as configurations being collected and the base directory from which they are collected.
 - The remaining leaf nodes in the tree represent the configuration data collected. The right pane displays the configuration data in both parsed and raw format.

You can also view the collected data from the target home page: from the target type menu, select **Configuration**, then select **Last Collected**.

Extending Configuration Data Collections

There are two options available to extend configuration data collections using a configuration extension specification:

- Add additional collection items to an existing target type
- Add a custom target type with new collection items

Extending Existing Target Collections

The following instructions describe how to extend the configuration data Cloud Control collects for an existing target type. The listener target type, for example, does not collect the `sqlnet.ora` file as provided by Oracle. To extend the listener data collection to include this item, take the following steps:

1. From the **Enterprise** menu, select **Configuration**, then select **Configuration Extensions**.
2. In the Configuration Extensions library, click the **Create** button.
3. Give the configuration extension an appropriate name and select **Listener** as the target type.
4. Click **Select Target** to choose a listener instance that is already deployed, so you can browse to the file location. Note that clicking this link selects a Sample Target for the configuration extension.

5. Set the **Default Base Directory** to **OracleHome**.
6. You are now ready to build the collection data specifications. Click **Add**, then click the search icon to log in to the remote file browser. Set the credentials appropriately.
7. In the Oracle home directory of the listener instance, browse to the network/admin subdirectory and select the `sqlnet.ora` file. Add it to the selection table and click **OK**.
8. With the file added to the **Files & Commands** tab, select an appropriate parser from the drop-down list, in this case, the Oracle ORA parser. Click **Preview** if you want to see the file attributes in parsed and raw form as it will appear in the collected data.

Click **Save** to complete creation of the configuration extension.

9. In the Configuration Extension library, select the new configuration extension and click **Manage Deployments**.
10. On the Manage Deployments page, click **Add**. In the dialog that opens, select the targets where you want to deploy the configuration extension.
11. When the status displays Selected for Deployment, click **Apply**. Refresh the view until status is successful, then click **Save**.
12. To verify the added data collection, go to the target instance home page. From the **Oracle Listener** menu, select **Configuration**, then select **Latest Collected**.

The Configuration Browser should display the configuration extension in the tree structure on the left, where you can drill down the directory structure to display the parsed and raw forms of the `sqlnet.ora` attributes and values on the right.

Use this description as a template for extending existing configuration data collections.

Adding New Target Data Collections

The following instructions describe how to extend the configuration data Cloud Control collects by adding a new target type. The example assumes to collect data for a custom Apache web server target type.

First, create a custom target type.

1. From the **Enterprise** menu, select **Configuration**, then select **Configuration Extensions**.
2. From the **Actions** menu, select **Create Custom Target Type**.
3. In the dialog that opens, enter a target type name, `MyApache`, for example. Click **OK**.
4. After a while, a message confirms target type creation. Click **Yes** to add a sample target instance.
5. Click the search icon to select a Management Agent on a host where the application (Apache Tomcat) already resides. Choose the Management Agent and click **Select** to close the dialog, then click **Add Target**.
6. In the target properties dialog that opens, enter the name (`MyApache`) and set the install home path to the start location of the application (Apache Tomcat) on the Management Agent. Click **OK**.
7. In the Configuration Extensions library, click the **Create** button.

- Enter a name (MyApache, for example).
 - Select the custom target type MyApache from the drop-down menu.
 - Click Select Target to select the MyApache sample target instance.
8. You are now ready to build the collection data specifications. Note that the {INSTALL_LOCATION} variable populates the **Default Base Directory** field. Click **Add**, then click the search icon to log in to the remote file browser. Set the credentials appropriately.
 9. In the Apache install home on the Management Agent, browse to the conf directory and select the httpd1.conf file. Add it to the selection table and click **OK**.
 10. With the file added to the **Files & Commands** tab, select an appropriate parser from the drop-down list, in this case, the Apache HTTPD parser. Click **Preview** if you want to see the file attributes in parsed and raw form as it will appear in the collected data.

Click **Save** to complete creation of the configuration extension.
 11. In the Configuration Extensions library, select the new configuration extension and click **Manage Deployments**.
 12. On the Manage Deployments page, click **Add**. In the dialog that opens, select the targets where you want to deploy the configuration extension, for example, the host on which the configuration extension was based.
 13. When the status displays Selected for Deployment, click **Apply**. Refresh the view until status is successful, then click **Save**.
 14. To verify the new data collection, do an all targets search and locate the custom target type under the **Others** category on the left and click it to display all deployments of that type on the right.
 15. Click a target instance (MyApache) in the deployments list on the right. The Configuration Browser should display the configuration extension in the tree structure on the left, where you can drill down the directory structure to display the parsed and raw forms of the httpd1.conf attributes and values on the right.

Use this description as a template for extending configuration data collections through custom target types.

Using Configuration Extensions as Blueprints

Specially formed configuration extensions called blueprints are available for download from Oracle. They are called blueprints because they lay out precisely the files and data to collect for a given platform. Platform support currently includes:

- Apache Tomcat
- Apache Web Server
- GlassFish
- iPlanet
- JBoss
- JRun
- Tuxedo

You can download these blueprints, also called configuration extensions, from the Configuration Management Best Practice Center, where you can also check for new platform support.

Overview of Parsers

A Parser takes raw configuration data and parses it into a nested attribute structure. This structure is a tree hierarchy where nodes are containers and leaves are name value pairs of attributes, or properties.

Configuration extensions include a host of parsers provided by Oracle. Each parser consists of a base parser and parser parameters. Some parsers also contain post-parsing rules. A base parser essentially is a category of parser capable of parsing data of a particular format. Parser parameters provide a way to tailor the base format to accommodate variations in data formatting. Post-parsing rules are a mechanism for aligning nodes in the tree that otherwise have no distinct identity. This is important when comparing configurations and tracking change history to avoid flagging "false positive" differences. It also aids in specifying search criteria and crafting SQL queries used in compliance rules.

There are four varieties of base parser:

- XML
- Format-specific
- Columnar
- Properties

Some parsers have default rules provided by Oracle. These rules address well-known instances where nodes need to be aligned. Specifically, the WebLogic and WebSphere parsers contain default rules to address such instances. You can leave these rules as is, execute a subset of them, or replace them with your own custom rules.

This section covers the following topics:

- [Managing Parsers](#)
- [About XML Parsers](#)
- [About Format-Specific Parsers](#)
- [About Columnar Parsers](#)
- [About Properties Parsers](#)
- [Using Parsed Files and Rules](#)

Managing Parsers

While creating, editing, or viewing configuration extensions, you can peruse the list of available parsers, their default parameters, and post-parser rules, if applicable. Parser parameters dictate formatting such as comment character, delimiters, start and end characters, and so forth. You cannot edit these parameters, but you can export a parser as an XML file, edit the file, and import it back into the application under a new name. Some parsers also have default rules that serve to align nodes in the parsed tree for purposes of comparison, for example.

1. In the Configuration Extensions library, select **Manage Parsers** from the **Actions** menu. A list of available parsers appears in a table. The column on the right (Base Parsers) denotes a general parser category, Properties for example, which implies file types that contain name/value pairs.
2. Select a parser and click **Details**. This dialog also shows default rules, if any.
 - Click the **Parameters** tab to see the parameter defaults in effect. You can then judge if you need to edit the parser to conform with your file format conventions.
 - Click the **Default Rules** tab to see the post-parsing rules that ship with certain parsers. This is a good way to get exposure to rules formation.
3. Assume you want to change the delimiter character in a given parser.
 - a. With the parser selected in the table, click **Export**.
 - b. In the dialog that opens click **Save** and navigate to a file system location. Save the XML file with an appropriate name.
 - c. In making your edits, be sure to change the parser ID and parser name in the XML, as you are creating a customized version of a parser provided by Oracle.
4. Assume you now want to import the new parser you saved for use in creating configuration extensions.
 - a. With the Parsers table open, click **Import**.
 - b. In the dialog that opens, browse to the file location where you saved the exported parser file. Select it and click **Import** on the dialog.

The new parser now appears in the Parsers table where it can be used in configuration extension creation.

About XML Parsers

Cloud Control has two XML parsers: a default (attribute-keyed) XML parser and a generic XML parser.

About the Default XML Parser

Parsing occurs as follows:

- XML elements with no XML attributes or child elements become parsed attributes; all other elements become containers.
- XML attributes become parsed attributes.
- Element text content becomes a parsed attribute, with its name dependent on whether or not the tag contains any XML attributes. If the tag contains XML attributes, the parsed attribute name takes the value specified in the STORE_CONTENT_AS parameter; otherwise, the parsed attribute name takes the tag name.

The default XML parser accepts the following parameters:

Parameter	Description
MULTIKEY_DELIMITER	Delimiter that separates a list of XML attribute names in the CONTAINER_NAME parameter; default is tilde (~)

Parameter	Description
STORE_CONTENT_AS	Name to give to parsed attributes derived from element text content, where the element contains XML attributes; default is <code>text_value</code>
CONTAINER_NAME	A list of XML attribute names delimited by the value of the <code>MULTIKEY_DELIMITER</code> parameter. If an attribute name in this list appears in a tag in the original file, the tag becomes a container named for the value of the XML attribute. All other XML attributes become parsed attributes as usual. The tag name itself is discarded. For example, the list includes attribute names Moe and Larry in this order. The original file contains an XML tag <code>Stooges</code> that has attributes Moe, Larry, and Curly. As Moe appears first in the delimited list, its value, <code>leader</code>, becomes the parsed container name; Larry and Curly become parsed attributes. The tag name <code>Stooges</code> is discarded. The original XML fragment might be as follows: <pre><?xml version="1.0" encoding="UTF-8"?> <Comedy> <Stooges Moe="leader", Larry="zany", Curly="bald"> </Stooges> </Comedy></pre>

WebLogic Attribute-keyed Parser

Cloud Control provides an attribute-keyed parser provided by Oracle specifically designed to parse the `WebLogic config.xml` file. It has the same parameters as the default XML parser and comes with 26 default post-parsing rules to uniquely identify nodes with the same name.

WebSphere Attribute-keyed Parsers

Cloud Control provides several attribute-keyed parsers provided by Oracle designed to parse specific WebSphere configuration files. Each parser has the same parameters as the default XML parser and comes with a set of default post-parsing rules to uniquely identify nodes with the same name. There are parsers for the following WebSphere configuration files:

- `node.xml` (1 default post-parsing rule)
- `plugin-cfg.xml` (7 default post-parsing rules)
- `resource.xml` (9 default post-parsing rules)
- `server.xml` (13 default post-parsing rules)
- `variables.xml` (1 default post-parsing rule)

About the Generic XML Parser

Parsing occurs as follows:

- All XML elements become containers.
- All XML attributes become parsed attributes.

- Element text content becomes a parsed attribute that takes the name `text_value`, where the text content becomes the parsed attribute value.

The generic XML parser accepts no parameters.

WebSphere Generic Parser

Cloud Control provides one generic parser provided by Oracle designed to parse the WebSphere `serverindex.xml` configuration file. It comes with three default post-parsing rules to uniquely identify nodes with the same name.

XML Parser Examples

This section contains three XML parser examples:

- As parsed using the default XML parser, with parameter values provided by Oracle
- As parsed using the default XML parser, with modified parameter values
- As parsed using the generic XML parser

Parsed examples derive from the following original XML file:

```
<?xml version="1.0" encoding="UTF-8"?>
<Application>
  <AppName>foo</AppName>
  <Server name="ajax" os="linux">production</Server>
</Application>
```

Default XML Parser (Parameter Values Provided by Oracle)

When parsed using the default XML parser with parameter values provided by Oracle, the parsed version appears as follows:

```
Application
  AppName = foo
  Server
    name = ajax
    os = linux
    text_value = production
```

Note the following about this parsed version:

- The element contents of the `AppName` and `Server` tags become parsed attributes.
- Since the `AppName` tag contains no XML attributes, the parsed attribute name takes the tag name.
- Contrast with the `Server` tag, which has XML attributes (`name` and `os`). This results in a container named for the tag (`Server`), with three parsed attributes, one for each of the XML attributes, and a third for the text content of the `Server` tag, which is set to the value of the `STORE_CONTENT_AS` parameter (`text_value`).

Default XML Parser (Modified Parameter Values)

To modify parameter values, you have to create a new parser by exporting the default XML parser, modifying the exported XML file, and importing the modified parser, using a new name and parser ID.

Assume you followed this process, making the following modifications:

- Set the STORE_CONTENT_AS parameter to the value myVal
- Set the CONTAINER_NAME parameter to the value name

When parsed using the default XML parser with modified parameter values, the parsed version appears as follows:

```
Application
  AppName = foo
  ajax
    os = linux
    myVal = production
```

Note the following about this parsed version:

- The AppName tag remains the same; that is, it has no XML attributes so it becomes a parsed attribute.
- Since the Server tag has an XML attribute that matches the value of CONTAINER_NAME, the container takes the value of the attribute (ajax), obviating the name=ajax parsed attribute. Remember that the CONTAINER_NAME parameter provided by Oracle has a placeholder but no actual default value; thus, the difference in this version of the parsed representation.
- The remaining Server tag attribute (os) becomes a parsed attribute as usual, and the text content associated with the tag becomes the value of the attribute myVal, per the edited STORE_CONTENT_AS parameter.

Generic XML Parser

When parsed using the generic XML parser (the one that takes no parameters), the parsed version appears as follows:

```
Application
  AppName
    text_value = foo
  Server
    name = ajax
    os = linux
    text_value = production
```

See [About the Default XML Parser](#) for a reminder of how parsing occurs.

About Format-Specific Parsers

Format-specific base parsers are applicable only to a particular data format. Format-specific parsers run the gamut from having no parameters to a few to many with which to tailor formatting.

Parser	Description
Blue Martini DNA	Parser for Blue Martini DNA files (no parameters).
Connect:Direct	Parser for Connect:Direct .cfg files (no parameters).

Parser	Description
Database Query (see Sample SQL Query Parsing and Rule Application for an example)	Parser for configuration extension database query output. Cloud Control automatically transforms query results into a format the parser accepts, organizing results into sections similar to a Windows .ini file. Each section represents one record; each line in a section contains a table column name and a value. See Database Query Parser Parameters .
Database Query Paired Column	Parser for configuration extension database query output. Cloud Control automatically transforms query results into a format the parser accepts, organizing results into sections similar to a Windows .ini file. Each section represents one record; each line in a section contains a name and value, where both the name and the value are values of returned columns. As such, the parser requires an even number of columns to be returned by the query in order to parse the data. A query which returns an odd number of columns will result in a parsing error. See Database Query Paired Column Parser Parameters .
Db2	Parser for the output of the DB2 GET DATABASE CONFIGURATION command (no parameters).
Directory	Parser for files containing multiple name value pairs on the same line, where each line may have varying numbers of pairs. For example, the first line might be a=b j=k, the second line c=d m=n y=z, and so forth. See Directory Parser Parameters .
E-Business Suite	Parser for E-Business Suite .drv files. The parser converts IF . . . THEN . . . ELSE structures in the file into containers in the parsed representation, and the rest of the lines into a container with a fixed number of parsed attributes. These lines can be of two types: directory specifications, whose parsed attribute names are specified in the DIR_HEADER parser parameter; configuration file specifications, whose parsed attribute names are specified in the HEADER parser parameter. See E-Business Suite Parser Parameters .
Galaxy CFG	Parser for Galaxy .cfg files. See Galaxy CFG Parser Parameters .
Introscope	Parser for Introscope files (no parameters).
MQ-Series	Parser for MQ-Series files. See MQ-Series Parser Parameters .
Odin	Parser for Odin files (no parameters).
Oracle ORA	Parser for Oracle .ora files, such as tnsnames.ora (no parameters).
Siebel	Parser for Siebel siebns files. The parser creates a container for each unique path in the file, and attributes for name value pairs, except where a line contains the string Type=empty, in which case the parser does not create a parsed attribute for the line. See Siebel Parser Parameters .
UbbConfig	Parser for BEA Tuxedo files (no parameters). The parser converts sections prefixed with an asterisk (*), and names in double quotes at the start of a new line, into containers. It converts all other data into attributes.
Unix Installed Patches	Parser for Unix installed patches data. The parser creates one container per (non-comment) line of the file. It treats every field ending with a colon (:) on each line as a property name field and the value following, if any, as the property value. Note that a property does not have to have a value. See Unix Installed Patches Parser Parameters .

Parser	Description
Unix Recursive Directory List	Parser for output of Unix recursive directory listing (<code>ls -l -R</code>). The parser converts each subdirectory line into a container, and each file information line into a container with a fixed set of attributes. See Unix Recursive Directory List Parser Parameters .

Remember, to modify a format-specific parser, you have to create a new parser by exporting the particular parser, modifying the exported XML file, and importing the modified parser, using a new name and parser ID.

Database Query Parser Parameters

The following table describes the parameters with which you can customize the Database Query parser:

Parameter	Description
CELL_DELIMITER	Character that separates name value pairs; default is =.
PROPERTY_DELIMITER	Character that separates the length of a name or value from the value itself; default is _.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.
SECTION_START	Character that denotes the start of a section; default is \[(backslash is escape character).
SECTION_END	Character that denotes the end of a section; default is \] (backslash is escape character).
USE_INI_SECTION	Flag that tells the parser to use Windows .ini type sections; default is true.

Database Query Paired Column Parser Parameters

The following table describes the parameters with which you can customize the Database Query parser:

Parameter	Description
CELL_DELIMITER	Character that separates name value pairs; default is =.
PROPERTY_DELIMITER	Character that separates the length of a name or value from the value itself; default is _.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.
SECTION_START	Character that denotes the start of a section; default is \[(backslash is escape character).
SECTION_END	Character that denotes the end of a section; default is \] (backslash is escape character).
USE_INI_SECTION	Flag that tells the parser to use Windows .ini type sections; default is true.

Directory Parser Parameters

The following table describes the parameters with which you can customize the Directory parser:

Parameter	Description
CELL_DELIMITER	Character that separates one property from another; default is a space.
EXTRA_DELIMITER	Character that separates a property name from its value; default is =.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.

E-Business Suite Parser Parameters

The following table describes the parameters with which you can customize the E-Business Suite parser:

Parameter	Description
DIR_HEADER	A tilde-delimited list of attribute names for directory specifications.
STRUCTURE_START	A tilde-delimited list of regular expressions denoting the start of a structure.
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
HEADER	A tilde-delimited list of attribute names for file specifications.
COMMENT	A tilde-delimited list of regular expressions denoting comments.
STRUCTURE_END	A tilde-delimited list of regular expressions denoting the end of a structure.
LAST_FREE_FORM	Flag that tells the parser to ignore cell delimiters in the last value of a directory or file specification; default is true.
ELEMENT_FIELD	A tilde-delimited list of file specification attribute names. The parser concatenates values of the specified attributes to form the name of the container associated with the file specification.
DIR_ELEMENT_FIELD	A tilde-delimited list of directory specification attribute names the parser uses to determine the name of the container associated with the directory specification.

Galaxy CFG Parser Parameters

The following table describes the parameters with which you can customize the Galaxy CFG parser:

Parameter	Description
COMMENT	Character that tells the parser to ignore the line that follows; default is !.
ADD_SUFFIX	Names of attributes whose values to append to a container name.
MONO_PROP_SECTION	Names of sections that have a single property.

Parameter	Description
MULTI_PROP_SECTION	Names of sections that have multiple properties.
NODES_SECTION	Names of section start and end elements

MQ-Series Parser Parameters

The MQ-Series parser has a single parameter that you can customize: `COMMENT`, which defaults to `*`.

Siebel Parser Parameters

The following table describes the parameters with which you can customize the Siebel parser:

Parameter	Description
LINES_TO_SKIP	Tells the parser the number of lines to ignore at the beginning of the file; default is 4.
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
COMMENT	A tilde -delimited list of regular expressions denoting comments.
SECTION_START	A tilde-delimited list of regular expressions denoting the start of a unique path specification section.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a unique path specification section.
USE_INI_SECTION	Flag that tells the parser to use Windows <code>.ini</code> type sections; default is true.

Unix Installed Patches Parser Parameters

The following table describes the parameters with which you can customize the Unix Installed Patches parser:

Parameter	Description
CELL_DELIMITER	Character that separates name value pairs; default is a space.
EXTRA_DELIMITER	Character that separates a property name from its value; default is <code>..</code> .
COMMENT	Character that tells the parser to ignore the line that follows; default is <code>#</code> .

Unix Recursive Directory List Parser Parameters

The following table describes the parameters with which you can customize the Unix Recursive Directory List parser:

Parameter	Description
LINES_TO_SKIP	Tells the parser the number of lines to ignore at the beginning of the file; default is 4.

Parameter	Description
CELL_DELIMITER	A tilde-delimited list of regular expressions denoting name value pair delimiters.
COMMENT	A tilde-delimited list of regular expressions denoting comments.
HEADER	A tilde-delimited list of attribute names.
LAST_FREE_FORM	Flag that tells the parser to ignore cell delimiters in the last value of a line; default is true.
SECTION_START	A tilde-delimited list of regular expressions denoting the start of a subdirectory section.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a subdirectory section.
ELEMENT_FIELD	A tilde-delimited list of attribute names. The parser concatenates values of the specified attributes to form the name of the container associated with the line.

About Columnar Parsers

Columnar parsers are inherently flexible owing to the parameters they can accept to tailor formatting. All columnar parsers use a subset of the same parameters.

Parser	Description
Cron Access	Parser for <code>cron.allow</code> and <code>cron.deny</code> files.
Cron Directory	Parser for <code>Unix etc</code> and <code>cron.d</code> files.
CSV	Parser for comma-separated-value (CSV) data. Because the number of columns in the CSV parser is unknown, use the CSV parser provided by Oracle as a template for the new CSV parser. Export the provided CSV parser, update the parameters, and re-import the new CSV parser tailored to your format. The parameter values provided by Oracle support CSV files with these characteristics: • Each line has the same number of values • The first parsed (that is, non-comment) line is a header line whose content is a comma-separated list of column names • Commas in double quotes are considered part of the value, not value delimiters • One of the column names is "name" whose value becomes the container name associated with each line Text inside double quotes is considered part of a value; to specify a value that contains a double quote, escape the double quote with a backslash (\). Use a backslash to escape the backslash character itself (\\).
Hosts Access	Parser for <code>hosts.allow</code> and <code>hosts.deny</code> files.
Kernel Modules	Parser for <code>kernel modules</code> files.
Linux Directory List	Parser for Linux directory listing data format (for example, output of a <code>ls -l</code> command).
PAM Configuration	Parser for <code>pam.conf</code> files.
PAM Directory	Parser for <code>Unix etc/pam.d</code> files.
Process Local	Parser for <code>process.local</code> files.
Secure TTY	Parser for <code>Unix etc/securetty</code> files.

Parser	Description
Solaris Installed Packages	Parser for Solaris installed packages files.
Unix Crontab	Parser for Unix crontab files.
Unix Directory List	Parser for Unix directory listing data format for example, the output of a <code>ls -l</code> command).
Unix Groups	Parser for Unix <code>etc/group</code> files. The parser ignores group name and password information.
Unix GShadow	Parser for Unix <code>etc/gshadow</code> files.
Unix Hosts	Parser for Unix <code>etc/hosts</code> files.
Unix INETD	Parser for Unix <code>etc/inetd.conf</code> files.
Unix Passwd	Parser for Unix <code>etc/passwd</code> files. The parser ignores password values.
Unix Protocols	Parser for Unix <code>etc/hosts</code> files.
Unix Services	Parser for Unix <code>etc/services.conf</code> files.
Unix Shadow	Parser for Unix <code>etc/shadow</code> files.
Unix System Crontab	Parser for Unix system crontab files. System crontab files are very similar to crontab files, but may contain name value pairs such as <code>PATH=/a/b</code> .

Columnar Parser Parameters

This section describes all columnar base parser parameters. Although the base parser can accept values for any of these parameters, a given parser specification does not necessarily need to provide values for all of them. All parameters have default values, which are used in the absence of a specified value, although in some cases, parameters have explicit values.

Use quotes when delimiters or other special text such as comment characters or new lines are part of some value. The `QUOTE_DELIMITER` determines the character value to use. Prefix the quote delimiter with a backslash (`\`) if you need to escape the character. Use a backslash to escape the backslash character itself (`\\`) in quoted strings.

Parameter	Description
COMMENT	A tilde-delimited list of regular expressions that denote comment characters or sequences. For example, <code>#[^\r\n]*</code> specifies that everything on a line following the <code>#</code> character is a comment. Default is an empty list; that is, parse all file contents.
LINES_TO_SKIP	The number of initial lines (excluding blank or comment lines) to ignore for parsing purposes, treating them in effect as comments. Default is 0; that is, skip no lines.
CELL_DELIMITER	A tilde-delimited list of regular expressions that delimit line values. Default is an empty list; that is, no delimiters (it is unusual to use the default).
QUOTE_DELIMITER	A tilde-delimited list of regular expressions that define how quoted values begin and end (usually either a single or double quote character). The beginning and end quote delimiter must be the same. Default is an empty list; that is, parser does not recognize quoted values.

Parameter	Description
PROPERTY_DELIMITER	A tilde-delimited list of regular expressions that delimit property names and values. Default is an empty list; that is, no property delimiters. Rarely, a columnar file may contain name value pairs of the syntax a=b.
RESERVED_DIRECTIVES	A tilde-delimited list of property keywords. Some crontab files contain lines of simple name value pairs, separated by a delimiter (foo=bar), thus violating the requirement that each line have the same number of fields. This parameter provides a workaround to specify property keywords. In the example, the property keyword would be foo. This says, in effect, parse any line beginning with this keyword as a parsed attribute name value pair under the root container. Default is an empty list; that is, no property keywords.
ALTERNATE_DELIMITER	An alternate delimiter for property names and values. Default is '/' (used only if ALTERNATE_FIELD parameter is nonempty).
ALTERNATE_FIELD	A tilde-delimited list of fields separated by alternate delimiters. Default is an empty list; that is, no alternate delimiters.
HEADER_FLAG	A flag specifying whether or not the file has a header line that specifies the column names. Default is false.
HEADER	A tilde-delimited list of column names to use if there is no header line. Default is an empty list; that is, no column names (it is unusual to use the default).
ELEMENT_FIELD	A tilde-delimited list of column names whose values the parser concatenates to create the container name associated with a line. Default is an empty list; that is, no column names (it is unusual to use the default).
IGNORE_FIELD	A tilde-delimited list of column names to ignore. No parsing of values in these columns occurs. Default is an empty list; that is, ignore nothing.
LAST_FREE_FORM	A flag that specifies whether the last column is free form. The parser ignores all delimiters in a free form column value. Default is false.
USE_LINE_COMMENT	A flag that specifies whether to treat end of line comments as a value to appear in the parsed representation of the data. Default is false.

About Properties Parsers

Properties parsers are inherently flexible owing to the parameters they can accept to tailor formatting and handle disparate organizational elements. All properties parsers use the same set of basic and advanced parameters, as well as advanced constructs.

Parser	Description
AIX Installed Packages	Parser for AIX installed packages files.
Apache HTTPD	Parser for Apache HTTPD.conf files.
Autosys	Parser for Autosys.jil files.

Parser	Description
Custom CFG	Parser for custom .cfg files. This syntax defines an element with E = {} syntax, where the brackets may contain name value pairs, nested elements, or both.
Java Policy	Parser for java.policy files.
Java Properties	Parser for java.properties files.
LDAP	Parser for LDAP .cfg files.
Mime Types	Parser for mime.types files.
Radia	Parser for Radia .cfg files.
Sectioned Properties	Parser for files containing name value pairs organized into sections, such as a Windows .ini file.
SiteMinder Agent	Parser for SiteMinder agent files.
SiteMinder Registry	Parser for SiteMinder .registry files.
SiteMinder Report	Parser for SiteMinder SmReport.txt files.
SmWalker	Parser for SiteMinder SmWalker.dat files.
Sun ONE Magnus	Parser for Sun ONE magnus.conf files.
Sun ONE Obj	Parser for Sun ONE obj.conf files.
Tuxedo	Parser for Tuxedo files.
Unix Config	Parser for Unix etc/config files.
Unix Login	Parser for Unix etc/login.defs files.
Unix PROFTPD	Parser for Unix etc/proftpd.conf files.
Unix Resolve	Parser for Unix etc/resolve.conf files.
Unix SSH Config	Parser for Unix etc/ssh/sshd.conf files.
Unix System	Parser for Unix etc/system files.
Unix VSFTPD	Parser for Unix etc/vsftpd.conf files.
Unix XINETD	Parser for Unix etc/xinetd.conf files.
WebAgent	Parser for WebAgent files.
Windows Checksum	Parser for Windows checksum output generated with fciv.exe.

Basic Properties Parser Parameters

This section describes basic properties parser parameters that are required to parse simple property data formats. Simple property data formats specify a property as a name value pair, usually with a defined delimiter separating the name and the value: foo=bar. The basic data format is a list of properties, one property to a line, together with optional comments; a java.properties file, for example. All parameters have default values, which are used in the absence of a specified value.

Use quotes when delimiters or other special text such as comment characters or new lines are part of some value. The QUOTE_DELIMITER determines the character value to use. Prefix the quote delimiter with a backslash (\) if you need to escape the character. Use a backslash to escape the backslash character itself (\\) in quoted strings.

A comment character such as the pound sign (#), or a particular character sequence (//) usually denotes a comment. Special sequences such as a C style comment (/*...*/) might denote the beginning and end of a comment. Some files might have generic informational content in the first couple of lines. In this case, a parameter is available to tell the parser to ignore these lines.

Parameter	Description
COMMENT	A tilde-delimited list of regular expressions that denote comment characters or sequences. For example, #[^\r\n]* specifies that everything on a line following the # character is a comment. Default is an empty list; that is, parse all file contents.
LINES_TO_SKIP	The number of initial lines (excluding blank or comment lines) to ignore for parsing purposes, treating them in effect as comments. Default is 0; that is, skip no lines.
CELL_DELIMITER	A tilde-delimited list of regular expressions that delimit line values. Default is an empty list; that is, no delimiters (it is unusual to use the default).
QUOTE_DELIMITER	A tilde-delimited list of regular expressions that define how quoted values begin and end (usually either a single or double quote character). The beginning and end quote delimiter must be the same. Default is an empty list; that is, parser does not recognize quoted values.
ALLOW_NAME_ONLY_PROPERTIES	A flag that indicates whether the parser allows property names without a delimiter or a value. Default: false.
REVERSE_PROPERTY	A flag that indicates whether the parser allows the value to come before the delimiter and property name. Default: false.

Advanced Properties Parser Parameters

This section describes advanced properties parser parameters that are required to parse more complex property data formats. All parameters have default values, which are used in the absence of a specified value.

Parameter	Description
PROPERTY_DELIMITER	A tilde-delimited list of regular expressions denoting property delimiters. For example, the text "a=b : x=y" could be interpreted in either of two ways: - As a single property "a" with value "b : x=y" - As two separate properties, "a=b" and "x=y" If a colon (:) is the property delimiter, the parsing engine interprets the text as containing two separate properties. Default is an empty list; that is, parser does not recognize property delimiters.
LINE_END_DELIMITER	A tilde-delimited list of regular expressions denoting line end sequences. When the parser encounters a line end delimiter, it assumes a new property or construct starts on the next line. Default is an empty list; that is, parser does not recognize line end delimiters.

Parameter	Description
CONTINUE_LINE	A tilde-delimited list of regular expressions denoting continue line sequences. When the parser encounters a continue line pattern, it interprets data on the following line as a continuation of the construct or property on the previous line, as opposed to interpreting the new line as the beginning of a new property or construct. For example, the parser must encounter a line continuation pattern to recognize property values that span multiple lines. Default is an empty list; that is, parser does not recognize line continuation patterns.
SECTION_START	A tilde-delimited list of regular expressions denoting the beginning of a section. Sections cannot be nested. Default is an empty list; that is, parser does not recognize sections.
SECTION_END	A tilde-delimited list of regular expressions denoting the end of a section. Default is an empty list.
STRUCTURE_START	A tilde-delimited list of regular expressions denoting the beginning of a structure. Structures can be nested. Default is an empty list; that is, parser does not recognize structures.
STRUCTURE_END	A tilde-delimited list of regular expressions denoting the end of a structure. Default is an empty list.
XML_STYLE_TAG	A flag that indicates whether structures in the file are XML style tags. Default: false.
USE_INI_SECTION	A flag that indicates whether INI style sections are present. Default: false.
RESERVED_DIRECTIVES	A tilde-delimited list of reserved names indicating the start of a reserved directive. Default is an empty list; that is, parser does not recognize reserved directives.
RESERVED_FUNCTIONS	A tilde-delimited list of reserved names indicating the start of a reserved function. Default is an empty list; that is, parser does not recognize reserved functions.
DIRECTIVE_PROPERTIES	A tilde-delimited list of reserved directive-implicit property names. Default is an empty list.
FUNCTION_PROPERTIES	A tilde-delimited list of required reserved function-explicit property names. Default is an empty list.
SECTION_PROPERTIES	A tilde-delimited list of section-implicit property names. Default is an empty list.
STRUCTURE_PROPERTIES	A tilde-delimited list of structure-implicit property names. Default is an empty list.
ELEMENT_FIELD	A keyword to be ignored by the parser when parsing properties. This typically applies to data formats that specify a keyword before a name value pair; "set a=b" for example. Default is an empty list; that is, parser ignores nothing.
ALLOW_ELEMENT_CELL	A flag that indicates whether the file format supports element cell structures. Default: false.
SECTION_EXPLICIT_PROPERTIES	A flag that indicates whether sections support explicit properties. Default: false.
STRUCTURE_EXPLICIT_PROPERTIES	A flag that indicates whether structures support explicit properties. Default: false.
NEWLINE_CONTINUE_LINE	A flag that indicates whether newlines can be line continuation sequences. Default: false.

Parameter	Description
KEYWORD_FIELD	A tilde-delimited list of regular expressions denoting keywords that precede properties that use a whitespace delimiter. Default is an empty list; that is, parser does not recognize keywords.

Advanced Properties Parser Constructs

Properties files come in variety of file formats. To accommodate the widest possible range of formats, the generic properties base parser uses combinations of constructs found in most files.

The constructs fall into two categories:

- Container constructs, which can be reserved functions, reserved directives, XML structures, structures, delimited structures, INI sections, delimited sections, sections, and element cells
- Property constructs, which can be simple properties, reverse properties, keyword properties, keyword name properties, bracket properties, implicit properties and explicit properties

Of the element constructs, section constructs cannot be nested, but can contain any other construct. Structure constructs can be nested, and can contain any construct except a section. Element cells can be nested, but can only contain element cells and simple properties. Reserved directives and reserved functions cannot be nested, nor can they contain any other constructs.

The rest of this section describes the constructs the base properties parser supports.

Simple Property

A simple property consists of a property name, cell delimiter, property value, and new-line sequence, in that order. A simple property may take up more than one line, although properties that span multiple lines usually contain a line continuation character or sequence. The parser ignores whitespace such as tabs and spaces, unless a parameter specifies whitespace as having some significance (cell delimiter, for example).

Example: `name=value_that_wraps_to_next_line_/,` where the forward slash serves as a line continuation character. A Java Properties file typifies this data format.

Keyword Property

This construct is the same as a simple property, only with a keyword in front, which the parser ignores.

Example: `set name=value,` where `set` is the ignored keyword. A Unix System file typifies this data format.

Keyword Name Property

This construct is a simple property where the property name matches a regular expression specified in the `KEYWORD_FIELD` parser parameter. This is a special case property type specific to the Unix XINETD parser. The XINETD file uses an equal sign (=) as a cell delimiter except when the property begins with the keyword "include" or "includedir", in which case the cell delimiter is whitespace.

While added specifically for XINETD files, the property can be used for other file types where appropriate.

Example: `includedir /etc`, where `includedir` is the parser parameter regular expression and whitespace is the cell delimiter.

Explicit Property

An explicit property consists of a property name, a delimiter, and a property value. Unlike a simple or keyword property, an explicit property is bound to a container construct such as a section or a structure; an XML tag attribute, for example.

Examples:

```
[SectionName p1=v1 p2=v2]

<StructureName p1=v1 p2=v2>
...
</StructureName>
```

In these constructs, the name value pairs `p1 v1` and `p2 v2` are explicit properties. A Sun ONE Obj file typifies this data format.

Implicit Property

An implicit property is a property value without an associated property name. Like an explicit property, an implicit property is bound to a container construct, usually a reserved directive. The `DIRECTIVE_PROPERTIES` parser parameter contains the property names of implicit properties.

Examples:

```
[SectionName myName myPath]

<StructureName myName myPath>
...
</StructureName>
```

In these constructs, the implicit properties have the values `myName` and `myPath`, with the presumed property names `name` and `path`, as declared in the `DIRECTIVE_PROPERTIES` parser parameter. An Apache HTTPD file typifies this data format.

Reserved Function

A reserved function is a keyword followed by one or more explicit properties. The `RESERVED_FUNCTIONS` parser parameter specifies keywords that denote reserved functions.

Example: `Error fn="query-handler" type="forbidden"`, where `Error` is the reserved function keyword specified in the `RESERVED_FUNCTIONS` parser parameter. A Sun ONE Magnus file typifies this data format.

Reserved Directive

A reserved directive is a keyword followed by one or more implicit properties. The `RESERVED_DIRECTIVES` parser parameter specifies keywords that denote reserved directives.

Example: `LoadModule cgi_module "/bin/modules/std/cgi"`, where `LoadModule` is the reserved function keyword specified in the `RESERVED_DIRECTIVES` parser parameter. An Apache HTTPD file typifies this data format.

XML Structure

An XML structure is a standard XML tag that can contain a name only, a name followed by explicit properties, or a name followed by implicit properties.

Examples:

```
<Name>
...
</Name>

<Name p1=v1 p2=v2>
...
</Name>
<Name "implicit_property1" "implicit_property2">
...
</Name>
```

A WebAgent file typifies this data format.

Delimited Structure

A delimited structure consists of the following (in the specified order):

- Structure name
- Delimiter
- Start structure character or character sequence
- Structure contents
- End structure character or character sequence

Example:

```
StructureName = {
...
}
```

Explicit and implicit properties are not allowed. Java Policy and Custom CFG files typify this data format.

Structure

A structure consists of the following (in the specified order):

- Structure name
- Start structure character or character sequence
- Structure contents
- End structure character or character sequence

The only difference between a delimited structure and a structure is the delimiter; that is, a structure does not require a delimiter between the structure name and the start structure indicator.

Example:

```
StructureName {
  ...
}
```

Explicit and implicit properties are not allowed. A Unix XINETD file typifies this data format.

INI Section

And INI section resembles a section heading in a Windows .ini file, characterized by:

- Section start character or character sequence
- Section name
- Optional (explicit and implicit) properties
- Section end character or character sequence

Examples:

```
[SectionName]
```

```
[SectionName p1=v1 p2=v2]
```

```
[SectionName "implicit_property1" "implicit_ property2"]
```

SmWalker and Sectioned Properties files typify this data format.

Delimited Section

A delimited section is a line that begins with a common pattern, but otherwise resembles a simple property.

Examples:

```
HKEY_LOCAL_MACHINE\SOFTWARE\A\B\C=789
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\X\Y\Z=123
```

These are two delimited section headings where the common pattern is HKEY_. Site-Minder Registry and LDAP files typify this data format.

Element Cell

An element cell consists of an element cell name and a property name value pair of the form A = B = C. Element cells typically use line continuation sequences and nesting to clarify the structure. An element cell that has multiple properties uses a property delimiter to separate them.

Example 1:

```
EC = \
    B = C, D = F
```

This example is an element cell named EC with two property name value pairs, B = C and D = F, separated by a comma. The structure uses the backslash character (\) to indicate line continuation. The advanced properties parser parameters PROPERTY_DELIMITER and CONTINUE_LINE define the respective format characters.

Example 2:

```

EC = \
    EC2 = \
        A = B, \
        C = D

```

This example is an element cell named EC that has a nested element cell named EC2 that contains two property name value pairs, A = B and C = D. This example uses the same delimiter and line continuation characters.

Using Parsed Files and Rules

A collected configuration file is stored in raw form and, if a parser is specified, in a tree structure of nodes, or containers, and attributes, or properties. The file also is generated internally in XML format for the purpose of applying post-parsing rules, which consist of XPath conditions and expressions. Note that even non-XML files are generated in this internal format. Since the internal format must accommodate other file types, it introduces an additional root node in the XML to compensate for files such as Java properties files that have only attribute names and values.

Examples of how files are parsed and displayed, and the effects of post-parsing rules help to clarify:

- [Sample XML File Parsing and Rule Application](#)
- [Sample Non-XML File Parsing and Rule Application](#)
- [Sample SQL Query Parsing and Rule Application](#)

Sample XML File Parsing and Rule Application

Consider the following simple XML file:

```

<dir name="/a/b/c">
  <file name="file1" size=120/>
  <file name="file2" size=350/>
</dir>

```

Its parsed form, using the default XML parser, appears in the user interface in the following tree structure:

```

dir
  name      = /a/b/c
  file
    name = file1
    size = 120
  file
    name = file2
    size = 350

```

Notice that two containers have the same name (file), which makes it impossible to distinguish between the two, at the container level, at least. Thus, this file is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```

<root>
  <dir>

```

```

<name>/a/b/c</name>
<file>
  <name>file1</name>
  <size>120</size>
</file>
<file>
  <name>file2</name>
  <size>350</size>
</file>
</dir>
</root>

```

Given the problem in the parsed form of having two containers with the same name, a rule resolution might consist of the following:

Condition: /root/dir/file

Expression: name/text()

Effectively, this says: for each file evaluate `name/text()` to produce an identifier that distinguishes one file from another within the `dir` node.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```

dir
  name = /a/b/c
  file[file1]
    name = file1
    size = 120
  file[file2]
    name = file2
    size = 350

```

The rule resolves to an identifier appended in square brackets to the container name. The combination (`file[file1]`, for example) enables various operations such as compare, search, change history, and so forth, to distinguish between file containers.

Sample Non-XML File Parsing and Rule Application

Consider the following simple ORA file:

```

acme=
  (DESCRIPTION=
    (SOURCE_ROUTE=yes)
    (ADDRESS=(PROTOCOL=tcp)(HOST=host1)(PORT=1630))
    (ADDRESS_LIST=
      (FAILOVER=on)
      (LOAD_BALANCE=off)
    )
    (ADDRESS=(PROTOCOL=tcp)(HOST=host2a)(PORT=1630))
    (ADDRESS=(PROTOCOL=tcp)(HOST=host2b)(PORT=1630))
    (ADDRESS=(PROTOCOL=tcp)(HOST=host3)(PORT=1630))
    (CONNECT_DATA=(SERVICE_NAME=xyz.example.com)))

```

Its parsed form, using the Oracle ORA parser, appears in the user interface in the following tree structure:

```

acme
  DESCRIPTION
    SOURCE_ROUTE    yes

```

```

ADDRESS
  PROTOCOL      tcp
  HOST           host1
  PORT          1630
ADDRESS_LIST
  FAILOVER       on
  LOAD_BALANCE   off
  ADDRESS
    PROTOCOL      tcp
    HOST           host2a
    PORT          1630
  ADDRESS
    PROTOCOL      tcp
    HOST           host2b
    PORT          1630
ADDRESS
  PROTOCOL      tcp
  HOST           host3
  PORT          1630
CONNECT_DATA
  SERVICE_NAME   xyz.example.com

```

Notice that the address containers, both standalone and within ADDRESS_LIST are indistinguishable. Thus, this file is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```

<root>
  <acme>
    <DESCRIPTION>
      <SOURCE_ROUTE>yes</SOURCE_ROUTE>
      <ADDRESS>
        <PROTOCOL>tcp</PROTOCOL>
        <HOST>host1</HOST>
        <PORT>1630</PORT>
      </ADDRESS>
      <ADDRESS_LIST>
        <FAILOVER>on</FAILOVER>
        <LOAD_BALANCE>off</LOAD_BALANCE>
        <ADDRESS>
          <PROTOCOL>tcp</PROTOCOL>
          <HOST>host2a</HOST>
          <PORT>1630</PORT>
        </ADDRESS>
        <ADDRESS>
          <PROTOCOL>tcp</PROTOCOL>
          <HOST>host2b</HOST>
          <PORT>1630</PORT>
        </ADDRESS>
      </ADDRESS_LIST>
      <ADDRESS>
        <PROTOCOL>tcp</PROTOCOL>
        <HOST>host3</HOST>
        <PORT>1630</PORT>
      </ADDRESS>
      <CONNECT_DATA>
        <SERVICE_NAME>xyz.example.com</SERVICE_NAME>
      </CONNECT_DATA>
    </DESCRIPTION>
  </acme>
</root>

```

```

    </acme>
</root>

```

Given the problem in the parsed form of having containers with the same name, a rule resolution might consist of the following:

```

Condition: //ADDRESS
Expression: /HOST/text()

```

Effectively, this says: for each address element evaluate `/HOST/text()` to extract the host name as the address identifier.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```

acme
  DESCRIPTION
    SOURCE_ROUTE      yes
    ADDRESS[host1]
      PROTOCOL        tcp
      HOST             host1
      PORT             1630
    ADDRESS_LIST
      FAILOVER         on
      LOAD_BALANCE     off
      ADDRESS[host2a]
        PROTOCOL        tcp
        HOST             host2a
        PORT             1630
      ADDRESS[host2b]
        PROTOCOL        tcp
        HOST             host2b
        PORT             1630
      ADDRESS[host3]
        PROTOCOL        tcp
        HOST             host3
        PORT             1630
    CONNECT_DATA
      SERVICE_NAME     xyz.example.com

```

The rule resolves to an identifier appended in square brackets to the container name. The combination (ADDRESS[host2a], for example) enables various operations such as compare, search, change history, and so forth, to distinguish between address containers.

Sample SQL Query Parsing and Rule Application

Consider the following three-column database table SERVER_DETAILS:

SERVER_NAME	ENVIRONMENT	HOSTED_APPLICATIONS
webserver-100	QA	5
webserver-200	PERFORMANCE	6
webserver-500	PRODUCTION	3

The SQL query expressed as part of the configuration extension creation is as follows:

```
select * from SERVER_DETAILS
```


This query returns the following raw output:

```
[row]
11_SERVER_NAME=13_ webserver-100
11_ENVIRONMENT=2_ QA
19_HOSTED_APPLICATIONS=1_5
[row]
11_SERVER_NAME=13_ webserver-200
11_ENVIRONMENT=11_ PERFORMANCE
19_HOSTED_APPLICATIONS=1_6
[row]
11_SERVER_NAME=13_ webserver-500
11_ENVIRONMENT=10_ PRODUCTION
19_HOSTED_APPLICATIONS=1_3
```

The Configuration Browser Source tab renders the data the same way.

Its parsed form, using the Database Query parser, appears in the user interface in the following tree structure:

```
row
  SERVER_NAME=webserver-100
  ENVIRONMENT=QA
  HOSTED_APPLICATIONS=5
row
  SERVER_NAME=webserver-200
  ENVIRONMENT=PERFORMANCE
  HOSTED_APPLICATIONS=6
row
  SERVER_NAME=webserver-500
  ENVIRONMENT=PRODUCTION
  HOSTED_APPLICATIONS=3
```

Notice that the row containers are indistinguishable. Thus, this query result is a candidate for a post-parsing rule. As mentioned, there is a special internal XML format against which to apply a rule's XPath condition and expression. This format treats nodes and attributes as XML elements, and converts attribute values into corresponding element text content. It also adds a root element that doesn't appear in the original file:

```
<root>
  <row>
    <SERVER_NAME>webserver-100</SERVER_NAME>
    <ENVIRONMENT>QA</ENVIRONMENT>
    <HOSTED_APPLICATIONS>5</HOSTED_APPLICATIONS>
  </row>
  <row>
    <SERVER_NAME>webserver-200</SERVER_NAME>
    <ENVIRONMENT>PERFORMANCE</ENVIRONMENT>
    <HOSTED_APPLICATIONS>6</HOSTED_APPLICATIONS>
  </row>
  <row>
    <SERVER_NAME>webserver-500</SERVER_NAME>
    <ENVIRONMENT>PRODUCTION</ENVIRONMENT>
    <HOSTED_APPLICATIONS>3</HOSTED_APPLICATIONS>
  </row>
</root>
```

Given the problem in the parsed form of having three containers with the same name, a rule resolution might consist of the following:

Condition:/root/row/SERVER_NAME

Expression:SERVER_NAME/text ()

Effectively, this says: for each row evaluate SERVER_NAME/text () to produce an identifier that distinguishes one row from another within the tree structure.

After applying the post-parsing rule, the parsed tree structure appears as follows:

```
row[webserver-100]
  SERVER_NAME=webserver-100
  ENVIRONMENT=QA
  HOSTED_APPLICATIONS=5
row[webserver-200]
  SERVER_NAME=webserver-200
  ENVIRONMENT=PERFORMANCE
  HOSTED_APPLICATIONS=6
row[webserver-500]
  SERVER_NAME=webserver-500
  ENVIRONMENT=PRODUCTION
  HOSTED_APPLICATIONS=3
```

The rule resolves to an identifier appended in square brackets to the container name. The combination (row[webserver-100], for example) enables various operations such as compare, search, change history, and so forth, to distinguish between row containers.

Overview of Relationships

Relationships define the associations that exist among targets, or more extensively, among managed entities. In general, relationships are inherent to a target type definition. But not all relationships can be anticipated at target type creation. Thus, Cloud Control supports creation of supplemental relationships. There are two methods available to create new relationships:

- Manually, by adding a generic system target
- Interactively, within the Configuration Topology Viewer

This section describes the manual process. For information on creating relationships within the Configuration Topology Viewer, see [Creating a Relationship to a Target](#).

There are two ways to access the generic system wizard:

- From the **Setup** menu, select **Add Target**, then select **Generic System**
- From the **Targets** menu, select **Systems**, then click the **Add** button

General

Provide general details of the generic system:

- Specify a meaningful target name
- Indicate whether this is to be a privilege-propagating system
- Set system properties such as cost center and life cycle status
- Add system members; there should a logical correspondence to the selections
- Review member dependencies and indicate whether to include them
- Set the time zone appropriately (defaults to Greenwich Mean Time)

When done, click **Next**.

Define Associations

Select the check box to display any associations (relationships) that Cloud Control automatically detects based on the members added to the system. Add additional associations as follows:

1. Click **Add**.
2. Complete the dialog that opens as follows:
 - a. Select a member target in the left table. This populates the right table.
 - b. Select an associated target in the right table. This populates the association drop-down list.
 - c. Select the association you want to create.
 - d. Click **OK**. The new association appears in the associations table.
3. Click **Add** and repeat to create additional associations.

When done, click **Next**.

Availability Criteria

Use this page to declare the key members of the system; that is, the members that must be running for the system to be considered available. You can select any one, some, or all members, but you must select at least one.

When done, click **Next**.

Charts

Customize how you want charts to appear on the System Charts page:

- Supplement suggested charts with charts you add
- Edit certain suggested charts to fit your needs
- Deselect the suggested charts check box and customize the page entirely
- Alter the appearance of the Members page by adding and removing columns and abbreviations

When done, click **Next**.

Review

Verify the makeup of the generic system target. If everything appears in order, click **Finish**.

Upon confirmation that the target was successfully created, use the Configuration Topology Viewer to review and traverse the relationships you created.

Overview of Configuration Topology Viewer

The Configuration Topology Viewer provides a visual layout of a target's relationships with other targets.

This section covers the following topics:

- [About Configuration Topology Viewer](#)
- [Examples of Using Topology](#)
- [Viewing a Configuration Topology](#)
- [Determining System Component Structure](#)
- [Determining General Status of Target's Configuration Health](#)
- [Getting Configuration Health/Compliance Score of a Target](#)
- [Analyzing a Problem and Viewing a Specific Issue in Detail](#)
- [About Dependency Analysis](#)
- [About Impact Analysis](#)
- [Creating a Custom Topology View](#)
- [Deleting a Custom Topology View](#)
- [Excluding Relationships from a Custom Topology View](#)
- [Including Relationships to a Target in a Custom Topology View](#)
- [Creating a Relationship to a Target](#)
- [Deleting a Relationship from a Target](#)
- [Controlling the Appearance of Information on a Configuration Topology Graph](#)

About Configuration Topology Viewer

The Configuration Topology Viewer provides a visual layout of a target's relationships with other targets. To access the Configuration Topology Viewer from a target's home page, select **Configuration**, then select **Topology** in the dynamic target menu. A topology graph appears for the current target. Using the viewer, you can:

- Determine the source of a target's health problems, that is, detect which targets might be causing the failure. For example, a database is down because its host is down.
- Analyze the impact of a target on other targets. For example, the payroll and finance applications will be impacted if the database goes down.
- Determine the system's structure by viewing the members of a system and their interrelationships.
- Add additional relationships between targets. These relationships will be reflected in other Cloud Control tools.
- Customize your configuration topology views to focus on the targets for which you have responsibility.
- Share custom topology views that you have created with other Cloud Control users.

Examples of Using Topology

The following are examples of when to use the topology feature:

- Determine a system's component structure (see [Determining System Component Structure](#))

- Analyze dependencies in relationships (see [About Dependency Analysis](#))
- Analyze the impact of relationships (see [About Impact Analysis](#))

Viewing a Configuration Topology

The Configuration Topology Viewer provides a visual layout of a target's relationships with other targets.

In the situations where the topology you are viewing is larger than your browser window, you can adjust the view by:

- Clicking the small arrow icon in the bottom right corner of the window to bring up a navigator, which allows you to select which portion of the topology is in view.
- Decreasing the size of the nodes in the display using the zoom control in the top left of the display.

Perform the following steps:

1. Access the Configuration Topology Viewer.

From the **Targets** menu on the Cloud Control home page, select **All Targets**. In the table, click the appropriate target. On the resulting page, select **Configuration** then select **Topology** from the dynamic target menu.

2. From the **View** list, select any of the following:

- Uses

This view helps you determine the targets that the selected target depends on. If a target is having problems, this view can be useful in helping you determine whether its problems have been caused by another target it depends on.

- Used By

This view shows you the targets that depend on the selected target. This can be useful, for example, if you are planning on shutting down the target and need to know what other targets will be affected

- System Members

This view shows the members of the system (available only for targets that are systems).

- Custom views that have been defined and shared by end users (custom views must be explicitly shared before they are available to others).

The Uses, Used By, and System Members views are topology views provided by Oracle. They cannot be modified.

3. The following operations are available on the Topology page:

- Create a custom topology view (see [Creating a Custom Topology View](#))
- Delete a custom topology view (see [Deleting a Custom Topology View](#))
- Exclude relationships from a custom topology view (see [Excluding Relationships from a Custom Topology View](#))
- Include relationships to a target in a custom topology view (see [Including Relationships to a Target in a Custom Topology View](#))

Determining System Component Structure

To determine which components (targets and target components) comprise your IT system and their interrelationships, use the Configuration Topology Viewer.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. In the View menu, select **System Members** (available only if the target is a system). The view displays the relationships between the targets. The target type controls the default view that is shown.

To see the specific relationship between two targets, hover over the link between them and the relationship name will pop up.

Note the following:

- The topology feature is available any time you are in the context of a target: select **Configuration** from the target type menu, then select **Topology**.
- Not all target types have configuration data. For these target types, the Configuration menu and topology graphs are not available.

Determining General Status of Target's Configuration Health

Topology enables you to view system health by displaying relationships among system entities, structure of a target, and target components, thus enabling you to analyze configuration health and status of the configuration.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. In the Uses view on the Configuration Topology Viewer page, icons indicate whether the target is down. You can choose a particular view, for example, Uses or Used By. In addition, icons indicate whether targets have associated incidents.

Getting Configuration Health/Compliance Score of a Target

To determine the configuration health and compliance score of a target, perform the following steps:

1. Access the Configuration Topology Viewer.
2. Zoom in on the target that has problems. Problems are represented by icons indicating a problem target status, and icons indicating target incidents. The target you selected in the All Targets page will always be highlighted.
3. When you click a target, properties for the target are available in the Configuration tab in the Properties section. The Configuration tab shows information about target compliance, configuration changes in the past week, and recommended patches. Links from these values lead to more detailed reports.

If incidents are reported in the Incident Summary tab, resolve the reported events and incidents. Compliance information is available through the Configuration tab. If the target is not compliant, resolve the issue. Also if patches are missing, apply them.

4. Repeat the process of analyzing the various targets until all targets are functioning properly.

Analyzing a Problem and Viewing a Specific Issue in Detail

When you drill-down in topology graphs, you can have a detailed view of the specific issue that could be the cause of the problem.

Perform the following steps:

1. Access the Configuration Topology Viewer.

From the **Targets** menu on the Cloud Control home page, select **All Targets**. In the table, click the appropriate target. On the resulting page, select **Configuration** then select **Topology** from the dynamic target menu.

To view target data, place the mouse over the node and continue to move the mouse to >>. The popup containing data appears. For additional information, select **Properties** located at the right. The links associated with the data lead to the detail pages.
2. View configuration history changes.

From the dynamic target menu, select **Configuration**, then select **History**. On the Configuration History page, determine whether there has been a history change in the last 24 hours. If so, view those changes in detail for that particular target.

Another way to access Configuration Changes from a node is to select the node, click on **Properties**, click the **Configuration** tab, and click the value associated with Configuration Changes.
3. View compliance violations, incidents, and unauthorized changes available from Properties.
4. View critical or warning incidents generated for a particular target: select **Properties**, then select **Incidents**.
5. Determine whether there are patch recommendations.

On the Topology page, select a node. Select **Properties** then select **Configuration**. Click the value associated with Patch Advisory.

About Dependency Analysis

Dependency analysis, also known as root cause analysis, traverses the relationships top to bottom to see if there is cause of a problem due to an issue with an asset on which the item is dependent.

To find the source of a target's health problem, perform the following steps:

1. Access the Configuration Topology Viewer.
2. In the **View** list, select **Uses**. This shows a topology of the targets that the selected target depends on.

Paths to the target or targets *potentially* causing the problem are colored.

If your target is not up, paths to the target or targets that may be causing the problem are colored. Red links lead from your target to targets that are down, and yellow links lead to targets whose status is not known.

By default the topology includes all depths of the tree, including the dependency relationships between those targets.

About Impact Analysis

Impact analysis traverses the relationships from the bottom to the top of the tree to see if a problem will occur if changes are made to the element (target or system) in which I'm interested. It answers the question: What items are dependent on my element that would be effected should I do something to my element. For example, if I shut down a listener, what databases would be affected?

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. On the **Topology** page, analyze the **Used By** view. The topology will show the targets that depend on the selected target.

Creating a Custom Topology View

Create a custom topology view to include only those targets of interest, perhaps for a specific task or report. From a custom view, you can also augment the relationship data provided by Cloud Control.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **Customize** menu, select **Create Custom View....** Provide the name and description for the topology and select one of the Initial Contents:
 - **Copy Current View** to create a topology view similar to the one you are viewing.
 - **Create Empty View** to create a topology view that starts with the root node.

Also, choose one of the following expose options:

- Expose the custom view for all targets of the current target type. For example, if you are creating a topology view for a database target, the new view will be available for all database targets.
- Expose the custom view for the current target only.

To share the view, click **Share this custom view with other users**.

Click **OK**.

3. Reduce the unwanted information in the topology by highlighting the target and selecting **Hide Relationships...** in the **Customize** menu.

You can also display relationships that are not being displayed by selecting a target. From the **Customize** menu, select **Target**, then select **Show More Relationships to Target Type....**

Privileged users can also choose to share their custom views with other users. To share a custom view, select the checkbox labeled **Share this custom view with other users**.

4. Click **OK**.

Deleting a Custom Topology View

When a custom topology view is no longer of use, delete it so it no longer clutters the View list. **Note:** System owned views cannot be deleted.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **View** list, select the topology view you want to delete.
3. From the **Customize** menu, select **Delete Custom View...**
4. Click **Delete Custom View** in the confirmation popup.

Excluding Relationships from a Custom Topology View

After you create a topology view, you may want to remove some of the targets displayed in the custom view. Note that you cannot modify the topology views provided by Oracle: Uses, Used By, and System Members.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **View** list, select the topology view you want to change then select the target.
Note: System created views cannot be modified.
3. From the **Customize** menu, select **Hide Relationships....**
4. The list of relationships that are displayed in the graph are listed in the Hide Relationships page. You can multi-select the relationships to exclude from the graph. Click **OK**.

Including Relationships to a Target in a Custom Topology View

After you create a topology view, you may find it necessary to include more relationships in the custom view. This will add targets to your custom view if they are related to the currently displayed targets using relationships you include.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **View** list, select the custom topology view you own or have the privileges to change. System views such as Uses, Used By, and System Members cannot be modified.
3. Highlight a target from which to expand the topology. From the **Customize** menu, select **Target**, then select **Show More Relationships to Target Type....**
4. The resulting dialog shows a list of the relationships that the selected target type can participate in. Select the relationships of interest, and click **OK**. Any targets that are related to the selected target type using the selected relationships will be added to the topology view.

Creating a Relationship to a Target

In cases where you find that Cloud Control has incomplete information about your systems, you can create relationships between targets.

Note: Once new relationships are created, any topology showing the specified relationships and containing the targets will automatically show the new relationships.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **View** list, select the topology view you want to change then select the target.
3. Select a target in the topology to be one end of the relationship.
4. On the Create Custom View page, provide a name and description, choose the initial contents, and determine how this custom view should be exposed. Click **OK**.
5. From the **Customize** menu, select **Target**, then select **Create Relationship to Target...**
6. On the **Create Relationship to Target** page, select the related target and the relationship between targets. Only relationships that the target type can participate in are shown in the list. Not all target types can be related to each other.



Note:

Created relationships are independent of the view. You can see and use created relationships in other areas of Cloud Control, such as System templates, topology views, and configuration comparisons. Deleting a custom view will not delete the new relationship.

7. On the Confirmation page, click **Create**.

The related target will be added to the view.

Deleting a Relationship from a Target

If you have created a relationship between two targets, you may decide that the relationship no longer exists. Reflect this change by deleting extraneous relationships where appropriate. Note that once relationships are removed, they no longer show in any topology views.

Perform the following steps:

1. Access the Configuration Topology Viewer.
2. From the **View** list, select a custom view.
3. Select the link to the relationship you want to delete. You can either right click the node to view the context menu that allows you to delete the relationship or, from the **Customize** menu, select **Relationship**, then select **Delete Relationship...**
4. On the Confirmation page, click **Delete**.

Relationships are used in various places in Cloud Control, such as System templates, topology views, configuration comparisons, and so on. Deleting a relationship from this topology can impact these other areas.

If you create a relationship, you can later delete it by using the **Delete Relationship...** menu item.

Controlling the Appearance of Information on a Configuration Topology Graph

To control the way the targets are displayed in a custom topology, you can customize the tier in which a target type is shown, and you can group target types together.

The tier in which a target type is shown will affect its vertical or horizontal placement in the topology, depending on whether the layout is left-right or top-down.

To customize the appearance, perform the following steps:

1. Access the Configuration Topology Viewer.
2. Create or select an existing custom view.
3. To control highlighted paths to targets that are down, toggle the "Highlight 'Down' Root Cause" menu item.

When this menu item is selected and the root target is down, paths from the root node to other down targets are highlighted. By visually following the highlighted paths, you may determine which targets are causing the root target's down status.

 Note:

When this option is selected, you will not be able to group nodes together.

4. To manipulate tiers:
 - a. On the **Customize** menu, select **Select Tiers**.
 - b. Select either **Specify Tiers** or **Use Default Tiers**. If you choose to specify tiers, drag target types to their desired tier.
5. To turn the coloring of the links on and off, on the **Customize** menu, select **Highlight "Down" Root Cause**.
6. To group targets:
 - a. Select a link that represents one or more associations between the source and destination target types.
 - b. On the **Customize** menu, select **Relationship**, then select **Group Targets...** (Another way to select group targets is to right mouse click a link and select **Group Targets**.)

All matching associations are placed into group boxes.

 Note:

Grouping targets is not possible when "Highlight 'Down' Root Cause" is selected.

Managing Compliance

Compliance Management provides the ability to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage. This is accomplished by defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules. In addition, Compliance Management provides advice of how to change configuration to bring your targets and systems into compliance.

This chapter explains how Compliance Management verifies that applications in your enterprise comply with pre-established standards and how to manage the compliance structure. This chapter includes:

- [Overview of Compliance](#)
- [Evaluating Compliance](#)
- [Investigating Real-time Observations](#)
- [Configuring Compliance Management](#)
- [Real-time Monitoring Facets](#)
- [Examples](#)

Overview of Compliance

The Compliance Management solution provides the tools to evaluate targets and systems for compliance with business best practices in terms of configuration, security, storage, and so on. In addition, Compliance Management provides the capability to define, customize, and manage the entities used to evaluate compliance.

The compliance solution:

- Automatically determines if targets and systems have valid configuration settings and whether they are exposed to configuration-related vulnerabilities.
- Advises how to change configurations to bring targets and systems into compliance with respect to best practices.
- Provides real-time monitoring of a target's files, processes, and users to let Oracle Enterprise Manager Cloud Control (Cloud Control) users know where configuration change or unauthorized action are taking place in their environment.
- Provides Oracle provided compliance frameworks (for example, Oracle Generic Compliance Framework) and compliance standards to map to compliance standard rules. This mapping makes it possible to visualize how out-of-compliance settings and actions will affect any compliance framework an organization follows.
- Provides a compliance-focused view of IT configuration and change that is suitable for Line of Business Owners, IT Managers, and Compliance Managers to refer to regularly to check on their organization's compliance coverage.

Before you start using the compliance features, there are a few basics you need to know. See the following for details:

- [Terminology Used in Compliance](#)
- [Accessing the Compliance Features](#)
- [Roles and Privileges Needed to Use the Compliance Features](#)

Terminology Used in Compliance

The following terms are used throughout this chapter when discussing the compliance feature:

- **Compliance Framework**

A compliance framework is an organized list of control areas that need to be followed for a company to stay in compliance in their industry. Enterprise Manager uses compliance frameworks as a foldering structure to map standards and rules to the control areas they affect. Compliance frameworks are hierarchical to allow for direct representation of these industry frameworks.

A single framework control area maps to one or more compliance standards. The outcome of these compliance standard evaluations results in a score for the given framework area.

- **Compliance Standard**

A compliance standard is a collection of checks or rules that follow broadly accepted best practices. It is the Cloud Control representation of a compliance control that must be tested against some set of IT infrastructure to determine if the control is being followed. This ensures that IT infrastructure, applications, business services and processes are organized, configured, managed, and monitored properly. A compliance standard evaluation can provide information related to platform compatibility, known issues affecting other customers with similar configurations, security vulnerabilities, patch recommendations, and more. A compliance standard is also used to define where to perform real-time change monitoring.

A compliance standard is mapped to one or more compliance standard rules and is associated to one or more targets which should be evaluated.

- **Compliance Standard Rule**

A compliance standard rule is a specific test to determine if a configuration data change affects compliance. A compliance standard rule is mapped to one or more compliance standards.

Cloud Control has the following types of compliance standard rules.

- **Repository Rule**
Used to perform a check against any metric collection data in the Management Repository
- **WebLogic Server Signature Rule**
Used to check a WebLogic target to support best practice configurations.
- **Real-time Monitoring Rule**
Used to monitor actions to files, processes, and database entities in real-time as the changes occur. Also captures users logging in and logging out, and SU and SUDO activities.
- **Agent-Side Rule**

Used to perform configuration checks on the agent and upload violations into the Management Repository.

- Manual Rule

Checks that must be performed but cannot be automated. For example: "Plans for testing installations, upgrades, and patches must be written and followed prior to production implementation."

- Compliance Standard Rule Folder

Compliance standard rule folders are hierarchical structures that contain compliance standard rules.

- Importance

Importance is a setting that the user can make when mapping compliance frameworks, standards, and rules. The importance is used to calculate the affect a compliance violation will have on the compliance score for that framework control area or compliance standard.

For compliance frameworks, when mapping a compliance standard, the importance for this compliance standard indicates the relative importance to other compliance standards in this framework.

For compliance standards, when mapping a compliance standard rule, importance indicates the relative importance of a compliance standard rule to all other compliance standard rules in the compliance standard.

- Score

A target's compliance score for a compliance standard is used to reflect the degree of the target's conformance with respect to the compliance standard. The compliance score is in the range of 0% to 100% inclusive. A compliance score of 100% indicates that a target fully complies with the compliance standard.

- Real-time Facets

The real-time monitoring rule definition includes facets that specify what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type. For example, the networking configuration files for your operating system could be defined by one facet containing multiple file names or file patterns.

- Real-Time Observations

Observations are the actions that were seen on a host or target that were configured to be monitored through real-time monitoring rules. Each distinct user action results in one observation.

- Observation Audit Status

Every observation has an audit status that determines if the observation was authorized, or unauthorized, or neither (unaudited). The audit status can be set manually or automatically through the real-time monitoring compliance standard rule configuration.

- Observation Bundles

Single observations are not reported from the Management Agent to the server. They are instead bundled with other observations against the same target, rule, and user performing the action. Bundles help combine like observations and make it easier to manage the observations in Cloud Control.

Accessing the Compliance Features

To access the compliance features, navigate to the **Enterprise** menu, select **Compliance**, then select one of the following:

- **Dashboard**

The dashboard provides a very high level view of results that show how compliant or at risk your organization or your area is. The dashboard contains dials representing the compliance score for a selected framework, least compliant systems and targets, and unmanaged discovered hosts.

- **Results**

Compliance results include evaluation results and errors for compliance frameworks and compliance standards, as well as target compliance.

- **Library**

The Compliance Library page contains the entities used for defining standards. From the Compliance Library page you can manipulate compliance frameworks, compliance standards, compliance standard rules, and real-time monitoring facets.

Note: The real-time monitoring facets are only for real-time monitoring rules.

- **Real-time Observations**

Observations are the actions that were seen on a host or target that were configured to be monitored through real-time monitoring rules. Each distinct user action results in one observation. Observations are additionally bundled if there are multiple observations done in a short period of time by the same user on the same target and against the same real-time monitoring rule. Multiple UI-based reports are provided to allow users to analyze the actions that are being observed.

Roles and Privileges Needed to Use the Compliance Features

To use the compliance standard features, you need to have access to the following roles and privileges.

Role	Description	Contains the Privileges
EM_COMPLIANCE_DESIGNER	Enables you to create, modify, and delete compliance frameworks, compliance standards, compliance standard rules, and real-time monitoring facets.	Target Privileges: • Manage Any Target Compliance • Manage Any Target Metric • View any Target Resource Privileges: • Compliance Framework (Create Compliance Entity; Full any Compliance Entity) • Configuration Extensions (Manage Configuration Extensions owned by any user) • Job System (Create)
EM_COMPLIANCE_OFFICER	Enables you to view compliance framework definition and results.	No target privileges Resource Privilege • Compliance Framework (View any Compliance Framework)

The target and resource privileges used in compliance include:

Privilege	Type of Privilege	Included in Role	Description
Manage any Target Compliance	Target	EM_COMPLIANCE_DESIGNER	Allows you to manage the compliance of any target including the association of a compliance standard to a target.
Manage any Target Metric	Target	EM_COMPLIANCE_DESIGNER	Enables you to manage a metric for any target.
View any Target	Target	EM_COMPLIANCE_DESIGNER	Allows you to view all managed targets in Enterprise Manager.
View any Compliance Framework	Target	EM_COMPLIANCE_OFFICER EM_COMPLIANCE_DESIGNER	Allows you to view compliance framework definition and results NOTE: This privilege is part of the Compliance Framework resource privilege. This privilege is granted by default for EM_COMPLIANCE_OFFICER role but it is <i>not</i> granted by default for the EM_COMPLIANCE_DESIGNER role.
Create Compliance Entity	Resource	EM_COMPLIANCE_DESIGNER	Allows you to create compliance frameworks, compliance standards, compliance standard rules, and real-time monitoring facets. This privilege is part of the Compliance Framework resource privilege.
Full any Compliance Entity	Resource	EM_COMPLIANCE_DESIGNER	Allows you to edit and delete compliance frameworks, compliance standards, compliance standard rules, and real-time monitoring facets. This privilege is part of the Compliance Framework resource privilege.
Compliance Framework	Resource	EM_COMPLIANCE_DESIGNER EM_COMPLIANCE_OFFICER	Provides the capability to define, customize, and manage compliance frameworks, compliance standards, and compliance standard rules, and evaluate the compliance of targets and systems with regards to business best practices for configuration, security, storage, and so on. This privilege contains the following privileges: - Create Compliance Entity (granted by default in EM_COMPLIANCE_DESIGNER role) - Full any Compliance Entity (granted by default in EM_COMPLIANCE_DESIGNER role) - View any Compliance Framework (granted by default in EM_COMPLIANCE_OFFICER role)
Configuration Extensions	Resource	EM_COMPLIANCE_DESIGNER	Allows extending target configuration collections. This privilege contains the following privileges: - Manage Configuration Extensions owned by any user (granted by default) - Manage Configuration Extensions owned by the user
Job System	Resource	EM_COMPLIANCE_DESIGNER	Job is a unit of work that may be scheduled that an administrator defines to automate the commonly run tasks. This privilege contains the following privileges: - Create (granted by default) - Manage View Access

The following table lists the compliance tasks with the roles and privileges required.

Task	Roles and Privileges Required
Create compliance framework	Create Compliance Entity privilege View any Compliance Framework privilege
Edit and delete compliance framework	Full any Compliance Entity privilege View any Compliance Framework privilege
Create, edit, and delete compliance framework	EM_COMPLIANCE_DESIGNER role EM_COMPLIANCE_OFFICER role
Associate a compliance standard to a target	Manage any Target Compliance privilege or MANAGE_TARGET_COMPLIANCE privilege on the target
Import or export a compliance framework	EM_COMPLIANCE_DESIGNER role EM_COMPLIANCE_OFFICER role
Create a real-time monitoring rule	EM_COMPLIANCE_DESIGNER role
Create a real-time monitoring facet	EM_COMPLIANCE_DESIGNER role

Note: In addition, ensure you have privileges to access the target you will be associating with a compliance standard. In particular, you need the Manage any Target Compliance privilege on the target.

Evaluating Compliance

Compliance evaluation is the process of testing the compliance standard rules mapped to a compliance standard against a target and recording any violations in the Management Repository.

By evaluating a target against a compliance standard, you are determining whether a target complies with the checks of the standard. In the case when a target does not meet the desired state, the test may suggest what changes are required to make that target compliant.

Compliance evaluation generates a score for a target based on how much the target is compliant with the standard. A 100% compliance score means that all checks of the compliance standard passed on the target. For real-time monitoring, the compliance score will drop as you have observations that have been marked as unauthorized either manually or through change request management integration. As these unauthorized observations are either cleared or changed to authorized, the score will improve.

Because target compliance is required to be monitored regularly, you need to associate a compliance standard with targets. Evaluation is automatically performed for any associated targets, when the target state refreshes, that is when new data has been collected from the target. For repository rules, when new data for the target gets loaded into the Management Repository, evaluation happens again. For Real-time Monitoring, evaluation happens every time an observation of a user action is seen.

What You Can Do To Ensure Compliance

When using Cloud Control to evaluate your compliance, you should regularly perform the following actions:

- Regularly monitor the compliance dashboard to find areas that may indicate your organization has a low compliance score or is at risk.
- View the results of an evaluation

Study the results of the evaluations and make the needed changes to the targets

Only results from the targets for which you have View privilege will be available. The compliance standard rule evaluation results are rolled up in order to produce a compliance standard evaluation state as well as a compliance summary.

- Study Oracle provided reports

Regularly monitor real-time monitoring observation UI reports to see if detected observations are normal or abnormal. Set abnormal observations to unauthorized until any unauthorized change can be reverted or until the actions can be investigated to the level required by your auditors.

- Study the trend overview as a result of the evaluation

Use the graphs in the Trend Overview pages to visually determine whether the targets are adhering to or distancing themselves from the compliance best practices.

To access the Trend Overview pages for compliance standards:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. From the **Compliance Standards** tab, choose **Evaluation Results**.
3. On the Evaluation Results page, choose the compliance standard you want to investigate and click **Show Details**.
4. On the resulting details page, click the **Trend Overview** tab.

Note: You can also review Trend Overview pages for compliance frameworks.

- Ensure your environments match baselines (or each other) by creating rules on top of configuration compare capabilities. Then monitor for configuration drift using real-time monitoring.
- Evaluate validity of configuration settings.
- Evaluate exposure to configuration-related vulnerabilities, storage, and security.
- Modify targets and systems to be compliant.
- Verify authorization of configuration changes or user actions.
- Continually test your systems, services, and targets, ensuring the best possible protection and performance your system can have.
- Use Oracle provided compliance standards and compliance standard rules to determine compliance.
- Keep an eye on hosts in your environment that are not monitored for compliance as these introduce a large amount of compliance risk in your environment.

The following sections provide additional details:

- [Accessing Compliance Statistics](#)
- [Viewing Compliance Summary Information](#)
- [Viewing Target Compliance Evaluation Results](#)
- [Viewing Compliance Framework Evaluation Results](#)
- [Managing Violations](#)

- [Investigating Compliance Violations and Evaluation Results](#)
- [Investigating Evaluation Errors](#)
- [Analyzing Compliance Reports](#)
- [Overview of Compliance Score and Importance](#)

Accessing Compliance Statistics

Compliance statistics are available throughout the interface in Compliance Summary regions located on pages such as the Compliance Dashboard, the Enterprise Summary page, and a target's home page.

These regions report the violations and compliance scores for the particular targets. However, the region only reports that there is a violation; it does not give the details. For example, a violation can be against the Secure Port compliance standard rule that is part of the Secure Configuration for Host compliance standard. But you will not know the details just by looking at the Compliance Summary regions.

Using the Compliance Dashboard Effectively

The compliance dashboard is a top level view of the Cloud Control compliance features. The dashboard includes several regions which give you a very good insight into how compliant your IT environment is according to the standards you have configured.

To access the Compliance Dashboard:

1. From the Enterprise menu, select **Compliance**.
2. Select **Dashboard**.

The Compliance Dashboard is also one of the pages available from the "Select Your Home" page and can be set as your home page when you log in to Cloud Control.

The Compliance Dashboard includes the following regions:

- Compliance Framework Summary

This region lets the user choose one Compliance Framework and it shows the compliance score for each second-level folder under that Compliance Framework. The needle on the dial shows the current compliance score for that given framework element. The score is based on the targets that the logged-in Enterprise Manager user is allowed to see.

Clicking on the dial will take you to the Compliance Results page for the given second-level framework folder giving you more details on the next framework folders down and/or the compliance standards belonging to this folder.

- Compliance Summary

This region has a view for frameworks and a view for standards. In the Framework view, this region shows you the list of all defined compliance frameworks and their overall score and violation details. In the standard view, this region will list the worst scoring compliance standards along with their violation details. Clicking on a framework or standard name will take you to a screen showing you more details of that framework or standard.

From this region, you can also click on the View Trends link to see a historic trend graph of the compliance score

- Least Compliant Generic Systems

This region shows the generic systems that have the lowest compliance score. The score for a given system is calculated by including all rules that are associated with all elements of that system. A generic system is used to define your IT Business Applications, such as HRIS, Payroll, and so on. Reporting these systems that have the lowest score can help identify which business units have compliance risk leading up to audit time.

- Most Recent Discovered Unmanaged Hosts

This region shows hosts that have been discovered recently using the Cloud Control automatic host discovery feature that have not been promoted to managed hosts. These hosts represent a specific compliance risk in that unmanaged hosts in an IT environment can lead to many access control and data access risks. The intent of this region is to highlight the hosts that have recently been discovered but may not be under compliance control.

- Least Compliant Targets

This region is similar to the Least Compliance Generic Systems except it shows you all targets (including the generic systems again). This region is less useful for an IT management or auditor perspective since it may not be clear what these individual targets are used for. It however can be used as another data point to find the areas where you are at highest risk leading up to an IT compliance audit.

Viewing Compliance Summary Information

Compliance summary information is available from the Cloud Control Compliance Results page and individual target home pages.

To view compliance summary information from the Cloud Control home page, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.

To view compliance summary information from a target's home page, follow these steps:

1. From the **Targets** menu, select the target type, and click the target.
2. On the target's home page, select the target menu located at the top-left of the page.
3. Select **Compliance**, then select **Results**. On the Results page, click **Target Compliance**.

Viewing Target Compliance Evaluation Results

Target-specific compliance evaluation results are available on the Cloud Control home page and individual target home pages. By evaluating compliance rules and standards, the possible evaluation results will be:

Evaluation Results	Description
Compliant	Target meets the desired state and there are no unauthorized real-time monitoring observations.

Evaluation Results	Description
Non-Compliant	Target does not meet the desired state. At least one test in the compliance standard detected a deviation from the desired state or there is at least one unauthorized real-time monitoring observation.
Error	No results returned due to an error. The error may be an unexpected internal error or an error in the test. Examples of errors in the test include attempts to: • Divide by zero • Invoke a function with incorrect parameter values

To view results using Cloud Control home page, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Target Compliance** tab. The page displays the targets with their Average Compliance Score.

To view compliance evaluation results from a target's home page, follow these steps:

1. From the **Targets** menu, select the target type.
2. Click the name of the target in which you are interested.
3. On the target menu located at the top-left of the page, select **Compliance**, then select **Results**.
4. Click the **Target Compliance** tab. The page displays the targets with the Average Compliance Score.

Use the page or region to get a comprehensive view about a target in regards to compliance over a period of time. Using the tables and graphs, you can easily watch for trends in progress and changes.

Note: Trend overview data might take up to six hours after initial compliance standard to target association to display in the time series charts.

Viewing Compliance Framework Evaluation Results

To effectively use a compliance framework, organize the frameworks to reflect the compliance framework control areas you use in your organization. The hierarchical structure of the framework should map directly to the control areas of the frameworks you follow.

Oracle provides a number of frameworks, for example, Oracle Generic Compliance, Fusion Applications Compliance, and Security Technical Implementation Guide (STIG). These frameworks can be used as a starting point for you to create your own frameworks to match your needs or can be used to understand how best to organize your own frameworks based on internal standards or based on SOX, HIPAA, NIST-800, or other common frameworks.

To view the results of a compliance framework evaluation, use the Evaluations Results page accessed through the Compliance Frameworks tab.

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. On the Compliance Results page, click the **Compliance Frameworks** tab and highlight the compliance framework of interest.

Since compliance frameworks are a hierarchical structure, each folder or node of the framework will have its own score. The bottom most children of the hierarchy will have their score roll up to the parent folder and so on. If one person viewing these reports is primarily interested in one control area of the framework they follow, they can focus on the score for that specific control area as represented by the folder they look at under the framework.

Managing Violations

Using the Compliance Results feature you can suppress and unsuppress violations, as well as clear manual violations.

Suppressing a violation enables you to acknowledge an existing violation while removing the violation from the compliance score calculation. Suppressing a violation prevents the violation from negatively impacting the compliance score but not delete it from the list of violations. Suppression can be indefinite or for a specified period of time.

Unsuppressing a violation causes the compliance score to be recomputed accounting for the violations that were unsuppressed.

Clearing of manual rule violations causes the violations to be cleared, and the compliance score to go up for the corresponding compliance standard or target. Clearing a manual rule violation can be indefinite or for a specified period or time.

Accessing the Managing Violations Feature

To access Managing Violation feature:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Highlight a compliance standard and click **Manage Violations**.

The following tabs are available:

- Unsuppressed Violations
- Suppressed Violations
- Manual Rule Violations

Unsuppressed Violations Tab

Use this tab to suppress violations.

1. Select one or more violations.
2. Click **Suppress Violations**.
3. On the Violation Suppressed Confirmation popup, you can suppress the violation indefinitely or provide a date by which the suppression will end. Optionally, you can provide an explanation for the suppression.
4. Click **OK**.

This submits a job to do the suppression asynchronously and returns you to the Result Library page. A suppression adds an annotation to the underlying event stating that the violation is suppressed along with the reason (if a reason was provided). **Note:** The job results are not instantaneous. It may take a few minutes for the results to be displayed.

Suppressed Violations Tab

Use this tab to unsuppress violations.

1. Select one or more violations.
2. Click **Unsuppress Violations**.
3. On the Violation Unsuppressed Confirmation popup, you can provide an explanation for the unsuppression.
4. Click **OK**.

This submits a job to do the unsuppression asynchronously and returns you to the result library. An unsuppression adds an annotation to the underlying event that the violation is unsuppressed along with the reason (if a reason was provided). **Note:** The job results are not instantaneous. It may take a few minutes for the results to be displayed.

Manual Rule Violations Tab

To clear a manual rule violation:

1. Select one or more manual rule violations.
2. Click **Clear Violations**.
3. On the Clear Violations Confirmation popup, you can clear the violation indefinitely or provide a date by which the clear will end. Optionally, you can provide an explanation for the clear.
4. Click **OK**.

This submits a job to do the manual rule violations clearing asynchronously and returns you to the Result Library page. Clearing manual rule violations also clears the underlying violation event. **Note:** The job results are not instantaneous. It may take a few minutes for the results to be displayed.

Investigating Compliance Violations and Evaluation Results

Here are a few suggestions for investigating compliance violations. Attend to the most critical violations or those that have the biggest impact on your overall IT enterprise compliance.

- Monitor the compliance framework scores along with the systems and targets that have the lowest scores on the compliance dashboard.
- Ensure that recently discovered hosts are either being monitored using Cloud control for compliance risk or are not possibly introducing risk in your IT compliance.
- Study the statistics on the Enterprise Summary Home page. In particular, look at the statistics in the Compliance Summary region. The compliance violations with "Critical" severity should be dealt with first.
- Address generic systems (IT business applications) and targets that have the lowest compliance scores.
- For the compliance violations of a particular target, examine the home page for that target. The Compliance Standard Summary region provides overview information, but it also gives you access to the Trend for that target.
- Review compliance violation-related events in the Incident Management area of Cloud Control.

- Navigate to the Results page for a particular compliance standard. In the navigation tree, click the name of the compliance standard and a summary page lists all the targets along with the number of violations.
- Navigate to the Trend Overview page to see charts relating to the number of targets evaluated, the average violation count per target, number of targets by compliance score, and the average compliance score.

 Note:

Only results from those targets for which you have View privilege will be available for viewing.

Investigating Violations of Repository Compliance Standard Rules and Targets Causing Violations

If you are looking at the Enterprise Summary page and you notice that there are critical violations against the Secure Configuration for Host compliance standard, you need to find what targets are causing the violations. Follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. In the Evaluations Results tab for Compliance Standards, highlight the Secure Configuration for Host compliance standard. Click **Show Details**.
3. In the Summary tab on the Compliance Standard Result Detail page, you can look at the results either by target or compliance standard rule. For this example, we will use Result by Compliance Standard Rule.
4. In the navigational list, click the **Secure Ports** compliance standard rule. In the resulting Secure Ports Summary tab, you will get a list of all the targets that are violating the Secure Ports rule. This is a security issue that needs to be addressed.

Viewing All the Violations Reported for Your Enterprise

If you want to see all the targets that are not compliant with the compliance standards:

- From the Enterprise menu, select **Compliance**, then select **Results**.
You have the option of viewing violations associated with compliance standards and compliance frameworks.
 - Click the **Target Compliance** tab for a roll-up view of all violations across all targets, that is, all those targets that are out of compliance.
 - Click the **Compliance Standards** tab to view the list of compliance standards against which there are violations. From this tab, you can also access the Errors tab to view the errors against the compliance standard.
- Navigate to the Home page for a particular target. The Compliance Standard Summary region lists the compliance violations according to severity level. Click the name of the compliance standard of interest to view the details of the violations.

Examples of Viewing Violations

As noted in the previous sections, the compliance feature provides violation details that help you resolve compliance issues. There are a number of ways to access violation details.

Violations are available from the following:

- **Compliance Summary** region located on the Enterprise Summary page.
You can easily see the violations against compliance frameworks and compliance standards.
- Compliance Results page. From the **Enterprise** menu, select **Compliance**, then select **Results**.

The following are examples of how to find violation details.

Example 1 - Accessing Violation Details of a Compliance Framework

To see the violations of a compliance framework, click the **Compliance Frameworks** tab then the **Evaluation Results** tab. The Violations columns list how many violations exist for each framework. When you click the number in a Violations column, all the targets with their associated compliance standards are listed.

In turn, when you click the number in the Violation Count column, the resulting Violations page lists the compliance standard rule that is violated. Again when you click the number in the Violation Count column, the resulting Violation Details page lists all metrics for a particular compliance standard rule that are responsible for the violations.

Example 2 - Accessing Violation Details of a Compliance Standard

When you click the **Compliance Standards** tab then the **Evaluation Results** tab, the Violations columns report how many violations exist for each compliance standard.

When you click the number in a Violations column, the Violations pop-up appears listing all the targets violating the standard.

Figure 50-1 Violations for a Compliance Standard

Violations
Compliance Standard: Security Recommendations For Oracle Products

Target Name	Violation Count
target1.mycompany.com	1
target2.mycompany.com	1

Background Table: Target Evaluations and Violations

Target	Target Evaluations	Violations
target1.mycompany.com	2	0
target2.mycompany.com	0	0
target3.mycompany.com	0	0
target4.mycompany.com	0	1
target5.mycompany.com	0	0
target6.mycompany.com	0	0
target7.mycompany.com	0	0
target8.mycompany.com	0	0
target9.mycompany.com	0	0
target10.mycompany.com	0	0
target11.mycompany.com	0	0
target12.mycompany.com	0	0
target13.mycompany.com	0	0
target14.mycompany.com	0	0
target15.mycompany.com	0	0
target16.mycompany.com	0	0
target17.mycompany.com	0	0
target18.mycompany.com	0	0
target19.mycompany.com	0	0
target20.mycompany.com	0	0
target21.mycompany.com	0	0
target22.mycompany.com	0	0
target23.mycompany.com	0	0
target24.mycompany.com	0	0
target25.mycompany.com	0	0
target26.mycompany.com	0	0
target27.mycompany.com	0	0
target28.mycompany.com	0	0
target29.mycompany.com	0	0
target30.mycompany.com	0	0
target31.mycompany.com	0	0
target32.mycompany.com	0	0
target33.mycompany.com	0	0
target34.mycompany.com	0	0
target35.mycompany.com	0	0
target36.mycompany.com	0	0
target37.mycompany.com	0	0
target38.mycompany.com	0	0
target39.mycompany.com	0	0
target40.mycompany.com	0	0
target41.mycompany.com	0	0
target42.mycompany.com	0	0
target43.mycompany.com	0	0
target44.mycompany.com	0	0
target45.mycompany.com	0	0
target46.mycompany.com	0	0
target47.mycompany.com	0	0
target48.mycompany.com	0	0
target49.mycompany.com	0	0
target50.mycompany.com	0	0
target51.mycompany.com	0	0
target52.mycompany.com	0	0
target53.mycompany.com	0	0
target54.mycompany.com	0	0
target55.mycompany.com	0	0
target56.mycompany.com	0	0
target57.mycompany.com	0	0
target58.mycompany.com	0	0
target59.mycompany.com	0	0
target60.mycompany.com	0	0
target61.mycompany.com	0	0
target62.mycompany.com	0	0
target63.mycompany.com	0	0
target64.mycompany.com	0	0
target65.mycompany.com	0	0
target66.mycompany.com	0	0
target67.mycompany.com	0	0
target68.mycompany.com	0	0
target69.mycompany.com	0	0
target70.mycompany.com	0	0
target71.mycompany.com	0	0
target72.mycompany.com	0	0
target73.mycompany.com	0	0
target74.mycompany.com	0	0
target75.mycompany.com	0	0
target76.mycompany.com	0	0
target77.mycompany.com	0	0
target78.mycompany.com	0	0
target79.mycompany.com	0	0
target80.mycompany.com	0	0
target81.mycompany.com	0	0
target82.mycompany.com	0	0
target83.mycompany.com	0	0
target84.mycompany.com	0	0
target85.mycompany.com	0	0
target86.mycompany.com	0	0
target87.mycompany.com	0	0
target88.mycompany.com	0	0
target89.mycompany.com	0	0
target90.mycompany.com	0	0
target91.mycompany.com	0	0
target92.mycompany.com	0	0
target93.mycompany.com	0	0
target94.mycompany.com	0	0
target95.mycompany.com	0	0
target96.mycompany.com	0	0
target97.mycompany.com	0	0
target98.mycompany.com	0	0
target99.mycompany.com	0	0
target100.mycompany.com	0	0

Again, click the number in the Violation Count column and the Violations pop-up appears. All the Compliance Standard Rules, for example Security Recommendations, are listed.

You continue the process by clicking the number in the Violation Count column again in the Violations pop-up. The subsequent pop-up displays the Violations Details. For example, the Violations Details pop-up displays the name of the patch that is causing the problem.

Example 3 - Accessing Violations of a Target

When you click the **Target Compliance** tab, the Violations columns report how many violations exist for each target.

When you click the number in a Violations column, the Violations pop-up appears listing all the targets violating the standard. See [Figure 50-2](#).

Figure 50-2 Violations Using the Target Compliance Tab

The screenshot shows the 'Compliance Results' interface with the 'Target Compliance' tab selected. A table on the left lists targets s1 through s11. Target s6 is highlighted. A 'Violations' pop-up window is open, displaying a table with two columns: 'Compliance Standards' and 'Violation Count'. The first row shows 'Test1' with a count of 64. The second row shows 'Security Recommendations For Oracle Products' with a count of 1. A mouse cursor is clicking on the number 1 in the Violation Count column.

Again, click the number in the Violation Count column and the Violations pop-up appears. All the Compliance Standard Rules, for example Security Ports, are listed.

You continue the process by clicking the number in the Violation Count column again in the Violations pop-up. The subsequent pop-up displays Violations Details. For example, the Violations Details pop-up displays the numbers of the ports violating the compliance standard.

Example 4 - Violations Using Show Details on Compliance Standards Page

You can also drill-down on violations using the Show Details option on the Compliance Results page. Highlight a standard and click **Show Details**. See [Figure 50-3](#).

Figure 50-3 Show Details Page

The screenshot shows the 'Compliance Results' interface with the 'Compliance Standards' tab selected. A table lists compliance standards. The 'Security Recommendations For Oracle Products' standard is highlighted. A 'Show Details' button is visible next to the standard name.

Compliance Standards	Applicable To	Compliance Standard State	Target Evaluations	Violations
Security Recommendations For Oracle Products	Host	Production	2 0 3	2 0 0

On the resulting page, you have the option of seeing violations by target or by compliance standard rule.

When you click the **Violations** tab, details regarding the compliance standard are listed including Event Details and Guided Resolution. See [Figure 50-4](#).

Figure 50-4 Event Details and Guided Resolution

Security Recommendations For Oracle Products (Compliance Standards)

Summary Trend Overview **Violations**

This table lists information about events/violations of this compliance standard. Select an individual event/violation to view a detailed impact statement as well as recommended actions for quick remediation.

[Search](#) [Advanced](#)

Rule	Target Name	Applicable To	Severity	Keywords	Recommendation
Security Recommendation: s1	Host	Host	Critical	Configuration,Security	Apply one of the identified security patches to the corresponding target in
Security Recommendation: s2	Host	Host	Critical	Configuration,Security	Apply one of the identified security patches to the corresponding target in

The target t1 in host s1 is vulnerable. The security patch 14,038,787 is applic... [Edit](#)

General My Oracle Support Knowledge Updates History

Event Details

- Root Compliance Standard: Security Recommendations For Oracle Products
- Root Compliance Standard Author: ORACLE
- Root Compliance Standard Version: 1
- Rule Name: Security Recommendations
- Rule Type: Repository
- Target: s1 (Host) ⓘ
- Event Reported: Jul 17, 2012 7:05:46 PM GMT

Guided Resolution

Recommendations
Apply one of the identified security patches to the corresponding target in your host.

Diagnostics
View topology
View recent configuration changes

Actions
Disable rule for this target

☒ This event will be automatically cleared when the underlying issue is resolved.

☒ To stop repeat notifications, create an incident for this event and acknowledge it.

Example 5: Accessing Violations from Enterprise Summary Page

When you click the name of a compliance standard in the Compliance Summary region of the Enterprise Summary page, the Compliance Standard Result Detail page appears. By clicking the Violations tab, you can view all the targets that violate the particular compliance standard. See [Figure 50-5](#).

Figure 50-5 Compliance Summary Region on Enterprise Summary Page

Compliance Summary

Compliance Frameworks **Compliance Standards**

[View Trends](#)

Name	Target Evaluations			Violations			Average Compliance Score (%)
	✖	⚠	✔	✖	⚠	✔	
Test1	0	1	0	64	2	0	67
Security Recommendations For Oracle Products	2	0	3	2	0	0	80
GaylenTest	0	0	1	0	0	1	99
GaylenTest	0	0	6	0	0	6	99

On the Compliance Standard Result Detail page, when you click the Summary tab then the Result By Target tab, the number of violations against the target display.

When you click a number in the violations columns, the Violations pop-up appears listing the compliance standard rules that are causing the violation. In turn, when you click the number in the Violation Count column, the name of the offending metric or patch displays.

Note: Similar drill-downs are available from the Target Compliance tab.

Tip: To get to the end result of a Violation, continue clicking the number in the Violation Count column. More and more details are presented, narrowing the cause of the problem.

Note: Target evaluations are only one level while Violations are multi-level.

Investigating Evaluation Errors

The Evaluation Errors page reports statistics about the problems encountered during the evaluation. On initial display, the Evaluation Errors page shows all the evaluation errors.

- Use the Evaluation Errors page to view the errors that occurred as a result of metric collection, as well as those that occurred during the last evaluation.
- Use the search filter to view only those evaluation errors that meet a set of search criteria that you specify.
- Click the message in the Message column to decide what your course of action should be to resolve the error.
- Normally the results of an evaluation overwrite the previous evaluation's results. However, in the case of evaluation failure or data provider collection failure, the previous results are left untouched.

After the underlying problem is fixed, the error will no longer be reported.

Search Filter for Evaluation Errors

By default, all the evaluation errors in your enterprise configuration appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the evaluation errors that meet those criteria in the results table.

For example, if you choose Host in the Target Type list, contains in the Target Name list, and "-sun" in the adjacent Target Name text field, and then click **Go**, Cloud Control displays, in the results table, only the compliance standard rule evaluation errors for the hosts that contain "-sun" in their names.

Analyzing Compliance Reports

Cloud Control provides reports specific to compliance. To access these reports:

1. From the **Enterprise** menu, select **Reports**, then select **Information Publisher Reports**.
2. Scroll to the Compliance section

Compliance reports include the following:

- Descriptions reports

The Descriptions reports list all the available compliance standards, compliance frameworks, and compliance standard rules available in the Compliance Library.

These reports enable you to decide whether additional compliance standards and compliance frameworks need to be defined for your enterprise to attain and maintain its compliance to the standards.

- Results reports

The Results reports provide details of the various evaluations against compliance standards and compliance frameworks. Using the Results reports you can view, in one place, all the statistics regarding the compliance of your enterprise against the defined standards. To view the target that is most likely in need of your immediate attention, view the Target with Lowest AVG COMPLIANCE SCORE report. The following are examples of the reports provided:

- Compliance Standard Results Details

Displays the compliance summary for all the compliance standards evaluated against a target. Data includes compliance score, compliant and non-compliant rules, violations, and last evaluation date.

- Compliance Standard Result Summary

Displays the compliance summary of a particular compliance standard. For example, if there are three targets each reporting on Security Recommendations for Oracle Products compliance, the Result Summary rolls up the information into one report. Data includes average compliance score, the number of targets that need immediate attention, and the number of rules that are non-compliant.

Cloud control also provides a set of reports using the BI Publisher integration. The following reports are available:

- Real-time Monitoring Violation Report

Shows current violations based on real-time monitoring rule type.

- Compliance Summary Report

Shows current compliance score, compliance trends, top 10 least compliant system targets and framework violation summary for a specific Compliance framework and all second-level framework folders.

- Observation Journal Report

Tabular report showing observations that have occurred over a period of time. The user can choose which targets and the start and end time for the report.

Note: To enable BI Publisher reports that include Compliance Frameworks to function, the user running the reports must have the EM_COMPLIANCE_OFFICER role.

Overview of Compliance Score and Importance

A target's compliance score for a compliance standard is used to reflect the degree of the target's conformance with respect to compliance standard. The compliance score is in the range of 0% to 100% inclusive. A compliance score of 100% indicates a target fully complies with the compliance standard.

During an evaluation, a target is found to be compliant or non-compliant with that compliance standard.

Types of Importance

Importance is a setting that the user can make when mapping compliance frameworks, standards, and rules. The importance is used to calculate the affect a compliance violation will have on the compliance score for that framework control area or compliance standard.

For compliance frameworks, when mapping a compliance standard, the importance for this compliance standard indicates the relative importance to other compliance standards in this framework.

For compliance standards, when mapping a compliance standard rule, importance indicates the relative importance of a compliance standard rule to all other compliance standard rules in the compliance standard.

However, just because a compliance standard rule has an importance of 'low' does not mean that it can safely be ignored. All compliance violations should be triaged and cleared once the risk has been removed through a fix or a compensating control.

Importance is used to weight compliance scores as they roll up in a compliance standard hierarchy.

The following sections provide examples of how the compliance score is calculated.

Compliance Score of a Compliance Standard Rule -Target

Note: This calculation is used for WebLogic Server Signature rules and Repository rules.

Compliance score of a compliance standard rule-target is calculated by taking the severity and importance of the compliance standard rule and multiplying the result by the total number of violations divided by the total number of rows evaluated for that target.

The formula is:

$$\text{hirange} - (\text{hirange} - \text{lorange}) * (\text{number of violations} / \text{number of rows evaluated})$$

The following table provides the combination of the severity and importance values used to calculate a compliance score.

Table 50-1 Importance and Severity Ranges

Importance	Critical Severity (1)	Warning Severity (1)	Minor Warning Severity (1)
High	0-25 (2)	66-75	95-96
Normal	26-50	76-85	97-98
Low	51-75	86-95	99-99

(1) low range and high range of the severity

(2) 0 is the lorange; 25 is the hirange

Real-time Monitoring Rule Compliance Score

The compliance score of a real-time monitoring rule is based on the number of observation bundles that have violations compared to how many observation bundles there have been over time. An observation bundle is a collection of all observations that happen over a short period of time (few minutes) by the same user against the same target. For instance, if user A is logged into a host and makes 10 file changes in 5 mi-

notes. These 10 observations will all belong to the same bundle. The bundling is handled automatically by Enterprise Manager.

When calculating the count of past observation bundles, the most recent bundles are weighted higher and they have a different weighting as they get older.

The score is calculated using the formula:

$$1 - V/T$$

where T is the sum of all the weighted bundle counts
and V is the count of the current bundles in violation

The result of the calculation of $1 - V/T$ will be a number around 1 as V is 0 (100% compliant) or will be a number near 0 when V is close to the value of T (0% compliant).

Compliance Score of a Compliance Standard for a Target

The compliance score of a compliance standard for each target is calculated by taking the individual compliance score of each rule - target and multiplying it by its importance. This multiplication is repeated for each rule then the resulting products are added. The sum of the products is then divided by the sum of the importance of each rule.

Figure 50-6 How Compliance Score of a Compliance Standard-Target Is Calculated

Key:

CS: compliance standard

Rule: compliance standard rule. There are 3 rules: Rule1, Rule2, and Rule3.

i: importance

i1: importance for Rule1

i2: importance for Rule2


i3: importance for Rule3

S: compliance score of the rule

S1: compliance score for rule1-target

S2: compliance score for rule2-target

S3: compliance score for rule3-target



Importance	Values
High	3
Normal	2
Low	1

$$\text{Compliance Score of Compliance Standard-Target} = \frac{(S_1 \times i_1) + (S_2 \times i_2) + (S_3 \times i_3)}{(i_1 + i_2 + i_3)}$$

Compliance Framework Compliance Score


The compliance framework score is a rolled up weighted average of all compliance standard-target scores across all compliance standards within the compliance frame-

work hierarchy. The weight is based on the importance of a compliance standard. In the following graphic, compliance framework CF has 2 standards CS1 and CS2. CS1 is associated and evaluated on targets t1 and t2 and CS2 is associated and evaluated on targets t3 and t4.

Figure 50-7 How Compliance Score of a Compliance Framework Is Calculated

Key:

CF: compliance framework
CS: compliance standard
CS₁: compliance standard 1
CS₂: compliance standard 2
t: target
i: importance
i_{cs1}: importance of CS1
i_{cs2}: importance of CS2
ST: compliance score of a compliance standard for a target
ST₁: compliance standard score for CS1-t1
ST₂: compliance standard score for CS1-t2
ST₃: compliance standard score for CS2-t3
ST₄: compliance standard score for CS2-t4



Importance	Values
High	3
Normal	2
Low	1

$$\text{Compliance Score of Compliance Framework} = \frac{(ST_1 \times i_{cs1}) + (ST_2 \times i_{cs1}) + (ST_3 \times i_{cs2}) + (ST_4 \times i_{cs2})}{(i_{cs1} + i_{cs1} + i_{cs2} + i_{cs2})}$$

Parent Node Compliance Score

The compliance score of a hierarchy node/parent node is calculated as shown in [Figure 50-8](#). Compliance standards are hierarchical, thus the top node in the tree is known as the parent node.

Figure 50-8 Compliance Score of Parent Node

$$\text{Compliance Score of Parent} = \frac{\sum_{\forall i} S_i \times I_i}{\sum_{\forall i} I_i}$$

In [Figure 50-8](#):

- i represents the number of children

- S is the score of the child node
- I is the importance of the child node

Investigating Real-time Observations

As previously described, observations are the actions that were seen on a host or target that were configured to be monitored through real-time monitoring rules. Each distinct user action results in one observation.

Observations can have one of many audit statuses. The basic audit status "unaudited" means that the observation was detected, there just is no indication that this action was good or bad. The authorized status means that some review has happened for the observation and it should be treated as expected to occur (it was a good change). The unauthorized status means that this observation has been reviewed and has been found to be against policy. This may result in either a corrective fix, a change to policy, or a compensating control being put in place. The audit status for observations can be automatically set by a rule so that all observations triggered by the rule get a default audit status. The status can also be set manually through the UI reports discussed below. The most advanced capability involves integrating with a Change Management Request server through a Cloud Control connector to automatically determine on a per-observation basis if that action was supposed to happen.

The following sections provide additional details regarding real-time monitoring observations:

- [Viewing Observations](#)
- [Operations on Observations During Compliance Evaluation](#)

Viewing Observations

There are four key ways to see what real-time monitoring observations have occurred in your environment:

- [Viewing Observations By Systems](#)
- [Viewing Observations By Compliance Framework](#)
- [Viewing Observations By Search](#)
- [Viewing Details of an Incident](#)

The first three observation screens are available from the Enterprise menu by selecting **Compliance**, then selecting **Real-time Observations**. This page that lets you choose which of the three reports to look at and also shows any Management Agent warnings related to configuration of Real-time monitoring rule configuration. These warnings are reported from the Management Agent and could impact observations from being delivered to the Cloud Control server. If you are missing observations that are expected, review these warnings and address any configuration issues that is causing them.

Viewing Observations By Systems

When observations occur, they can be marked as authorized or unauthorized automatically. This provides one way you to find observations that are important for you to look into. However, if a rule is not configured to reconcile observations with a change management server, it can be difficult to find the observations that are important to you

through only an attribute search. Being able to view observations by business application (generic systems) and drilling down into observation details allows you to discover where there may be issues that should be investigated regardless of the observation's audit status.

Typically, IT managers and line of business owners must identify when unwanted configuration drift occurs in their business applications. By browsing observations by systems, you can easily see which changes affect specific business applications. Observations can be filtered by whether they are authorized, unauthorized, unaudited or both. They can also be filtered by time.

This begins with you choosing one or more business applications and being able to see the relative counts of observations. This report starts at the business application level (generic systems) because an IT manager and compliance auditor may not know what a target is used for. A business application is modeled in Cloud Control as a generic system.

If you are more technical, you still may want to start at this business application level if this is the business application you are working on.

To view observations by systems, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Real-time Observations**.
2. Click **Browse Observations by System Targets**.

Cloud Control displays the Select Root Target(s) page that lists the Target Name for each system target. There is also a link for all targets not belonging to a system target.

3. You can begin viewing a report for a given system target by selecting one or more system targets and clicking on the View Details for Selected Systems button.

You will see counts for each system target selected by the time range selected. For instance if you are looking at the monthly time range, each column in the table will represent one day from the month. The count will be the count of observations for that day and system target.

Click on the system target name to drill down and show the counts by each target that comprises the system target. You can continue to click on the links in the first column of the table to drill down until you get to the entities that had observations (for example: file names, process names, user accounts, and so on).

Clicking on the count displays a screen that shows the actual observations that occurred during that time period.

Viewing Observations By Compliance Framework

The ability to view observations as they relate to a compliance standard structure is something that is typically done by a non-technical role such as an IT Manager, Line of Business Owner, Compliance Manager, or Executive.

You can identify some set of Compliance Frameworks that reflect the IT compliance framework that the organization follows. Observations can be filtered by whether they are authorized, unauthorized, unaudited or both. They can also be filtered by time.

To view observations by Compliance Framework, follow these steps:

1. From the Enterprise menu, select **Compliance**, then select **Real-time Observations**.
Clicking on the count displays a screen that shows the actual observations that occurred during that time period.
2. Click **Browse Observations by Compliance Frameworks**.
Cloud Control displays the Select Compliance Frameworks page that lists each defined Compliance Framework.
3. You can begin viewing a report for a given framework by selecting one or more frameworks and clicking on the View Details for Selected Frameworks button.
You will see counts for each framework selected by the time range selected. For instance if you are looking at the monthly time range, each column in the table will represent one day from the month. The count will be the count of observations for that day and framework. Click on the framework name to drill down and show the counts by each second-level framework folder that is in the selected framework. You can continue to click on the links in the first column of the table to drill down until you get to the entities that had observations (for example: file names, process names, user accounts, and so on).
4. Clicking on the count displays a screen that shows the actual observations that occurred during that time period.

This drill-down capability provided by these screens makes it easy for you to easily find where observations are occurring. When you have an environment with tens of thousands of targets across hundreds of business applications, it is impossible to view observations simply using a table and search unless you know exactly the search conditions they are looking for. In a matter of an hour, with this large of an environment even with little activity, there can be thousands of observations.

Viewing Observations By Search

For cases when the two browse by screens cannot provide the best view of what observations have happened in your environment, Cloud Control also provides a search capability to find observations.

To search observations, follow these steps:

1. From the Enterprise menu, select **Compliance**, then select **Real-time Observations**.
2. Click **Search Observations**.
Cloud Control displays the Search observation page which has search filters on the top half of the page and search results on the bottom half
3. You can set any number of filters in the search area. You can also click on the Add Fields button to add any fields that are available in the search results table.
4. With the options available in search, you can find observations performed over a time range, by a specific user, against a specific target, changes to a specific entity, and so on. Nearly every use case for finding observations can be solved using a combination of search fields.

Viewing Details of an Incident

Observations are logically bundled together based on the compliance standard rule, target and user that performed the action. This bundling is discussed in more detail in [Creating a Real-time monitoring Rule](#) section.

When one or more observations of a bundle are unauthorized, the bundle is considered to be in violation. This violation will lead to an event being created in Cloud Control Incident Management. The event name will be based on the message field defined in the real-time monitoring rule. When viewing this event in the incident management UI, several fields will show details of the bundle; the target type, entity type, number of observations in the bundle, observations by audit status, and so on. You can click on the [Update Audit Status](#) link to go to the bundle observations page.

This Observations page shows the list of observations in the observation bundle for this event. You can filter on various attributes for each observation, including but not limited to the authorized/unauthorized status, user, time, and so on.

Operations on Observations During Compliance Evaluation

The following sections describe how a real-time monitoring observation's audit status can be adjusted and how notifications can help in evaluating compliance results.

- [Manually Setting an Observation As Authorized Or Not Authorized](#)
- [Notifying a User When an Observation Occurs](#)
- [Notifying a User When an Authorized Observation Occurs](#)

Manually Setting an Observation As Authorized Or Not Authorized

Any time a user is viewing the details of a real-time observation, the user can change the audit status for the observation. You can override the audit status of an observation if you investigate the user action and determine that the activity should have resulted in a different audit status. Based on the real-time monitoring rule, all observations will either have a pre-set audit status or will have an audit status determined by an integration with a Change Request Management server. The available audit statuses are:

- **Unaudited:** No evaluation has happened to determine if the observation was good or bad.
- **Authorized:** The observation has been determined to be good, some action that was desired to occur.
- **Unauthorized:** The observation has been determined to be bad, some action that was not wanted.
- **Unauthorized-Cleared:** The observation had previously been determined to be bad, some action that was not wanted, but it has been handled through a fix, a policy change, or a compensating control and has now been cleared.

To change the audit status of an observation, view the observation from either of the browse by UI pages, the observation search page, or the incident manager UI. Select the observation and click [Update Audit Status](#). A popup will come up allowing you to select the new audit status and a comment describing the reason for the status change. The history of all audit status changes is maintained for each observation.

If the Cloud Control instance is using the Change Request Management server connector for integration, there are some special considerations:

- If you change an unauthorized observation into an authorized observation, then you have the option of entering a change request ID that is known to authorize the change. This change request ID should match a request that already exists in your change request management system. You can also enter a comment. If a change request ID is provided, then the change request is annotated with the change just as if the system had automatically authorized it. If an incident had been created for the observation bundle, then the event/incident is updated with the new number of unauthorized observations.
- If you change an authorized observation into an unauthorized or unaudited observation, any annotations that were made to any change requests are rolled back. If there was already an incident raised for the observation bundle, then the annotation is changed to update the number of unauthorized observations in the incident. If this is the first unauthorized observation in a group, then an event is created and an incident is raised. You can provide a comment for the change.
- When you manually set the observation to be authorized and enter a change request ID and the rule has change management integration enabled, no attributes of the change request are compared with the observation. The change request is simply updated with the observation details.
- When rolling back annotations in the change management server, the observation annotations are marked as rolled-back instead of actually removing the annotation. This occurs to avoid user confusion not knowing possibly why the annotations were removed. Also, if the observation later becomes authorized again, the rolled-back marking can simply be removed to bring the annotation back.

Notifying a User When an Observation Occurs

If a compliance standard rule is created and you do not use change management reconciliation with the rule, then there will be no automated authorized/unauthorized check done on the observations. You can specify for this rule that each observation bundle should result in informational event being generated for the observation bundle. Details on how to configure this is in the section [Creating a Real-time Monitoring rule](#).

The event will have a notation. From the Incident Management console the user can look at events and incidents. When looking at a single event, there is a link available to see the observations associated with this observation bundle's event. Each observation bundle can only have one event. If at least one observation in the bundle is unauthorized, then the bundle is considered to be in violation which results in the event being generated.

Since this notification does not require user intervention or follow-up action, it is treated as informational. If at a later time, someone changes one of these unaudited observations into an authorized or unauthorized one, a new informational event for the unaudited observations will not be re-delivered. It is delivered only once for the observation bundle. However if one of the observations is manually set to unauthorized, then a violation is raised for the entire observation bundle.

When at least one observation in a bundle is in an unauthorized state, a violation is created. That violation becomes an event in the Incident Manager Console. Use the Incident Manager feature to set up a notification. For more information about this, on the Incident Manager page, click on the online help link, [Setting Up Notifications With Rules](#) under the [Setting Up Notifications](#) section under [Getting Started](#).

Notifying a User When an Authorized Observation Occurs

When an authorized observation occurs, it is not typical for you to receive a notification on these observations since the activity that caused the observation was expected. If you are using change management reconciliation, you have an option to annotate the authorizing change request with the observation details. The updates to the change request is one way customers can learn of authorized activity. You can set filters in their change management system to let them know that a change request has had authorized activity against it.

Configuring Compliance Management

Before you can use the compliance features, compliance frameworks, compliance standards, and compliance standard rules must be defined for your enterprise.

The following sections describe how to define and maintain these compliance entities.

- [About Compliance Frameworks](#)
- [Operations on Compliance Frameworks](#)
- [About Compliance Standards](#)
- [Operations on Compliance Standards](#)
- [About Compliance Standard Rule Folders](#)
- [About Compliance Standard Rules](#)
- [Operations on Compliance Standards Rules](#)

About Compliance Frameworks

A compliance framework is a hierarchical structure where any node can be mapped to one or more compliance standards, compliance standard rule folders, and compliance standard rules. Compliance frameworks provide a way to map your standards to a structure similar to the regulatory or standards-based compliance structure you use in your company.

Managing Compliance Frameworks

To manage compliance frameworks, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to manage and choose the action you want to perform.

Frameworks Provided by Oracle and User-Defined Compliance Frameworks

There are compliance frameworks provided by Oracle and user-defined compliance frameworks.

- Compliance frameworks provided by Oracle include
 - Oracle Support Compliance is a collection of controls that check for expected environment compliance for Oracle Supportability.

- Oracle Generic Compliance Framework is a standard set of compliance standards and associated controls for tracking changes and events taking place across your IT infrastructure for determining how well your organization is in compliance with your IT policies.
- Security Technical Implementation Guide (STIG) is a set of standards to ensure Security Technical Implementation Guide (STIG) compliance.
- User-defined compliance frameworks

You can define a compliance framework to satisfy the needs of your organization.

Compliance frameworks provided by Oracle cannot be deleted or edited. However, if you want to extend these frameworks, use the Create Like functionality to create your own user-defined frameworks based on the Oracle provided frameworks and then edit the new frameworks.

Recommendation: It is highly recommended that you create a top level compliance framework like the ones provided for STIG and Oracle Generic compliance.

Benefits of Using Compliance Frameworks

Compliance standards are defined to perform tests on targets. Examples include: testing if a configuration value is set properly, test to see if real-time file changes are occurring, and so on. A compliance framework is a way to map how different control areas of your compliance initiative are going to be affected by the results of those tests.

An organization may choose to define a compliance framework that extends an Oracle provided compliance framework. This is accomplished by creating a new compliance framework like the Oracle provided compliance framework and include new or existing compliance standards. Then each compliance standard is mapped to an appropriate framework hierarchy folder so that any violation against the standard is also mapped to that framework folder. Each folder in the framework represents one control area.

Reasons for Using Compliance Frameworks

There are a number of reasons for creating compliance frameworks including:

- Mapping underlying IT violations to the regulatory and standard compliance controls used by your company so you can easily identify the compliance control areas that will be affected by the violations
- Compliance auditing at compliance specification level
- Auditing, security evaluation, and trend analysis

What Compliance Frameworks Can Do

A compliance framework can:

- Represent industry-standard compliance control areas or can be created to match your internal frameworks in use.

Many companies may start by using an industry-standard framework, but modify it according to their own needs and auditing requirements.

- Help in IT audits by identifying which compliance controls are at risk and may need compensating controls based on the violations. Without mapping your compliance checks to the control areas affected, it is hard to identify what the real impact would be in a compliance audit.

- Since compliance frameworks can contain compliance standards of different types (Repository, WebLogic Server Signature, Real-time monitoring), they provide a good way of grouping similar checks of different types for reporting purposes.

Usage Note

Evaluation Results for a repository rule may become invalidated if a compliance standard rule within a compliance framework is modified or deleted. Evaluation of a compliance standard always references the current compliance standard rule definition for each compliance standard rule within the compliance standard.

Operations on Compliance Frameworks

You can perform the following operations on a compliance framework:

- [Creating a Compliance Framework](#)
- [Creating Like a Compliance Framework](#)
- [Editing a Compliance Framework](#)
- [Deleting a Compliance Framework](#)
- [Exporting a Compliance Framework](#)
- [Importing a Compliance Framework](#)
- [Browsing Compliance Frameworks](#)
- [Searching Compliance Frameworks](#)

The following sections explain these operations.

Note: Before you perform any of the operations on compliance frameworks, ensure you have necessary privileges. For example, when creating a compliance framework, ensure you have access to the compliance standards you will be including during the definition of the framework. (See [Roles and Privileges Needed to Use the Compliance Features](#).)

Creating a Compliance Framework

To make the creation for the compliance framework easier, ensure that the compliance standards, which will be referred to by the compliance framework, are already defined in the Cloud Control. You can add system out-of-the-box and user-defined compliance standards to any hierarchical element of the compliance framework. If you do not define the compliance standards before hand, you must add them later.

To create a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. Click **Create** button.
4. Provide the Name and Author and click **OK**.
5. Once you have provided the information on the definition page, look at the options available when you right-click the name of the compliance framework (located at the top-left of the page). From this list you can create subgroups, include compliance standards, and so on.
6. Click **Save**.

Usage Notes

- Lifecycle status can be either Development or Production.
 - Development

Indicates a compliance framework is under development and that work on its definition is still in progress. While in development mode, all management capabilities of compliance frameworks are supported including editing of the compliance framework and deleting the compliance framework. Results of development compliance standards will NOT be viewable in target and console home pages, and the compliance dashboard.

Lifecycle status default is Development. It can be promoted to Production only once. It cannot be changed from Production to Development.
 - Production

Indicates a compliance framework has been approved and is of production quality. When a compliance framework is in production mode, its results are rolled up into a compliance dashboard, target and console home page.

Production compliance frameworks can only refer to Production compliance standards. A production compliance framework can be edited to add/delete references to production compliance standards only.

Lifecycle status cannot be changed from Production to Development.
- All compliance frameworks with the same keyword will be grouped together when sorted by the Keyword column.
- If you modify a repository or WebLogic Server signature compliance standard that has been added to a compliance framework, either by editing the compliance standard directly, or by using Import to overwrite the compliance standard with new settings, the existing evaluations become invalid. That is, if this modified compliance standard was included in a compliance framework that was previously evaluated, and has evaluation results, these results are no longer viewable.

Adding a Compliance Standard to a Compliance Framework

Click on a framework folder element that you want to map a compliance standard to. Right click and select Add Standards to bring up a popup to allow you to select the standards to map to this folder.

Use the search criteria to minimize the number of compliance standards that display in the select list.

Once you make your selections, click **OK**. The framework hierarchy screen refreshes and shows your newly included compliance standards under the framework folder element.

Editing Importance

After you map the compliance standards that are to be part of the selected compliance framework folder, you can edit the importance of each compliance standard for this specific folder.

The importance impacts the way the compliance score is calculated for this compliance standard in this framework folder.

See Overview of Compliance Score and Importance for details on how this score is computed.

Creating Like a Compliance Framework

To create a compliance framework like another compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. On the Compliance Framework Library page, highlight the compliance framework you want to use as the base and click the **Create Like** button.
4. Customize the fields as needed.

Ensure that the Compliance Framework name is different from the original compliance framework and any other existing compliance frameworks.

5. Click **Save**.
6. You can then edit this newly created framework and add or remove standards, subfolders, or modify importance levels.

Editing a Compliance Framework

Use the edit compliance framework feature to add new compliance standard rules to a compliance framework, or edit details of existing compliance frameworks, or remove compliance standards from the compliance framework.

To edit a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to edit and click the **Edit** button.
4. Update the properties as needed.

To add standards and subgroups, right-click the name of the framework located at the top left of the page.

5. Click **Save**.

Usage Notes

- Changing a compliance framework definition may impact trend analysis.
- The compliance standards you add to a compliance framework may be system-defined and user-defined compliance standards as displayed on the Compliance Standard Library page.
- If you modify a repository or WebLogic Server signature compliance standard that has been added to a compliance framework, either by editing the compliance standard directly, or by using Import to overwrite the compliance standard with new settings, the existing evaluations become invalid. That is, if this modified compliance standard was included in a compliance framework that was previously evaluated, and has evaluation results, these results are no longer viewable. The compliance framework evaluation results will again become visible after the next evaluation happens. The new evaluation includes the changes to the compliance standard within the compliance framework.

- The importance impacts the way the compliance score is calculated for this compliance standard in this framework folder.
- A compliance standard can be added to more than one compliance framework, and can have a different importance when added to a different compliance framework. For example, you could have a compliance standard called Check Password Expired which flags user accounts with expired passwords. This compliance standard may be a member of two compliance frameworks: All System Passwords Secure and 30-day Password Validation. The All System Passwords compliance framework verifies a password's security, whereas the 30-day Password Validation compliance framework checks the date that this password was last set.
 - The Check Password Expired compliance standard could have Extremely High importance for the 30-day Password Validation compliance framework, since this check is warning users that their passwords are about to expire.
 - In the All System Passwords Secure compliance framework, the Check Password Expired compliance standard could have a Normal importance, and other added compliance standards that do security checks could have a higher importance within the compliance framework.

Deleting a Compliance Framework

To delete a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to delete, click **Delete** button.
4. Confirm that you want to delete the compliance framework by clicking **OK**.

Usage Notes

- You can delete a single compliance framework or a list of compliance frameworks. When you delete a compliance framework, the associated metadata and evaluation results are also deleted.
- **YOU CANNOT DELETE COMPLIANCE FRAMEWORKS DEFINED BY ORACLE.** These are indicated by the presence of a lock icon in front of the compliance framework name on the compliance framework listing page.

Exporting a Compliance Framework

The Export feature provides a mechanism for transporting user-defined compliance framework definitions across Management Repositories and Cloud Control instances. The export stores the definitions in an operating system file. Because the exported compliance framework definitions are in XML format, they conform to the Oracle Compliance Standard Definition (XSD) format. You can then change the definition of the compliance framework and re-import the generated compliance framework definitions into another Management Repository.

To export a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. Highlight the compliance framework you want to export.

4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the compliance framework definition is to be exported. All leaf level rules and compliance standards are exported.

The system generates an XML representation of the compliance framework in the directory and file you specify.

Importing a Compliance Framework

Importing allows you to re-use a compliance framework that you already have, share framework definitions across multiple instances of Cloud Control, or enable offline editing of the framework.

Before you import a compliance framework, ensure the compliance framework to be imported is defined in a file. The file should be locally accessible to the browser you are using to access Cloud Control. Also ensure that you have privileges to access the compliance framework definition XML file to be imported.

Note:

When importing a compliance standard containing rules (or a framework containing standards) from the UI or command-line interface, import the xml file with `<ComplianceContent>` as root. This root file might have a list of rules, standards, frameworks, and standard groups.

This ensures that the framework and standard definition will be successfully imported. Also all associated targets will be re-evaluated based on the definition change made.

To import a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. From **Actions** menu, select **Import**.
4. Provide the file name from which the compliance framework definition (as per Compliance Framework XSD) will be imported. Specify whether to override an existing definition if one already exists. Specify whether to import referring content as well where all leaf level rules and compliance standards are imported. Real-time monitoring facets are also imported for real-time monitoring type of rules.
5. Click **OK**.

Browsing Compliance Frameworks

To browse a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. To view the details of a particular compliance framework, highlight the compliance framework and click **Show Details**.

Searching Compliance Frameworks

To search for a compliance framework, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Frameworks** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Browsing Compliance Framework Evaluation Results

To browse compliance framework evaluation results, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Frameworks** tab and then the **Evaluation Results** tab.
3. Highlight the compliance framework and click **Show Details** to view the details of a particular compliance framework.

Results include the following:

- Average compliance score for different targets evaluated for compliance standards referred to by the compliance framework
- Count of target evaluations (critical, warning, compliant) for different compliance standards referred to by the compliance framework
- Count of violations (critical, warning, minor warning) related to compliance standards referred to by the compliance framework

Searching Compliance Framework Evaluation Results

To search compliance framework evaluation results, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Frameworks** tab and then the **Evaluation Results** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Browsing Compliance Framework Errors

To browse compliance framework errors, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Frameworks** tab and then the **Errors** tab.

Usage Notes

The error may be an unexpected internal error or an error in the test.

Evaluation errors can often be due to configuration and installation issues. See the following manuals for information:

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*

- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*

If the installation and configuration are correct and the errors persist, call Oracle for assistance.

Searching Compliance Framework Errors

To search for compliance framework errors, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Frameworks** tab and then the **Errors** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Usage Notes

The error may be an unexpected internal error or an error in the test.

Evaluation errors can often be due to configuration and installation issues. See the following manuals for information:

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*
- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*

If the installation and configuration are correct and the errors persist, call Oracle for assistance.

Verifying Database Targets Are Compliant with Compliance Frameworks

For auditors to verify that database targets are in compliance with the compliance frameworks, the Cloud Control structure needs to be defined. The steps to provide this structure include the following:

1. Super Administrator creates three Cloud Control users: Compliance Author, IT Administrator, and Compliance Auditor.
2. Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.
3. Super Administrator assigns the same target privileges to IT Administrator and Compliance Auditor.
4. Compliance Author logs in to Cloud Control and views Oracle provided compliance frameworks, compliance standards, and compliance standard rules.
He then enables and disables the appropriate compliance standard rules and creates new compliance standard rules.
5. IT Administrator logs in to Cloud Control and associates the targets for which he has target privileges with the appropriate compliance standards.
6. IT Administrator sets up the correct configuration parameters and settings for the compliance frameworks, compliance standards, and compliance standard rules for a particular target.

He then creates a monitoring template from this target and applies it to the other targets, to which he has privileges, that require compliance standards.

7. Compliance Auditor logs in to Cloud Control to view the violations and errors at the Enterprise level, for which he has view privileges, and at each target level.

He would then take the necessary actions to rectify the errors and violations.

About Compliance Standards

A compliance standard is a collection of checks or rules. It is the Cloud Control representation of a compliance control that must be tested against some set of IT infrastructure to determine if the control is being followed.

Compliance standards are made up of the following in a hierarchical structure (see [Figure 50-9](#)):

- Compliance standard rules
- Rule folders that can include nested rule folders and individual compliance standard rules.


Rule Folders are hierarchical structures that contain compliance standard rules. A rule folder has an importance attribute that denotes the importance of the rule folder relative to its siblings at the same level. This importance is considered when determining compliance scores being rolled up from other sibling rule folders. A certain rule folder may have multiple tests that occur, in this way a certain test can be given more weight than other tests.

- Included compliance standards. A compliance standard can include other compliance standards.

Figure 50-9 Compliance Standard Definition

Key:

CS - compliance standard



What Compliance Standards Can Do

- Can represent industry-wide standards. A compliance standard is applicable to a single target type.
- Be used as a reference configuration or a certified configuration
- Be a collection of compliance standard rules describing best practices in an enterprise

For example, when a target fails to adhere to a compliance standard, the target is not in compliance with the compliance standard.

Accessing Compliance Standards

The compliance standards, including those provided by Oracle, are available on the Compliance Standard Library page. To access this page, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.

To view the compliance standard rules associated with the compliance standard, click the name of the compliance standard and click **Show Details**. Once the Compliance Standard Detail page appears, right click the name of the standard located at the top left of the page, and select either **Collapse**, **Expand All Below** or **Collapse All Below**.

Note: The compliance standards defined by Oracle cannot be changed. However, you can create a standard similar to the one provided by Oracle by using the Create Like feature.

General Usage Notes for Compliance Standards

You can override an existing compliance standard by checking the Overwrite existing compliance standards check box. As a result, evaluations of compliance standards require that the compliance standard is associated to one or more targets.

- For repository compliance standards, evaluation starts after the standard is associated with a target based on data collected from that target in the Management Repository.
- For WebLogic Server compliance standards, evaluation happens when the Management Agent-side evaluation metric is refreshed. The refresh occurs once every 24 hours for Oracle WebLogic Domain, Oracle WebLogic Java EE Server, and Oracle WebLogic Cluster targets.
- For Real-time Monitoring compliance standards, monitoring at the Management Agent starts when a compliance standard is associated to a target. A violation occurs when an observation bundle contains at least one observation that is unauthorized

Usage Note Specific to Repository Rules

If you manually type a WHERE clause in the compliance standard rule XML definition, then the < (less than) symbol must be expressed as <, to create a valid XML document. For example:

```
<WhereClause>:status &lt; 100</WhereClause>
```

Example of How to Set Up Compliance Standards for Auditing Use

For auditors to verify that database targets are in compliance with the compliance frameworks, the Cloud Control structure needs to be defined. The steps to provide this structure includes the following:

1. Super Administrator creates three Cloud Control users: Compliance Author, IT Administrator, and Compliance Auditor.
2. Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.

3. Super Administrator assigns the same target privileges to IT Administrator and Compliance Auditor.
4. Compliance Author logs in to Cloud Control and views Oracle provided compliance frameworks, compliance standards, and compliance standard rules.
He then enables and disables the appropriate compliance standard rules and creates new compliance standard rules.
5. IT Administrator logs in to Cloud Control and associates the targets for which he has target privileges with the appropriate compliance standards.
6. IT Administrator sets up the correct configuration parameters and settings for the compliance frameworks, compliance standards, and compliance standard rules for a particular target.
He then creates a monitoring template from this target and applies it to the other targets, to which he has privileges, that require compliance standards.
7. Compliance Auditor logs in to Cloud Control to view the compliance dashboard, violations and errors at the Enterprise level, for which he has view privileges, and at each target level.
He would then take the necessary actions to rectify the errors and violations.

Operations on Compliance Standards

You can perform the following operations on a compliance standard:

- [Creating a Compliance Standard](#)
- [Creating Like a Compliance Standard](#)
- [Editing a Compliance Standard](#)
- [Deleting a Compliance Standard](#)
- [Exporting a Compliance Standard](#)
- [Importing a Compliance Standard](#)
- [Browsing Compliance Standards](#)
- [Searching Compliance Standards](#)

The following sections explain these operations.

Note: Before you perform any of the operations on compliance standards, ensure you have necessary privileges. For example, when creating a compliance standard, ensure you have access to the compliance standard rules you will be including during the definition of the compliance standard. (See [Roles and Privileges Needed to Use the Compliance Features](#).)

Creating a Compliance Standard

You can use the compliance standards provided by Oracle, for example, Security Configuration for Oracle Database, or create your own standard.

Before creating a compliance standard, ensure the compliance standards and compliance standard rules, which will be referred to by the compliance standard, are defined in the Management Repository.

To create a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Click the **Create** button. You will be prompted for the Name, Author, target type to which the standard is applicable, and the standard type. The standard types are:
 - Repository
 - WebLogic Server Signature
 - Real-time Monitoring
 - Agent-side

Click **Continue**.

4. On the resulting Properties tab, provide the property values.
Click **Add** to either add a keyword by which this standard is identified or use an existing keyword.
5. To further define the compliance standard, right-click the name of the compliance standard located at the top left of the page. From this menu, you can create rule folders, add rules, and include compliance standards.

By using rule folders, you can view the summary of results, categorized by the targets that were evaluated against the selected rule folder and the Compliance Standard Rules evaluated for the selected rule folder.

6. Click **Save**.

Once you define the compliance standard, associate the standard with a target and define the target type-specific settings.

1. While on the Compliance Standards Library page, ensure the correct compliance standard is highlighted.
2. Click the **Associate Target** button.
3. On the **Target Association for Compliance Standard** page, click **Add** to choose the target to be evaluated against the standard.
4. In the **Search and Select: Targets** popup, choose the appropriate targets.
5. Click **Select**.

After you associate the targets with the compliance standard, you can edit the parameters associated with the target.

1. While on the **Target Association for Compliance Standard** page, click **Edit**.
2. On the **Customize Compliance Standard Parameters** page, change the parameters as needed.

 Note:

You can also associate a compliance standard with a target from the target home page. At the top left of the target's home page, right click the name of the target. On the resulting menu, select **Compliance**, then select **Standard Associations**.

Including a Compliance Standard into Another Compliance Standard

Use the Include Compliance Standard page to select one or more compliance standards to be included into the compliance standard. This list is prefiltered by the target type of the compliance standard.

To include a compliance standard into another compliance standard:

1. From the Compliance Standard Library page, highlight the compliance standard to which you want to add another compliance standard.
2. Click the **Edit** button.
3. On the Properties page, right-click the node, located at the top left of the page.
4. On the resulting menu, select **Add Standards**.
5. Select the compliance standard to include. Click **OK**.

When you include a compliance standard within another top level compliance standard, the included standard must be of the same target type as the top level compliance standard. For composite target types, one of the member target types of the composite target type of the top level standard is a member target type within the top level composite target type.

Note that a root compliance standard is associated to a root target (of composite target type). Compliance standards are associated to member targets of the same applicable target type and target filter criteria.

6. On the **Properties** page, choose the **Importance** for the compliance standard you just included. Click **Save**.
7. After the compliance standard is included, highlight the root compliance standard. The Properties page displays a set of parameters.

A parameter is a variable that can be used by one or more compliance standard rules contained in that compliance standard. When a compliance standard rule references a parameter, the parameter's actual value is substituted at compliance standard rule evaluation time. It is through the use of parameters that customizations of compliance standards is supported.

Usage Notes

- Because compliance standards are hierarchical, the top node in the tree is known as the root node.
- When you create a compliance standard, the version is 1.
- Lifecycle status default is Development. It can be promoted to Production only once. It cannot be changed from Production to Development.

- Development

Indicates a compliance standard is under development and that work on its definition is still in progress. While in Development mode, all management capabilities of compliance standards are supported including complete editing of the compliance standard, deleting the compliance standard, and so on. However, while the compliance standard is in Development mode, its results are not viewable in Compliance Results nor on the target or Cloud Control home page.

- Production

Indicates a compliance standard has been approved and is of production quality. When a compliance standard is in production mode, you have limited editing capabilities, that is, you can add references to production rules, and you

can delete references to rules ONLY from a compliance standard. All other management capabilities such as viewing the compliance standard and deleting the compliance standard will be supported. Results of production compliance standards are viewable in target and console home pages, and the compliance dashboard. Production compliance standards can only refer to production compliance standards and production compliance standard rules.

Once the mode is changed to Production, then its results are rolled up into compliance dashboard, target home page, and Cloud Control home page. Production compliance standards can only refer to other production compliance standards and production compliance standard rules. A production compliance standard can be edited to add and delete references to production compliance standards and production compliance standard rules only.

Creating Like a Compliance Standard

To create a compliance standard like another compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Click the **Create Like** button.
4. Customize the fields as needed.
The name has to be different than an existing Compliance Standard.
5. Click **Save**.

Editing a Compliance Standard

You can customize compliance standards by editing the existing compliance standard rule settings. You can change the added rules' importance for the compliance score calculation, prevent template override, override default parameter values (when possible), and exclude objects from a compliance standard rule's evaluation (when possible).

Note: You cannot edit an Oracle provided compliance standard, that is, a compliance standard defined by Oracle.

To edit a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Highlight the standard you want to edit and click the **Edit** button.
4. Update the parameters as needed.
5. Click **Save**.

Deleting a Compliance Standard

Before you delete a compliance standard, ensure the compliance standard is not in use by a compliance framework. You must remove any references to the compliance standard in all compliance frameworks.

Note: You cannot delete an Oracle provided compliance standard, that is, a compliance standard provided by Oracle.

To delete a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Highlight the compliance standard you want to delete, click **Delete** button.
4. Confirm that you want to delete the standard by clicking **OK**.

Exporting a Compliance Standard

The Export feature provides a mechanism for transporting user-defined compliance standard definitions across Management Repositories and Cloud Control instances. The export stores the definitions in an operating system file. Because the exported compliance standard definitions are in XML format, they conform to the Oracle Compliance Standard Definition (XSD) format. You can then change the definition of the compliance standard and re-import the generated compliance standard definitions into another Management Repository.

Before you export a compliance standard, ensure that you have privileges to access the compliance standard to be exported.

To export a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Highlight the standard you want to export.
4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the standard definition is to be exported. All leaf level rules and compliance standards are exported.
6. The XML representation of the compliance standard is generated. The file is located in the directory you specify.

Importing a Compliance Standard

The Import feature uploads an XML-based compliance standard definition file containing definitions of a single user-defined compliance standard or a list of user-defined compliance standards. This upload creates a new user-defined compliance standard or a list of user-defined compliance standards. This compliance standard must have been previously exported.

The compliance standard xml definition must comply with the compliance standard XML Schema Definition (XSD) as defined in User-Defined Compliance Standard XML Schema Definition.

Before importing a compliance standard, ensure the compliance standard to be imported is defined in a file. The file should be locally accessible to the browser you are using to access Cloud Control. Also ensure that you have privileges to access the compliance standard definition XML file to be imported.

To import a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.

3. From the **Actions** menu, select **Import**.
4. Provide the file name from which the compliance framework definition (as per Compliance Framework XSD) will be imported. Specify whether to override an existing definition if one already exists. Specify whether to import referring content as well.
5. Click **OK**.

You can override an existing compliance standard by checking the Overwrite existing compliance standards check box. As a result:

- If you override a compliance standard, the override deletes all target and template associations, as well as evaluation results for that compliance standard.
- If the overwritten compliance standard is part of a compliance framework, the compliance standard is updated in the compliance framework. However, the evaluation results for that compliance standard within the compliance framework are invalidated.

Browsing Compliance Standards

To browse a compliance standard, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. To view the details of a particular standard, highlight the standard and click **Show Details**.

Searching Compliance Standards

To search for compliance standards, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Browsing Compliance Standard Evaluation Results

To browse compliance standard evaluation results, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Standards** tab and then the **Evaluation Results** tab.
3. Highlight the compliance standard and click **Show Details** to view the details of a particular standard.

Results include the following:

- Average compliance score for different targets
- Count of target evaluations (critical, warning, compliant)
- Count of violations (critical, warning, minor warning)

Searching Compliance Standard Evaluation Results

To search for compliance standard evaluation results, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Standards** tab and then the **Evaluation Results** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Browsing Compliance Standard Errors

To browse compliance standard evaluation errors, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Standards** tab and then the **Errors** tab.

Searching Compliance Standard Errors

To search for compliance standard errors, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Results**.
2. Click the **Compliance Standards** tab and then the **Errors** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.
4. Click **Search**.

Usage Notes

- Use the Evaluation Errors page to view the errors that occurred as a result of metric collection, as well as those that occurred during the last evaluation.
- Use the search filter to view only those evaluation errors that meet a set of search criteria that you specify.
- Click the message in the Message column to decide what your course of action should be to resolve the error.
- On initial display, the Evaluation Errors page shows all the evaluation errors.
- Normally the results of an evaluation overwrite the previous evaluation's results. However, in the case of evaluation failure or data provider collection failure, the previous results are left untouched.

Once the underlying problem is fixed, the error is no longer reported.

Example of Search Filter

By default, all the evaluation errors in your enterprise configuration appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the evaluation errors that meet those criteria in the results table.

For example, if you choose Host in the Target Type list, contains in the Target Name list, and "-sun" in the adjacent Target Name text field, and then click **Go**, Cloud Control displays, in the results table, only the compliance standard rule evaluation errors for the hosts that contain "-sun" in their names.

Associating a Compliance Standard with Targets

After you create a compliance standard, you can associate the standard with one or more targets. As part of the association, you can customize parameters, that is, the importance of the standard in relation to the target, status of the compliance standard evaluation, reason for changing the evaluation status, and the thresholds.

Before you associate a compliance standard with a target, ensure you have privileges to access the targets you want to associate compliance standards to.

To associate a compliance standard with a target, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Highlight the compliance standard you want to associate with various targets. Click the **Associate Target** button.
4. Select the targets you want to associate with this compliance standard. Click **OK**.
5. With the compliance standard still highlighted, click the **Override Target Type Settings** button.
6. Customize the critical and warning thresholds and importance as needed.

By changing critical and warning thresholds, you signify how the Compliance standard score event is generated. For example, if the actual score is less than the critical threshold, then a critical score event is raised.

Changing the importance can change the compliance score. The importance denotes how important the compliance standard is in the hierarchy.

7. Click **OK**.

To further customize the evaluation of a compliance standard against a target, you can alter compliance standard parameters: importance, critical threshold, and warning threshold. Customizations can also be made on the compliance standard rules used within the compliance standards. For example, for the Secure Ports compliance standard rule, DFLT_PORT is an override parameter. You can change the default value of the port. You can also exclude objects from the evaluation, for example a particular port from the evaluation.

Note: For real-time monitoring, you can change parameters that are used in facet patterns. You can also change Automatic Change Management reconciliation settings.

By changing critical and warning thresholds, you signify how the Compliance standard score event is generated. For example, if the actual score is less than the critical threshold, then a critical score event is raised.

Best Practices

You can perform compliance association in two ways: for testing and editing, and production and mass associations.

- For testing and editing a standard/target and standard rule, or rule folder/target association settings purposes, associate the target with a compliance standard as previously described in this section.

Using the Compliance UI, you can:

- Test the association and remove it after testing is complete.

- Edit the association for importance, evaluation status, and thresholds.

Note: You *cannot* edit an association using the Administration Groups and Template Collections page.

- For production and mass associations, associate the target using the Administration Groups and Template Collections page:

From the **Setup** menu, select **Add Target**, then select **Administration Groups**. Click the **Associations** tab.

Because each Administration Group in the hierarchy is defined by membership criteria, a target is added to the group only if it meets the group's membership criteria. Therefore, when a target is successfully added to a group, it is automatically associated with the eligible compliance standards for that group. This makes it easier to associate a target to a large number of compliance standards.

Associating a Compliance Standard with a Group Target

After you create a compliance standard, you can associate the standard with a group target. This enables the association of key standards to targets when they are part of the group.

Before you associate a compliance standard with a group target, ensure the following prerequisite is met:

- You have privileges to access the group target you want to associate the compliance standards to

See [Privileges and Roles Needed to Use the Compliance Feature](#).

Perform the following steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. Highlight the compliance standard you want to associate with the group target. Click the **Associate Groups...** button.
4. Select the group target you want to associate with this compliance standard. Click **OK**.

After you click OK, the group target is associated to the compliance standard and all eligible targets with the group are associated to the compliance standard. In the future when new targets are added to the group target, and if they have the same target type and match the target property filter criteria, they will then be automatically associated to the compliance standard.

Viewing Real-time Monitoring Compliance Standard Warnings

When you associate a real-time monitoring compliance standard to targets, there is a chance that there are setup steps that were not followed on the target to enable real-time monitoring or there could be inconsistency with the configuration. Any warnings will be shown on the Associate Targets screen. This screen is reached by selecting a compliance standard and selecting **Associate Targets** button. If there are any warnings, there will be a warning icon with a link above the table of target associations. Clicking on this link will take you to a screen that lists all current warnings for this compliance standard.

All warnings can be fixed by correcting some configuration problem on the host/target you are monitoring or by fixing rule/facet content. Once the underlying problem is fixed, these warnings will be cleared automatically.

This list of warnings is also available on the Real-time Observations page (from the Enterprise menu, select **Compliance**, then select **Real-time Observations**) where you can pick one of three types of reports to view your observations. The bottom half of the screen shows all active warnings across all targets and compliance standards related to real-time monitoring.

Enabling Security Metrics

Because security collections are disabled by default, they must be enabled before using security features like security compliance standards, reports, and so on. To enable Security metrics, follow these steps:

1. From the **Enterprise** menu, select **Monitoring**, then select **Monitoring Templates**.
2. In the Search area, select **Display Oracle provided templates and Oracle Certified templates** and click **Go**.
3. Select **Oracle Certified-Enable Database Security Configuration Metrics** and click **Apply**.
4. In the Destination Targets region on the Apply Monitoring Template Oracle Certified-Enable Database Security Configuration Metrics: General page, click **Add**.
5. On the Search and Select: Targets page, select the database instances in which you are interested and click **Select**.
6. In the Destination Targets region of the Apply Monitoring Template Oracle Certified-Enable Database Security Configuration Metrics: General page, select the database instances in which you are interested and click **OK**.

After you click **OK**, a confirmation message on the Monitoring Templates page appears.

Considerations When Creating Compliance Standards

A compliance standard will refer to one or more Compliance Standard Rules. When creating a compliance standard, the standard should be granular enough that it can be appropriately mapped to one or more related Compliance Frameworks. For example, consider this Compliance Framework structure that exists in the Oracle Generic Compliance Framework:

- Change and Configuration Management (compliance framework subgroup)
 - Database Change (compliance framework subgroup)
 - * Configuration Best Practices for Oracle Database (compliance standard)
 - * Configuration Best Practices for Oracle RAC Database (compliance standard)
 - * Configuration Best Practices for Oracle Pluggable Database (compliance standard)

Many compliance standards will exist that should be mapped to this part of the Compliance Framework structure, each with their own rules to address this specific require-

ment. One may check that configuration settings are set properly. Another may be used to check in real-time if anyone changes a configuration setting.

In this example, the "Database Change compliance framework subgroup" can relate to many different types of targets. Oracle Database, Oracle RAC Database, and Oracle Pluggable Database all have their own types of configurations that all need to be secured. Any Standards created to monitor these target-specific configurations would map to the same "Database Changes subgroup".

If compliance standards are structured in a granular way so that they can map to existing and future compliance frameworks, then violations in a rule can be rolled up to impact the score of the compliance framework properly.

About Compliance Standard Rule Folders

Rule Folders are optional hierarchical structures used to group similar compliance standard rules within a compliance standard. You can add individual compliance standard rules to a compliance standard, or group them if you have a large number of rules in a standard. A compliance standard rule can be added to multiple Rule Folders within a compliance standard, each with different importance settings. Rule Folders can be nested within a compliance standard.

A rule folder has an importance attribute that denotes the importance of the rule folder relative to its siblings at the same level. This importance is considered when determining compliance scores being rolled up from other sibling rule folders. A certain rule folder may have multiple tests that occur, in this way a certain test can be given more weight than other tests.

The following topics address compliance standard rule folders:

- [Creating Rule Folders](#)
- [Managing Rule Folders in a Compliance Standard](#)

Creating Rule Folders

To create a rule folder, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standards** tab.
3. On the Compliance Standard Library page, highlight the compliance standard and click **Edit**.
4. On the **Properties** page, right-click the name of the compliance standard. The name of the standard is located in the top-left corner of the page.
5. Select **Create Rule Folder**.
6. Type the name of the folder and click **OK**.
7. On the **Properties** page, provide a description, ReferenceUrl, and importance. See [Overview of Compliance Score and Importance](#).

Managing Rule Folders in a Compliance Standard

After you create a rule folder and populate it with compliance standard rules, you can perform the following actions on the folder:

- Edit the tree structure by re-ordering the Rule Folder, Rule Reference, and Compliance Standard Reference nodes in the tree or by deleting any of these nodes.
- Select any node (except the top-level Compliance Standard node) object and then click **Remove** menu item from context menu. The Remove option is disabled on the root node. You can also select multiple objects and click **Remove** to delete multiple nodes.

About Compliance Standard Rules

A compliance standard rule is a test to determine if a configuration data change affects compliance. Based on the result of the test, a compliance score is calculated. These rule compliance scores are rolled up to compute the compliance standard score and then this score can be rolled up and reported along with the compliance framework scores.

Types of Compliance Standard Rules

There are three types of compliance standard rules are:

- **Agent-side Rules**
Used for detecting configuration problems on the agent. This enables the implementation of the Security Technical Implementation Guide (STIG) security specifications. Agent-side rules generate violations for a target which is based on the results data collected for the underlying configuration extension target.
- **Configuration Consistency rule**
Determines the consistency of targets of similar target types within a composite target. For example, a user has a Cluster Database made up of 15 databases. He can use the Cluster Database Comparison Template for configuration consistency to flag databases that may have changed within the cluster.
- **Configuration Drift rule**
Determines the deviation of targets of similar target types. For example, a user has 10 databases that he is monitoring. He needs to ensure that the Initialization Parameter File Permission compliance standard rule is the same across all the databases. This deviation can occur when the database configuration has been updated.
- **Manual rule**
Enables you to account for checks that cannot be performed automatically, thus allowing you to account for these types of checks in the compliance framework.
For example, a common security check is "To ensure secure access to the data center". When a standard is associated to a target, each manual rule will have one violation. A user must manually attest to the positive status of the rule. In other words, a person responsible for the task ensures he has performed the task. The compliance framework records when and who clears the violation of the manual check so it can be reported.
- **Missing Patches rule**
Used for detecting patches that have not been applied to the appropriate targets. This rule generates violations which appear on the compliance results UI and subsequent compliance dashboard regions. A rolled up violation count appears on the dashboard regions. The user can drill down to examine violation details and then correct the issue by applying the missing patches to the appropriate targets.

- If the rule is based on a list of patches, then the rule checks if none of the patches are applied to the target. If any of the patches are applied, then no violation is generated. If none of the patches are applied, then one violation is generated listing the patches that are not applied.
- The patch numbers can refer to Oracle recommended patches or manually entered patches.
- After a patch is applied, the corresponding ORACLE_HOME configuration is uploaded. Oracle then reevaluates all associated missing patches rule for the target.
- After you create the Missing Patches rule, you can add missing patches rules to compliance standards of type Repository. You can then associate the standard to targets by selecting a standard, and clicking the Associate Target button. Upon association, the missing patch rule will be evaluated on the applied targets.
- If a standard with the missing patches rule is associated to a group, when new targets are added to the group, the new target is automatically evaluated for missing patches.

- Real-time Monitoring rule

Monitors operating system and database level entities that store configuration data. Real-time monitoring rules define the entities to monitor, user actions to watch for, and any types of filters to apply to the monitoring. Monitoring can be filtered by: when changes occurred, who made the changes, and what process made the changes.

The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type. For example, you may choose to define a facet that lists all of the critical configuration files for the Host target type. These configuration files would be the ones that, if changed, would most likely result in instability of the host. You may also create a facet that lists all users which are DBA users.

The real-time monitoring rule can be part of a compliance standard that is associated with one or more targets. The monitoring can occur on any operating system level entity, for example, file, process, user, registry, and so on. Real-time monitoring rules can additionally specify whether observations captured by the rule are automatically reconciled. This reconciliation determines whether the actions observed were authorized or not.

Change Request Management reconciliation compares open change requests to actions performed on targets. If there is a match of expected actions to actual actions, then those actions are authorized, otherwise they are unauthorized. Authorizations can also be done manually. All observations are captured and bundled by rule, target and user. Attributes can be set on the frequency of observation data collection.

- Repository Rules

Used to perform a check against any metric collection data in the Management Repository.

Used for checking the configuration state of one or multiple targets. A rule is said to be compliant if it is determined that the configuration items do in fact meet the desired state and the rule test failed to identify any violations. Otherwise, a rule is said to be non-compliant if it has one or more violations. The data source that is

evaluated by a compliance standard rules test condition can be based on a query against the Cloud Control Management Repository. A compliance standard rules test condition can be implemented using a threshold condition based on the underlying metrics (or queries) column value or SQL expression or a PLSQL function. To use a rule, it must be associated to one or more compliance standards. The compliance standard then will be associated to one or more targets. This effectively enables this rule to be evaluated against these targets.

- WebLogic Server Signature rule

Preemptively identifies WebLogic Server configuration problems. The purpose of the WebLogic Server Signature rules is to evaluate at the WebLogic Server if certain configuration data satisfies some conditions (or checks) and the evaluation results are sent as violation information to the Management Service.

Detailed information about how to identify problems is specified in the WebLogic Server Signature rule definition. The WebLogic Server Signature rule definition includes Dataspec and XQuery logic that are used to determine what is important to collect and evaluate for a given target type and target properties. A Dataspec is a group of MBeans used to collect from a WebLogic Server. The XQuery logic contains the check on the collected data (by the MBeans). The WebLogic Server Signature rule can be associated with one or more specific Web Logic targets: Web Logic Domain, Web Logic Java EE Server, and Web Logic Cluster.

Version-specific details include:

- To enable data collection for the WebLogic Server signature-based rules on WebLogic Server targets earlier than v10.3.3, you need a copy of bea-guardian-agent.war. You can find a copy of this war file in your OMS installation's work directory: \$T_WORK/middleware/wlserver_10.3/server/lib/bea-guardian-agent.war
- For WebLogic Server v9 and v10.0
Install and deploy bea-guardian-agent.war to all servers in the domain. Do not change the context root. See <http://<host>:<port>/console-help/doc/en-us/com/bea/wlserver/core/index.html> for more information on installing a web application.
- For WebLogic Server v10.3 up to and including v10.3.2
Copy the war file from your OMS installation into each target's \$WL_HOME/server/lib directory. Restart all the servers in the target domain.
- For WebLogic Server v.10.3.3 and higher
No action is required.

Operations on Compliance Standards Rules

The following sections explain the operations you can perform on compliance standard rules.

- [Creating a Repository Compliance Standard Rule](#)
- [Creating a WebLogic Server Signature Compliance Standard Rule](#)
- [Creating a Real-time Monitoring Compliance Standard Rule](#)
- [Creating an Agent-side Rule](#)
- [Creating a Manual Rule](#)

- [Creating Like a Compliance Standard Rule](#)
- [Editing a Compliance Standard Rule](#)
- [Deleting a Compliance Standard Rule](#)
- [Exporting a Compliance Standard Rule](#)
- [Importing a Compliance Standard Rule](#)
- [Browsing Compliance Standard Rules](#)
- [Searching Compliance Standard Rules](#)

Note: Before you perform any of the operations on compliance standard rules, ensure you have the necessary privileges. (See [Roles and Privileges Needed to Use the Compliance Features](#).)

Creating a Repository Compliance Standard Rule

To create a repository compliance standard rule to check if a target has the desired configuration state based on collected configuration data, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select Repository Rule as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - **Rule Name**
Provide a unique name for the rule.
 - **Compliance Rule State**
Set whether the state of this rule is development or production. Development means that the rule is still being defined or tuned and is not yet ready to be used on targets yet. After you promote a rule to production, you cannot change it back to development.
 - **Severity**
The rule can have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score along with the importance that may be set for this rule when it is added to a compliance standard.
 - **Applicable To**
Target type this rule works against.
 - **Target Property Filter**
You can specify specific target properties that determine which targets this rule can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this rule, for instance for Linux OS, it will only be applicable to targets on Linux Operating System.
 - **Description**

Description of the rule

- Rationale

Text describing what this rule is checking and what the effect of a violation of this rule may be.

- Recommendation

Recommendation text describing how to fix a problem when a violation occurs.

- Reference URL

URL to a document that describes the compliance control in more details. Many times these documents may be stored in a content management system.

- Keywords

Keywords can be assigned to a rule so that you can control how data is organized in various reports.

7. Click **Next**.

8. On the next screen, you need to provide a SQL query that will execute against the Cloud Control Management Repository. You can directly enter the SQL query, or click the Model Query button to enter a screen that will guide you through choosing the query content.

9. Enter Compliant and Non-Compliant Message. These are the messages that will be shown in regards to the evaluation. When a violation occurs, the Non-Compliant message will be the string describing the event under the Incident Management capabilities.

10. Enter the Recommendation. The recommendation describes how to fix a problem when a violation occurs.

11. Click **Next**.

12. On the next screen, you will see the columns that will be returned from this query as part of the evaluation results. You can modify the display name of each column as needed.

13. On this screen, you also need to set the condition you are checking against the returned query results to look for a violation. Your condition check can be a simple one based on the column name and a comparison operator of the value. Or you can compose a SQL condition by providing parameter names and providing a where clause to add to the evaluation query.

14. If you are using the SQL condition, you can click the **Validate Where Clause** button to check for any issues with your condition.

15. Click **Next**.

16. The next screen will allow you to test your rule. You can choose a target in your environment and click the Run Test button. Any issues with the rule will be displayed and you can resolve them before saving the rule.

17. Click **Next**.

18. The final page allows you to review everything you have configured for this rule. Ensure that everything is correct and click the Finish button to save the rule.

Additional Notes for Repository Rules

- All rules are visible in the global rule library and are visible to all users.

- Once the compliance standard rule is created, it is not automatically evaluated. Users must associate a rule to a compliance standard before it can be used. Only when a compliance standard is associated with one or more targets will a rule evaluation occur. Rules cannot be evaluated directly.
- One rule can be associated to multiple compliance standards.
- Various attributes of a rule can be customized through the compliance standard this rule is associated with. These customizations occur in the Compliance Standard screens. One of these attributes that can be customized per compliance standard is the importance of the rule in relationship to this standard.
- Because the user-defined compliance standard rule is defined by a privileged user, only privileged users can modify the compliance standard rule. Violation results are available to all users.
- To share this user-defined compliance standard rule with other privileged users, provide the XML schema definition (using the Export feature) so they can import the compliance standard rule to their Management Repository.
- You can minimize scrolling when reading the Description, Impact, and Recommendation information by restricting the text to 50 characters per line. If more than 50 characters are needed, start a new line to continue the text.
- Look at the context-sensitive help for information for each page in the Compliance Standard Rule wizard for specific instructions.
- If you manually type a WHERE clause in the compliance standard rule XML definition, then the < (less than) symbol must be expressed as <, to create a valid XML document. For example:

```
<WhereClause>:status &lt; 100</WhereClause>
```

Creating a WebLogic Server Signature Compliance Standard Rule

There are several hundred out-of-the box WebLogic Server signature rules designed to uncover compliance violations known to occur in WebLogic installations based primarily on in-depth knowledge of common pitfalls and best practices. You can also create your own rules to extend the checks that are performed.

A signature describes a potential problem in a WebLogic installation. It consists of categorization metadata, a user-readable description of the problem, and an XQuery expression for evaluating whether the problem exists at the target.

A WLS Signature rule is an Management Agent-side rule that checks a signature definition against an associated target for the existence of the problem the signature defines. WebLogic Server targets include: WLS Domain; WLS Cluster; WebLogic Managed Server. The first two are composite target types: logical groupings of instances of simple WebLogic Server targets. Rules must be evaluated against the whole domain or cluster to render meaningful violation results.

WLS Signature rules, like other compliance rules, are grouped into Compliance Standards, which are logical groupings based on signature metadata such as severity and remedy.

To create a WebLogic Server Signature compliance standard rule to evaluate if certain configuration settings satisfy known good configurations, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.

3. Click the **Create** button.
4. In the Create Rule popup, select the WebLogic Server Signature rule type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - Rule Name
Provide a unique name for the rule.
 - Compliance Rule State
Set whether the state of this rule is development or production. Development means that the rule is still being defined or tuned and is not yet ready to be used on targets yet. After you promote a rule to production, you cannot change it back to development.
 - Severity
The rule can have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score along with the importance that may be set for this rule when it is added to a compliance standard.
 - Applicable To
Target type this rule works against.
 - Target Property Filter
You can specify specific target properties that determine which targets this rule can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this rule, for instance for Linux OS, it will only be applicable to targets on Linux Operating System.
 - Description
Description of the rule
 - Rationale
Text describing what this rule is checking and what the effect of a violation of this rule may be.
 - Recommendation
Recommendation text describing how to fix a problem when a violation occurs.
 - Reference URL
URL to a document that describes the compliance control in more details. Many times these documents may be stored in a content management system.
 - Keywords
Keywords can be assigned to a rule so that you can control how data is organized in various reports.
7. Click **Next**.
8. On the next screen, you select the method of providing the signature definition file. You can either load it by uploading a file, or enter the text directly into the UI.

9. Enter Compliant and Non-Compliant Message. These are the messages that will be shown in regards to the evaluation. When a violation occurs, the Non-Compliant message will be the string describing the event under the Incident Management capabilities.
10. Choose the columns that will be displayed along with violations. These columns should be defined as return columns in the signature definition.
11. Click **Next**.
12. The next screen will allow you to test your rule. You can choose a target in your environment and click the **Run Test** button. Any issues with the rule will be displayed and you can resolve them before saving the rule.
13. Click **Next**.
14. The final page allows you to review everything you have configured for this rule. Ensure that everything is correct and click the **Finish** button to save the rule.

This newly created rule does not function until it is associated to one or more compliance standards and those compliance standards are associated to targets. Once this association happens, the following is the workflow of this rule:

- The standard/rule combination gets transferred to and then evaluated on the Management Agent-side against a metric collected specifically for the Compliance Standard and target type to determine compliance.
- The evaluation generates violations (if any).
- Violations are uploaded to Cloud Control server, from where they are subsequently processed into violations in Management Repository tables.
- Violations are then viewable in compliance results pages and the Compliance Dashboard.

Example WebLogic Server Signature

Using the rule creation wizard makes it simple to add a new rule, but the important part of the WebLogic Server signature rule is the signature definition. A signature definition consists of a list of managed beans (MBeans) and an XQuery expression. Managed beans represent the configuration data to collect. They define a type and the attributes within the type to collect. They also declare which attributes to consider in determining whether there are violations. The XQuery expression defines the logic to use in evaluating the collected data for compliance. An XML example signature definition follows.

```
<SignatureDefinition> xmlns="http://www.oracle.com/DataCenter/ConfigStd">
  <MBeanList>
    <MBean scoreBase="true" mBeanType="ServerRuntime">
      <AttributeName>Name</AttributeName>
      <AttributeName>WeblogicVersion</AttributeName>
    </MBean>
  </MBeanList>
  <XQueryLogic>declare function local:getServerRuntimesEqualToVersionWith-
Patch($targetData, $major as xs:integer,
$minor as xs:integer, $servicePack as xs:integer, $crNumber as xs:string) {
  for $ServerRuntime in $targetData/DataCollection/ServerRuntime
  let $weblogicVersion := fn:replace($ServerRuntime/@WeblogicVersion, &quot;WebLogic
Server Temporary Patch&quot;;, &quot;&quot;&quot;);
  let $majorVersion :=
    let $spaceParts := fn:tokenize(fn:substring-after($weblogicVersion, &quot;Web-
Logic Server &quot;), &quot;; &quot;);
    let $majorVersionParts := fn:tokenize($spaceParts[1], &quot;.\&quot;);
  return
```

```

        $majorVersionParts[1] cast as xs:integer
let $SP_MP :=
    if ($majorVersion = 8) then
        "SP"
    else
        if ($majorVersion >= 9) then
            "MP"
        else ""
let $minorVersion :=
    let $spaceParts := fn:tokenize(fn:substring-after($weblogicVersion, "Web-
Logic Server "), " ")
    let $minorVersionParts := fn:tokenize($spaceParts[1], ".")
    return
        $minorVersionParts[2] cast as xs:integer
let $servicePackVersion :=
    let $spaceParts := fn:tokenize(fn:substring-after($weblogicVersion, "Web-
Logic Server "), " ")
    let $servicePackParts := fn:substring-after($spaceParts[2], $SP_MP)
    return
        if ($servicePackParts = "") then
            0
        else
            $servicePackParts cast as xs:integer
where $majorVersion = $major and $minorVersion = $minor and $servicePackVersion
= $servicePack and
    fn:contains(fn:upper-case($ServerRuntime/@weblogicVersion), fn:upper-
case($CrNumber))
return
    $ServerRuntime
};
for $server in local:getServerRuntimesEqualToVersionWithPatch(/,
10,0,1,"",CR366527)
|
local:getServerRuntimesEqualToVersionWithPatch(/,10,0,0,"",CR366527)
return <Server
Name="{fn:data($server/@Name)}" />
</XQueryLogic>
</SignatureDefinition>

```

Effectively, this definition collects the server name and WebLogic version of all runtime servers. Much of the definition iterates over the preciseness of the version-major and minor patch, service pack, CR number, and so forth. A violation occurs if any server has either of the stated patches (10.0.1 CR366527 or 10.0.0 CR 366527), in which case return the name of the server to be reported in violation. Hence, the rule definition must include a column to account for display of the server name. The version is irrelevant in the context of the display. Those alerted are interested only in which servers are in violation.

Important Prerequisite Steps to Use WebLogic Server Signature Rules

The following are some required steps that are specific to the version of WebLogic you are trying to monitor:

1. WebLogic versions earlier than 10.3.3: To enable data collection for the WebLogic Server signature-based rules on WebLogic Server targets earlier than v10.3.3, you need a copy of bea-guardian-agent.war. You can find a copy of this war file in your OMS installation's work directory:

```
$T_WORK/middleware/wlserver_10.3/server/lib/bea-guardian-agent.war
```

2. WebLogic Server v9 and v10.0: Install and deploy bea-guardian-agent.war to all servers in the domain. Do not change the context root. For more information on

installing a web application, see: <http://<host>:<port>/console-help/doc/en-us/com/bea/wlserver/core/index.html>

3. WebLogic Server v10.3 up to and including v10.3.2: Copy the war file from your OMS installation into each target's \$WL_HOME/server/lib directory. Restart all the servers in the target domain.
4. WebLogic Server v.10.3.3 and higher: No action is required.

Additional Notes for WebLogic Server Signature Rules

- All rules are visible in the global rule library and are visible to all users.
- After the compliance standard rule is created, it is not automatically evaluated. Users must associate a rule to a compliance standard before it can be used. Only when a compliance standard is associated with one or more targets will a rule evaluation occur. Rules cannot be evaluated directly.
- One rule can be associated to multiple compliance standards.
- Various attributes of a rule can be customized through the compliance standard this rule is associated with. These customizations occur in the Compliance Standard screens. One of these attributes that can be customized per compliance standard is the importance of the rule in relationship to this standard.
- Because the user-defined compliance standard rule is defined by a privileged user, only privileged users can modify the compliance standard rule. Violation results are available to all users.
- To share this user-defined compliance standard rule with other privileged users, provide the XML schema definition (using the Export feature) so they can import the compliance standard rule to their Management Repository.
- You can minimize scrolling when reading the Description, Impact, and Recommendation information by restricting the text to 50 characters per line. If more than 50 characters are needed, start a new line to continue the text.
- Look at the context-sensitive help for information for each page in the Compliance Standard Rule wizard for specific instructions.
- Look at the context-sensitive help for information for each page in the Compliance Standard Rule wizard for specific instruction.

Creating a Real-time Monitoring Compliance Standard Rule

To create a Real-time monitoring compliance standard rule to monitor for user actions that occur on a target such as file changes, user access, and process activity, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select Real-time Monitoring type.
5. Click **OK**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - Rule Name
Provide a unique name for the rule.

- **Compliance Rule State**
Set whether the state of this rule is development or production. Development means that the rule is still being defined or tuned and is not yet ready to be used on targets yet. After you promote a rule to production, you cannot change it back to development.
- **Severity**
The rule can have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score along with the importance that may be set for this rule when it is added to a compliance standard.
- **Applicable To**
Target type this rule works against.
- **Entity Type**
A type of object that is part of a target being monitored. For example, for the Operating System (OS), entity type may be OS File, OS Process, or OS User. For Database, an entity type may be Database Table, Database Function, Database Procedure, or Database User.
- **Target Property Filter**
You can specify specific target properties that determine which targets this rule can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this rule, for instance for Linux OS, it will only be applicable to targets on Linux Operating System.
- **Description**
Description of the rule
- **Rationale**
Text describing what this rule is checking and what the effect of a violation of this rule may be.
- **Details URL**
URL to a document that describes the compliance control in more details. Many times these documents may be stored in a content management system.
- **Message**
The message that will be used for the violation when an observation is determined to be unauthorized.
- **Clear Message**
The message that will be used for a previous violation after it is cleared.
- **Keywords**
Keywords can be assigned to a rule so that you can control how data is organized in various reports.

For additional information, see [Importance of Target Property Filters for a Real-time Monitoring Rule](#).

7. Click **Next**.

8. On the next page, you select the facets that are to be monitored for this rule. You can include facets that are already defined or create a new facet inline with this rule creation. A facet is simply a list of patterns to monitor. For instance, a list of files, user names, processes, and so on. Facets are discussed later in the section Real-time Monitoring Facets.
9. Click **Next** after including existing facets or adding new facets.
10. On the next screen, you will choose the actions you want to monitor. The actions you choose will depend on what entity type you chose for the rule. For instance, for OS File Monitoring, you can watch for actions such as file create, modify, delete, rename, and so on. For OS User monitoring, you can watch for actions such as login, logout, SU, SSH, and so on. You must choose at least one action to monitor for a rule.

For additional information, see [Selecting the Types of Actions You Want to Monitor](#).

11. Click **Next**.
12. On the next screen, you can optionally configure filters for monitoring. Filters are used to limit when or under what conditions you want an action to be observed. For instance, if you are monitoring a file facet FILES1, you can add a filter so that only file changes done by a specific list of users are captured, or if the change happens during a certain time window, or a certain process is used to modify the file. Filters are also facets, just of different entity types. If you are monitoring OS File entity type, you can apply an OS User, OS Process, or Time Window facet as a filter. You can include an existing facet, or create a new facet inline with the rule creation. If you cancel the rule wizard, any facet you created inline will still exist in the facet library.

For additional information, see [Using Facets as Filters in Real-time Monitoring Rules](#).

13. Click **Next**.
14. On the next screen, you can configure several settings related to how the observations are handled when detected at the Management Agent.
 - Authorize Observations Manually
 - Authorize Observations Automatically using Change Request Management System
 - Collection Settings

For additional information, see [Configuring Audit Status](#) and [Controlling Observation Bundle Lifetimes](#).

15. Click **Next**.
16. On this screen you can review the settings of the rule.
17. Click **Finish** to save the rule and return to the rule listing page.

Importance of Target Property Filters for a Real-time Monitoring Rule

When creating a rule, you must choose a target type for the rule. Since the Real-time monitoring capabilities on the Management Agent have some dependencies on operating system and versions of operating systems, you must be allowed to set the criteria for a rule. The target may be different on a target type, so patterns in the facets may be different. For instance, Oracle Database on Microsoft Windows is not the same as it is on the UNIX operating system.

If target property filters are not set, all rule options are available then at target-cs association time, if a target's settings do not match, then that rule and facet is ignored. If you only set, for example, the platform name, but not version, then only the options that are common across all versions of the platform are available.

The list of facets that are selectable when creating a rule are filtered by the target properties that are set when a facet is created. For instance if you have a facet, FACET1, that works on Linux or HP-UX and you create a rule for Windows, FACET1 will not be available to select for your rule. This applies both when selecting the monitoring facet or using a facet as a filter. However if you create a rule for either Linux or HP-UX, FACET1 will be available because the criteria for the rule at least overlapped with that of the facet.

Using Facets as Filters in Real-time Monitoring Rules

When creating a rule, facets can be used in two ways. The first is to use the facet to specify what entities to monitor in the rule. The second is to use the facet as a filter to apply on top of activities detected by the Management Agent.

You can use the same facet as a monitoring facet in one rule and a filtering facet in another rule. The benefit is once you define a collection of patterns, for example to define your administrative users, you can use that collection in many ways without having to redefine the collection again.

Filters in rules are set up to reduce the observations that are captured and reported to Cloud Control. If there are no filters defined, then all observations related to the monitoring facet(s) selected in the rule are captured. When selecting a facet as a filter, the default is to only include observations that have attributes that match. The following example IT compliance control demonstrates an example for the filtering:

IT Control: Monitor all changes to critical OS configuration files by administrators during production hours.

To implement this IT control, you can create a compliance standard rule with the following:

1. Create a rule and select the file facet "Critical OS configuration files" for the monitoring facet that has patterns covering all critical OS configuration files.
2. Select "content change" as the action types to capture.
3. Add an OS Users filter selecting facet "Administrators" that lists patterns describing all of the OS user accounts that are considered administrators.
4. Add a Time Window filter selecting facet "Production Hours" that lists patterns describing the times of the week that are considered to be production hours. For example, Every day 4am-2pm PST.

When the Management Agent sees any content change to the patterns in Critical OS configuration files, it will only report these changes back to Cloud Control if the change happened during production hours and if any user described in the Administrator's facet is the one making the change. Filters can also be inverted to monitor anyone not in the administrators group or for changes outside of production hours.

More details on how to use filters is described in the section above on Creating a Real-time monitoring rule.

Configuring Audit Status

Each observation can have an audit status. This audit status can change over time and be set manually or automatically by Cloud Control. The way audit statuses are managed is configured when creating or editing a real-time monitoring rule.

When creating a rule, on the settings page of the wizard, the user has an option of choosing whether all observations detected against this rule will get their audit status manually from the user or automatically using connector integration with a Change Request Management server.

When the user chooses to manually set audit status in a rule, there are two options available:

- Default Audit status can be set so that all observations that are found against this rule are by default unaudited, authorized, or unauthorized. Unaudited is the same as saying they have not been reviewed and there has been no determination of whether the observation is good or bad.
- The user can choose an informational event during manual authorizations. This is used to create a new event of informational class in the Incident Manager when a new observation bundle occurs. Based on this event, an event rule could be created to send a notification based on the observation bundle or perform any other action the Incident Manager can perform.

If the user chooses to use automatic reconciliation using a Change Request Management server, then steps must be taken to set up the Cloud Control connector for Change Management. This is explained in detail in the later section, Additional Setup for Real-time Monitoring.

Once the connector has been configured, there will be a drop down in this settings step of the rule creation wizard to choose which connector to use for this rule. Based on attributes of the observation and observations defined in any open change requests, the observation will be automatically determined to be authorized if there are open matching change requests, otherwise it will be considered unauthorized.

When using automatic reconciliation, an additional option is available to specify that the details of any authorized observations should be annotated back into the change request in the Change Request Management Server that allowed the observation to be authorized.

Multiple observations can belong to the same Observation Bundle. Even though an observation is part of group, the determination of authorized versus unauthorized is done for a single observation, not at the group level. If a group has at least one observation that is marked as "unauthorized", then the group is considered to be a "violation" and an event or incident can be raised for this group violation.

Controlling Observation Bundle Lifetimes

Observation bundles are logical groupings of observations that occur over a relatively short period of time against the same rule on the same target and by the same user. The last three factors cannot be configured by the user because they will be how the Management Agent groups observations before sending them back to the Cloud Control server.

The user creating the rule however does have three variables that they need to be able to configure:

1. Idle timeout: The amount of time after the user has no more activity from their last activity against a specific rule on a given target. The use case for this is that a user logs into a server, starts making a few file changes and then no more file changes

are made after 15 minutes. This 15 minute waiting period is the idle timeout. After this idle timeout period is reached, the current observation bundle is closed and sent to the Cloud Control server. The next time a new observation is detected, a new group will be started and the process starts over.

2. **Maximum lifespan of a group:** If a user were to set the idle timeout to 15 minutes and a user on a host was making one file change every 10 minutes for an indefinite period of time (say through a script or even manual), the observation bundle will never close and therefore never get sent to the Cloud Control server for reporting/processing. Setting the maximum lifespan of a group tells the Management Agent to only allow a group to accumulate for a maximum specific time. For example, this maximum lifespan may be 30 minutes or an hour.
3. **Maximum number of observations in a bundle:** If a rule is being triggered because of an activity that is causing a lot of observations to be detected, it may be desirable for the user to not bundle every observation together if there are too many. Bundles have a management lifecycle to them where observations can be set to authorized/unauthorized, after they arrive at the Cloud Control server. Having observation bundles with tens of thousands of observations could become hard to manage.

The user creating a rule cannot choose to turn off bundling, but if they desired to reduce delays in observation reporting to Cloud Control server, they could set the idle timeout and maximum lifespan of a bundle to be lower.

The event/incident subsystem will track only the observation bundles, not each individual observation. If one observation is marked as unauthorized, then the entire bundle will be in violation. This bundle is the entity that will be tracked by the Incident Management event.

Observation bundles are built at the Management Agent and will only be sent to the Cloud Control server when the bundle is complete according to the above criteria. In most compliance use cases, this is acceptable because you will not need to view the results immediately. Capturing and bundling results together is more important for understanding what is happening and making observations easier to manage.

When an observation becomes part of two or more bundles on the Management Agent because the same facet is used in multiple rules or multiple targets on the same host monitor the same facet with shared entities, then whenever the first bundle either hits its ending criteria (idle timeout, group maximum life, or maximum group entries), then all of the bundles containing these shared observations are closed at the same time.

To control observation bundle lifetimes, see the section above on how to create Real-time Monitoring Rules and set the appropriate settings on Settings page of the rule creation wizard.

Selecting the Types of Actions You Want to Monitor

When creating a rule, you can decide which types of observations or user actions are important to be monitored and reported back to Cloud Control. The Management Agent has a specific set of observations that are possible for each entity type. Some options may be specific to certain operating system platforms or versions. You can select one or more of these options.

The observation types that you may be able to select can also be limited by the target properties/criteria selected for the rule. For instance, some operating systems may not have every monitoring capability for files. When building the list of available observation types available, the target type, entity type, and target properties are all taken into consideration to come up with the resulting available observation types.

To select the type of observations you want to monitor in a rule, follow these steps:

1. If you want to select observations for a currently existing rule, click on the Real-time Monitoring rule in the Rules table and then click **Edit**.

Cloud Control opens the Edit Rule: Real-time Monitoring wizard and displays the Details page. Move to the Observations page.

If you want to select observations while creating a new rule, click **Create** to create a new rule. Cloud Control opens the Create Rule: Real-time Monitoring wizard and displays the Details page. After entering relevant information on the Details and Facets pages of the wizard, move to the Observations page.
2. On the Observations page, select one or more activities to be observed from the list that appears. During target association for this rule, auditing must be enabled to capture selected details. It is important to note that different operating systems and different capabilities have specific auditing requirements.
3. In the Parameters section, if there are additional observation parameters, you can review and update the parameters.

Additional Notes for Real-time Monitoring Rules

- All Rules are visible in the global rule library and are visible to all users.
- Once the compliance standard rule is created, it is not automatically evaluated. Users must associate a rule to a Compliance Standard before it can be used. Only when a compliance standard is associated with one or more targets will a rule evaluation occur. Rules cannot be evaluated directly.
- One rule can be associated to multiple compliance standards.
- Various attributes of a rule can be customized through the compliance standard this rule is associated with. These customizations occur in the Compliance Standard screens. One of these attributes that can be customized per compliance standard is the importance of the rule in relationship to this standard.
- Because the user-defined compliance standard rule is defined by a privileged user, only privileged users can modify the compliance standard rule. Violation results are available to all users.
- To share this user-defined compliance standard rule with other privileged users, provide the XML schema definition (using the Export feature) so they can import the compliance standard rule to their Management Repository.
- You can minimize scrolling when reading the Description, Impact, and Recommendation information by restrict the text to 50 characters per line. If more than 50 characters are needed, start a new line to continue the text.
- Look at the context-sensitive help for information for each page in the Compliance Standard Rule wizard for specific instructions.
- If you choose to monitor OS File entity type, you will notice one action type "File Content Modified (successful) - Archive a copy of the file [Resource Intensive]". If you select this option, every time a file modify action is observed, a copy of the file will be archived locally on the Management Agent. This can be used later to visually compare what changed between two versions of the file. There is an additional setting to set how many archived copies to store on the Actions to Monitor page of the rule creation wizard.
- When you add a facet inline with the create rule wizard either as a monitoring facet or as a filtering facet, if you cancel the rule wizard, the newly created facets will still exist and be usable in future rules. You can delete these facets by going to the

facet library. Real-time monitoring facets are discussed in a separate section later in this document.

Creating an Agent-side Rule

Note: Before you create an agent-side rule, you must create a configuration extension.

To create an agent-side compliance standard rule to check if a target has the desired configuration state based on collected configuration data, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select **Agent-side Rule** as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - **Rule Name**
Provide a unique name for the rule.
 - **Compliance Rule State**
Set whether the state of this rule is development or production. Development means that the rule is still being defined or tuned and is not yet ready to be used on targets yet. After you promote a rule to production, you cannot change it back to development.
 - **Severity**
The rule can have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score along with the importance that may be set for this rule when it is added to a compliance standard.
 - **Applicable To**
Target type this rule works against.
 - **Target Property Filter**
You can specify specific target properties that determine which targets this rule can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this rule, for instance for Linux OS, it will only be applicable to targets on Linux Operating System.
 - **Description**
Description of the rule
 - **Rationale**
Text describing what this rule is checking and what the effect of a violation of this rule may be.
 - **Recommendation**
Recommendation text describing how to fix a problem when a violation occurs.

- **Reference URL**
URL to a document that describes the compliance control in more details. Many times these documents may be stored in a content management system.
 - **Keywords**
Keywords can be assigned to a rule so that you can control how data is organized in various reports.
7. Click **Next**.
 8. On the Check Definition page, provide the configuration extension details by selecting the appropriate Configuration Extension-Alias Name from the drop-down list.
 9. Enter Compliant and Non-Compliant Message. These are the messages that will be shown in regards to the evaluation. When a violation occurs, the Non-Compliant message will be the string describing the event under the Incident Management capabilities.
 10. Click **Next**.
 11. The Text screen allows you to test your rule. You can choose a target in your environment and click the **Run Test** button. Any issues with the rule will be displayed and you can resolve them before saving the rule.
 12. Click **Next**.
 13. The final page allows you to review everything you have configured for this rule. Ensure that everything is correct and click the **Finish** button to save the rule.

Creating a Manual Rule

To create a manual compliance standard rule to check if a target has the desired configuration state based on collected configuration data, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select Manual Rule as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - **Rule Name**
Provide a unique name for the rule.
 - **Compliance Rule State**
Set whether the state of this rule is development or production. Development means that the rule is still being defined or tuned and is not yet ready to be used on targets yet. After you promote a rule to production, you cannot change it back to development.
 - **Severity**
The rule can have a severity level, which could be Critical (serious issue if this rule is violated), Warning (not a serious issue if violated), or Minor Warning (a minor issue if violated). Severity impacts the compliance score along with the

importance that may be set for this rule when it is added to a compliance standard.

- **Applicable To**

Target type this rule works against.

- **Target Property Filter**

You can specify specific target properties that determine which targets this rule can work against when it is associated with a compliance standard. These properties are Operating System, Target Lifecycle State, Version, and Platform. When you specify a target property filter for this rule, for instance for Linux OS, it will only be applicable to targets on Linux Operating System.

- **Description**

Description of the rule

- **Rationale**

Text describing what this rule is checking and what the effect of a violation of this rule may be.

- **Recommendation**

Recommendation text describing how to fix a problem when a violation occurs.

- **Compliant Message**

This message displays when the target is compliant.

- **Non-Compliant Message**

When a violation occurs, the Non-Compliant message will be the string describing the event under the Incident Management capabilities.

- **Reference URL**

URL to a document that describes the compliance control in more details. Many times these documents may be stored in a content management system.

- **Keywords**

Keywords can be assigned to a rule so that you can control how data is organized in various reports.

7. Click **Finish**.

Creating a Missing Patches Compliance Standard Rule

To create a missing patches compliance standard rule to detect patches that have not been applied to the appropriate targets, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select **Missing Patches Rule** as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - Rule

Provide a descriptive name for the rule, for example, DBMS Patches.

This is a required field.

- Compliance Rule State

- Development

Indicates a compliance standard rule is under development and that work on its definition is still in progress. While in development mode, a rule cannot be referred from production compliance standards. Use Development until the rule has been developed and tested.

- Production

Indicates a compliance standard rule has been approved and is of production quality.

You can edit a production rule to create a draft from a production rule and update the draft rule, test it, and then make it production and then overwrite/merge it back to the original production rule. This will make all the compliance standards, referring to the original production rule, to see the new definition of the rule (after overwrite).

- Severity

Minor Warning, Warning, Critical

- Applicable To

Type of target the rule applies to, for example, Database Instance. This is a required field.

- Target Property Filter

In addition, you can choose target properties by which to filter the data.

You can modify the target properties by selecting **Targets** on the Enterprise Manager menu, then the target type, for example, Database Instance. Choose the appropriate target. On the resulting page, expand the menu at the top left of the target's home page, select **Target Setup**, then select **Properties**.

- Version Name

- Platform Name

- Lifecycle State

- Description and Rationale

Provide complete and descriptive information for all explanatory fields, for example, description, rationale (reason for the rule), recommendations (how to fix the problem denoted when this rule is violated), and so on.

- ReferenceUrl

This URL should reference information that is pertinent to this rule.

- Keywords

Add Keywords to further categorize the compliance standard rules. Choose one or more keywords that closely match your rule's intent.

7. Click **Next**.

8. On the Define Patch Check page:

- Select recommended patches from a table or from a list of patches.

- Provide the text for the compliant and non-complaint messages.

Element	Description
Compliant Message	A compliance standard rule is compliant when the SQL query does not return result data. If a user has preferences to be notified when a compliance standard rule is cleared, this is the message he or she will receive for compliance. Default: Compliance standard rule <name of compliance standard rule> is compliant. You can override the default text.
Non-Compliant Message	A compliance standard rule is non compliant when the SQL query returns result data. If no data is returned, the compliance standard rule is compliant. This message is used in notification rules. If a user has preferences to be notified for compliance standard rule violations, this is the message he or she will receive for violation. Default: Compliance standard rule <name of compliance standard rule> is not compliant. You can override the default text

9. Click **Next**.

10. On the Test page, validate whether a patch was applied to a particular target. This test evaluation is not stored in the Management Repository and is a one-time run. If there are no errors, the compliance standard rule is ready for publication or production.

Note: You can have test results that intentionally show violations. For example, if you are testing `target_type` equal to host and you are evaluating a host target, then you will see violation results.

Rule Violations

Provides the details of a compliance standard rule violation. This is the same information you see on the Violation Details drill-down page in the Compliance Standard Rules Errors page.

11. Click **Next**.

12. On the Review page, verify that the information on the page reflects what you intended to supply in the definition.

If corrections are needed, click **Back** and make the needed corrections.

13. Click **Finish**.

Note: The compliance standard rule is not defined until you click **Finish**.

Tips

- Once the compliance standard rule has been created, it is not automatically evaluated. Consider adding the compliance standard rule to a compliance standard.
- Assign a corrective action to the rule after the rule has been created.
 - On the Compliance Standard Rules tab, highlight the rule you just created.
 - From the **Actions** menu, select **Assign Corrective Action**.

- From the **Assign Creative Action** popup, select an existing corrective action and click **OK**.

Creating a Configuration Consistency Rule

To create a configuration consistency compliance standard rule to determine the consistency of targets of similar target types within a composite target, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select **Configuration Consistency Rule** as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:
 - **Rule**
Provide a descriptive name for the rule, for example, DBMS Consistency.
This is a required field.
 - **Compliance Rule State**
 - **Development**
Indicates a compliance standard rule is under development and that work on its definition is still in progress. While in development mode, a rule cannot be referred from production compliance standards. Use Development until the rule has been developed and tested.
 - **Production**
Indicates a compliance standard rule has been approved and is of production quality.

You can edit a production rule to create a draft from a production rule and update the draft rule, test it, and then make it production and then overwrite/merge it back to the original production rule. This will make all the compliance standards, referring to the original production rule, to see the new definition of the rule (after overwrite).
 - **Severity**
Minor Warning, Warning, Critical
 - **Description**
Provide complete and descriptive information.
 - **Applicable To**
Type of target the rule applies to, for example, Database Instance. This is a required field.
 - **Comparison Template**
This is a required field.
 - **Target Property Filter**
You can choose target properties by which to filter the data.

You can modify the target properties by selecting **Targets** on the Enterprise Manager menu, then the target type, for example, Database Instance. Choose the appropriate target. On the resulting page, expand the menu at the top left of the target's home page, select **Target Setup**, then select **Properties**.

- Operating System
- Target Lifecycle State
- Version
- Platform

- Rationale

Provide complete and descriptive information about the importance of the rule.

- Keywords

Add Keywords to further categorize the compliance standard rules. Choose one or more keywords that closely match the rule's intent.

7. Click **Finish**.

Creating Configuration Drift Rule

To create a configuration drift compliance standard rule to determine the deviation of targets of similar target types, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Click the **Create** button.
4. In the Create Rule popup, select **Configuration Drift Rule** as the type.
5. Click **Continue**.
6. On the next screen, you are asked to fill out several key attributes of the rule:

- Rule

Provide a descriptive name for the rule, for example, DBMS Drift.

This is a required field.

- Compliance Rule State

- Development

Indicates a compliance standard rule is under development and that work on its definition is still in progress. While in development mode, a rule cannot be referred from production compliance standards. Use Development until the rule has been developed and tested.

- Production

Indicates a compliance standard rule has been approved and is of production quality.

You can edit a production rule to create a draft from a production rule and update the draft rule, test it, and then make it production and then overwrite/merge it back to the original production rule. This will make all the compliance standards, referring to the original production rule, to see the new definition of the rule (after overwrite).

- **Severity**
Minor Warning, Warning, Critical
- **Applicable To**
Type of target the rule applies to, for example, Database Instance. This is a required field.
- **Comparison Template**
This is a required field.
- **Source Configuration**
 - Latest Configuration
 - Saved Configuration
- **Target Property Filter**
You can choose target properties by which to filter the data.

You can modify the target properties by selecting **Targets** on the Enterprise Manager menu, then the target type, for example, Database Instance. Choose the appropriate target. On the resulting page, expand the menu at the top left of the target's home page, select **Target Setup**, then select **Properties**.
 - Operating System
 - Target Lifecycle State
 - Version
 - Platform
- **Description and Rationale**
Provide complete and descriptive information for all explanatory fields, for example, description, rationale (reason for the rule), recommendations (how to fix the problem denoted when this rule is violated), and so on.
- **Keywords**
Add Keywords to further categorize the compliance standard rules. Choose one or more keywords that closely match the rule's intent.

7. Click **Finish**.

Creating Like a Compliance Standard Rule

To create a compliance standard rule like another compliance standard rule, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to replicate.
4. Click **Create Like** button.
5. Customize the fields as needed.
6. Click **Save**.

Editing a Compliance Standard Rule

To edit a compliance standard rule, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to edit and click the **Edit** button.
4. Step through the screens of the rule creation wizard as previously described when creating a rule.
5. Click **Save**.

Usage Notes

- For repository rules, you can change all the rule properties except the Rule Name, State (if it is already production), and Applicable To.

For real-time monitoring rules, you cannot change Rule Name, State (it is already production), Applicable To, Target Property Filters, and Entity Type.
- If you change the critical rule properties for a repository rule, for example, rule query, violation condition, parameters, or severity, then editing the rule invalidates the results for compliance standards which refer to the rule. The compliance standards compliance score will be reevaluated at the next rule evaluation.
- For rules in production mode, you have a choice to create and save a draft of the rule or to overwrite the existing production rule. If you create a draft, you can edit the draft rule, at a later point in time, test it, and then overwrite and merge it back to the original production rule the draft was made from. **Note:** You cannot include a draft rule into any compliance standard.
- For WebLogic Server Signature rule or Real-time Monitoring rule, if the rule being edited is referred to by a compliance standard which is associated with a target, then the rule definition will be deployed to the Management Agent monitoring the target, so that the Management Agent can evaluate the latest definition of the rule. In the case where the Management Agent is down or unreachable, the rule definition changes will be propagated to the Management Agent as soon as the Management Agent is available.

Deleting a Compliance Standard Rule

Before you delete a rule, you must ensure that compliance standard rule references have been removed from compliance standards before deleting the compliance standard rule. You cannot delete a rule that is in use by a compliance standard.

To delete a compliance standard rule, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to delete, click **Delete** button.
4. Confirm that you want to delete the rule by clicking **OK**.

Exporting a Compliance Standard Rule

The Export feature provides a mechanism for transporting user-defined compliance standard rule definitions across Management Repositories and Cloud Control instances. The export stores the definitions in an operating system file. Because the exported compliance standard rule definitions are in XML format, they conform to the Oracle Compliance Standard Definition (XSD) format. You can then change the definition of the compliance standard rule and re-import the generated compliance standard rule definitions into another Management Repository.

To export a compliance standard rule, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. Highlight the rule you want to export.
4. From the **Actions** menu, select **Export**.
5. Provide the file name to which the standard rule is to be exported.
6. The XML representation of the compliance standard rule is generated and placed in the directory and file you specified.

Importing a Compliance Standard Rule

Importing allows you to re-use a compliance standard rule that you already have, share rule definitions across multiple instances of Cloud Control, or enable offline editing of the rule.

Before you import a compliance standard rule, ensure the compliance standard rule to be imported is defined in a file. The file should be locally accessible to the browser you are using to access Cloud Control. Also ensure that you have privileges to access the compliance standard rule definition XML file to be imported.

To import a compliance standard rule, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. From **Actions** menu, select **Import**.
4. Provide the file name from which the rule definition (as per Compliance Standard Rule XSD) will be imported. Specify whether to override an existing definition if one already exists. The override option is not available to Real-time monitoring rules.
5. Click **OK**.

Browsing Compliance Standard Rules

To browse compliance standard rules, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. To view the details of a particular standard rule, highlight the rule and click **Show Details**.

Searching Compliance Standard Rules

To search for compliance standard rules, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Click the **Compliance Standard Rules** tab.
3. In the Search portion of the page, provide criteria to use to narrow the search.

By default, all the compliance standard rules in the compliance standard rule library appear in the results table. However, you can specify a set of search criteria and then perform a search that will display only the compliance standard rules that meet those criteria in the results table.

For example, if you choose Security in the Category list, contains in the Compliance Standard Rule list, "port" in the adjacent Compliance Standard Rule text field, Host in the Target Type list, and then click **Go**, Cloud Control displays only the compliance standard rules for the host security category that contain "port" in their names.

4. Click **Search**.

Using Corrective Actions

A corrective action is a script that fixes the problem causing a violation to a compliance standard rule.

There are two types of corrective actions:

- Manual - Created in the context of the compliance standard rule.
- Automatic - Created in the context of an incident rule.

Manual Corrective Action

To create a corrective action manually, perform the following steps:

1. From the **Enterprise** menu, select **Monitoring**, then select **Corrective Actions**.
2. On the Job page:
 - a. Select **SQL Script** in the Create Library Corrective Action field, and click **Go**.
 - b. On the General tab, type a name for the corrective action (for example, CA1), provide a description, and select **Compliance Standard Rule Violation** as the Event Type. Select **Database Instance** as the Target Type.
 - c. On the Parameters tab, select the default: **WHENEVER SQLERROR EXIT FAILURE**; Click **Save to Library**.

Note: To enable intelligent remediation, pass parameters from the compliance violation to the corrective action. For example, to lock changes to Well Known Accounts, add the following SQL statement:

```
alter user %EVTCTX.dbuser% account lock;
```

where dbuser is the event context parameter

You can make similar changes to any parameter. Ensure that the parameter name matches the name of the column in the SQL query.

- d. Select the corrective action you just created and click **Publish**.
- e. On the confirmation page, click **Yes**.
3. From the **Enterprise** menu, select **Compliance**, then select **Library**. Choose a database compliance standard rule with the rule type of agent-side or repository. In the **Actions** menu, select **Assign Corrective Action**. Select a corrective action and click **OK**.

You will then see the corrective action in the Show Details page for the compliance standard rule.

Automatic Corrective Action

To create a corrective action that is automatically triggered when the violation occurs, follow these steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the **Incident Rules - All Enterprise Rules** page, click **Create Rule Set**. Provide a name for the rule, select **All targets** in the Targets region, and click **Create...** in the Rules region.
3. On the Select Type of Rule to Create dialog box, select **Incoming events and updates to events**. Click **Continue**.
4. For the type, select **Compliance Standard Rule Violation**.
5. Select either **All events of type Compliance Standard Rule Violation** or **Specific events of type Compliance Standard Rule Violation**.
6. In the Advanced Selection Options, select **Corrective action completed**. Click **Next**.
7. On the Create New Rule: Add Actions page, click **Add**. On the Add Conditional Actions page, click **Select corrective action**. Select the corrective action. Click **Continue**.
8. In the Create New Rule: Add Actions page, click **Next**. Provide a description on the Create New Rule: Specify and Description page and click **Next**.
9. Review the information and click **Continue**.
10. Click **Save**. Note that newly added rules are not saved until the **Save** button is clicked. After you click Save, verify that the rule set entity has added the new incident rule by reviewing the details.

Real-time Monitoring Facets

The real-time monitoring rule definition includes facets that are used to determine what is important to monitor for a given target type, target properties, and entity type. A facet is a collection of patterns that make up one attribute of a target type.

The following sections explain real-time monitoring facets in detail:

- [About Real-time Monitoring Facets](#)
- [Operations on Facets](#)

About Real-time Monitoring Facets

A target type has several facets to it. A target type will have a facet of which files are critical configuration files, which files are log files, which files are executables, which database tables have sensitive configuration data, and so on. The sum of all of these facets for a given target type makes up everything that is important to monitor for the given target type in terms of compliance.

For a given target type, you can create any number of facets. A facet is not only for a specific target type, but for a specific target type plus a combination of some number of target type properties. For instance, creating a facet for a Host Target Type on Windows is different than creating a facet for a Host Target type on Linux. A facet can have several target type properties or can be open to any target without specifying any properties.

Facets are reusable in many rules. The benefit is that you can add or remove entries from a facet without having to modify every rule. For instance, if today there are 5 log files you want to monitor, you can setup your rules to monitor a facet listing those 5 files. When a new log file should be added tomorrow, you only need to change the facet, not each rule.

Facets can be created on their own, or created inline with a Real-time Monitoring rule creation. No matter how they are created, they can be used again at a later time in any number of rules.

Real-time Monitoring facets based on target types are used to specify the entities to monitor in real-time monitoring rules. As an example, if monitoring a host for file changes, a facet can be a list of distinct single files, patterns with wildcards that would include many files, or simply an entire directory. These patterns can also include parameters that have a default, but can be overridden as needed for each target. Built-in parameters, such as `ORACLE_HOME` will be dynamically filled in for each target. If you wanted to specify monitoring the database configuration file *tnsnames.ora*, your pattern may be `{ORACLE_HOME}/network/admin/tnsnames.ora`.

Facets can be used in two totally distinct ways. Primarily, facets describe what to monitor. In the rule creation wizard, these facets are selected on the wizard step "Entities to Monitor". Facets also can be used to filter your monitoring results. These filtering facets are specified on the Filters step of the rule creation wizard. When monitoring an OS file entity type for instance, you can filter your results based on the user that made a file change, the time the file change happened, or the process used to make the file change.

When performing continuous real-time monitoring, it is important to scope your monitoring only to critical entities. Monitoring more activity than is important to the organization will result in higher CPU loads on the Management Agent as well as a very large amount of data to be processed/stored by the Oracle Enterprise Manager servers.

Facet Entity Types

Each facet has an *entity type* which defines what kind of entities the facet describes. For example, for OS level monitoring, there are OS File, OS Process, OS User, Windows Registry, and several Active Directory entity types. For database monitoring, the entity types include Table, View, Index, Procedure among others. The possible entity types are fixed by the continuous real-time configuration change monitoring capabilities available from the Management Agent.

Creation of facets is possible through the Facet Library screen. In this screen, you can add/edit patterns for facets, and see which facets are being consumed by rules.

Given below are the entity types Cloud Control supports for real-time monitoring:

- OS File
- OS Process
- OS User
- Microsoft Windows Registry
- Microsoft Active Directory User
- Microsoft Active Directory Computer
- Microsoft Active Directory Group
- Oracle Database Dimension
- Oracle Database Synonym
- Oracle Database Type
- Oracle Database Table
- Oracle Database View
- Oracle Database Procedure
- Oracle Database User
- Oracle Database Index
- Oracle Database Sequence
- Oracle Database Function
- Oracle Database Profile
- Oracle Database Public Synonym
- Oracle Database Role
- Oracle Database Package
- Oracle Database Library
- Oracle Database Trigger
- Oracle Database Tablespace
- Oracle Database Materialized View
- Oracle Database Cluster
- Oracle Database Link
- Oracle Database Public DB Link
- Oracle Database Segment
- Oracle Database SQL Query Statement

Facet Patterns

A facet contains one or more patterns. These patterns can express inclusion or exclusion filters. For instance, you may define a facet for critical configuration files that looks like the following:

Include c:\myapp1\config

Exclude c:\myapp1\config\dummy.cfg

In this case, everything under c:\myapp1\config will be considered to be a member of this facet except for the individual file c:\myapp1\config\dummy.cfg. In general there are some rules to how patterns work given the most common use cases listed below. Each entity type might have special cases or special formats of patterns.

- Patterns of the same specificity with one being include and one being exclude, the include will win.
- Patterns that are more specific override (like in the previous example, exclude dummy.cfg overrides the inherited include c:\dummy.cfg from the first pattern.)
- If there are no patterns at all, exclude * is assumed (for example, no entities in the facet)

For each pattern that you add to a facet, an optional description field is available to let you document their patterns.

Operations on Facets

The following sections explain the operations you can perform on facets:

- [Viewing the Facet Library](#)
- [Creating and Editing Facets](#)
- [Creating and Editing Facet Folders](#)
- [Deleting a Facet](#)
- [Using Create Like to Create a New Facet](#)
- [Importing and Exporting Facets](#)
- [Changing Base Facet Attributes Not Yet Used In a Rule](#)

Ensure you have the privileges to create, delete, and modify facets as these configurations relate to the compliance monitoring. See [Roles and Privileges Needed to Use the Compliance Features](#) for information.

Viewing the Facet Library

Any user who can view observation data is able to also view the facet library and see the facet history for any facet.

There are two ways to view the facet library, search mode and browse mode. In search mode, all facets meeting the search criteria are shown in a flat list. In browse mode, facets are shown along with a folder hierarchy that the facets belong to. This folder structure can help users manage a very large number of facets in Cloud Control.

To view the facet library in search mode, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the Real-time Monitoring Facets Library tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export if you have the audit author role.

3. Click the **Search Facets** tab.

The Facet Library page displays the Facet Name, Author, Target Type, Entity Type, Rules Using the facet, Description, and the Last Updated time of the facet. You can see the details of any facet by selecting it from the table and clicking Show Details.

4. You can choose which columns to display in the table by clicking **View** and then choosing **Columns**. You can either choose to **Show All** columns or you can select individually the columns you want to appear in the table. You can reorder the columns by clicking **Reorder** after you click **View** and then changing the order in which the columns appear by moving them up or down using the arrow keys.
5. You can expand the area of the page titled "Search" to choose the search criteria to apply to the view of facets.
6. You can view a history of a selected facet by choosing it from the table and then clicking History. The View History page appears.

To view the facet library in browse mode, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export if you have the audit author role.

3. Click the Browse Facets tab.

The Facet Library page that is shown is split into two views. The left side shows the facet folder hierarchy. The right side lists facets in the folder that is selected on the left. The table on the left displays the Facet Name, Author, Target Type, Entity Type, Rules Using the facet, Description, and the Last Updated time of the facet. You can see the details of any facet by selecting it from the table and clicking Show Details.

4. You can choose which columns to display in the table by clicking **View** and then choosing **Columns**. You can either choose to **Show All** columns or you can select individually the columns you want to appear in the table. You can reorder the columns by clicking **Reorder** after you click View and then changing the order in which the columns appear by moving them up or down using the arrow keys.
5. The only filtering allowed on this screen is by selecting a different folder. You will always see the facets that are in the selected folder only.
6. You can view a history of a selected facet by choosing it from the table and then clicking **History**. The View History page displays.

Creating and Editing Facets

When you create a facet and subsequently use a facet in a Real-time Monitoring Compliance Standard Rule, the compliance rule only references the facet. If the content changes, then the rule will use the new content automatically.

The content of the facet only begins being used when it is added to a rule that is part of a compliance standard that is associated to one more targets.

Each facet is assigned a description that allows you to document the facet. Each pattern also has an optional description field. Only begins being used when it is added to a rule that is part of a compliance standard that is associated to one more targets.

To create or edit a facet, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export. There are two views when looking at this page, search or browse. In the search view, all facets are listed in a flat list. In the browse view, facets are grouped in folders to make it easier to find facets.

3. Click **Create** to create a new facet.
4. Choose which facet folder this facet should belong to. If you have not yet created the folder for it, you can add it to the Unfiled folder. This folder always exists and cannot be removed. Later you can move the facet to a new folder you create using drag-and-drop in the UI from the Unfiled folder to the new folder.
5. Enter the name you want to assign to the facet in the **Facet Name** field, then choose the target type for the facet you are creating from the drop-down list in the **Target Type** field. Once you choose the Target Type, you can enter values in the Target Property Filter fields.

The target properties you add here limit which targets to which this facet can ultimately be assigned. For instance, you could define a facet to work only for Linux version 5 on 64-bit servers.

6. Choose the **Entity Type** from the drop-down. This list will be limited depending on the target type chosen previously.
7. Enter a description for the facet in the **Description** field.
8. The Create Facet page contains two tabs you can use to enter the patterns and parameters for the facet you create. Use the Patterns tab to add patterns to be either **Included** or **Excluded**. Use the **Add** or **Delete** buttons to add additional patterns or to remove a selected pattern from the facet definition. There is a bulk add button which will bring up a popup window where you can paste text listing patterns rather than entering each in the UI manually.
9. If you are defining a facet for the OS File entity type, there is an optional ability to browse a host to find the files you want to monitor. The right side of the page has an area where you can choose the host to use as the basis for looking for files. In the pattern area, you can click the Browse button to interactively browse the files on the selected host and select the files to include in the pattern. After selecting patterns from a host, you can continue to manually add more or edit existing ones.
10. Use the Parameters tab to view parameters that are part of the new facet. Oracle provides a set of predefined parameters based on target parameters (such as ORACLE_HOME) that are defined out of the box. These parameters do not require a default value and are always set according to the target's value. Parameters will appear under this tab when they are used in a pattern. To start using a new parameter, simply add the parameter to the pattern by enclosing it in curly brackets {}. For instance, a pattern of {INSTALL_DIR}\config\main.conf would result in a parameter of INSTALL_DIR being listed under this tab. All parameters must have a default value that will be automatically used for all targets against

which this facet is used. This value can be overridden when associating a compliance standard containing a real-time monitoring rule to one or more targets. The Parameters tab displays the Parameter Name, Default Value, Used in Pattern, and Description. Used in Pattern indicates that the parameter is currently in use. This parameter may have been defined at some point in a pattern and then removed. The pattern will still be available for use again at a later time even if the pattern is not currently in use. If the entity for which you are adding a pattern includes a "{" or "}", you can escape these characters by using "{{}" and "}}" in the pattern respectively. These will not be counted as parameters.

11. A third tab, Time Window is only available if the facet being created/edited is of entity type Time Window. A facet of this entity type is only usable as a filter in a Real-time monitoring rule. For instance, you can specify in the rule that you only want to monitor a facet during a specific time, for example, "Production Hours". In the Duration section, choose either a **24 Hour Interval** or **Limit Hours to**, which allows you to enter a Start time and an Interval in Hours and Minutes. In the Repeating section, you can choose either *All the time* or you can select **Repeat** and then choose which days of the week to repeat the operation.

12. Choose **OK** to create the facet.

Creating and Editing Facet Folders

When viewing the facets in Browse Facets mode, you will see two regions on the page. The left side will show the facet folders which exist. The right side will show the facets that exist in the currently selected folder. On the left side showing the folders, there are three actions available for folders.

- **Create:** Allow you to create a new folder. A popup will display asking for the folder name to create. You will also have the choice of making this new folder a top level folder or adding it as a child to the currently selected folder.
- **Rename:** Allows you to rename an existing user-defined folder
- **Delete:** Allows you to delete a user-defined folder. You cannot delete a folder that has facets or other folders inside of it.

You cannot delete, rename or move out-of-the-box folders that are populated by Oracle.

There is a default folder that exists called Unfiled. Anytime a facet is created or imported without specifying a folder, it will go into this Unfiled folder.

You can move facets into folders by simply finding the facet you want to move in the right side, selecting it and dragging it to the folder on the left where you want to place it. The facet will move to that folder. A facet can only belong to one folder at a time and it always must belong to a folder (even if it is just the Unfiled folder). You can also click on the facet and click on the MOVE button. A popup window will appear letting you choose which folder to move the facet to.

Folders have no impact on observation analysis or compliance score. They are only used in the Real-Time Monitoring Facets library screen to make it easier to manage a very large number of facets that exist.

Deleting a Facet

Deleting a facet is not possible as long as the facet is in use either as a monitoring facet in a rule or as a filter facet in a rule. If this facet is not in use in any rules, then the

facet can be deleted. If a facet is in use, the user is alerted to the current use and not allowed to delete the facet until the rules using it are modified to no longer include it.

When deleting a facet, any historic observation data will no longer be referenced to the facet and instead it will show "(Deleted Facet)" as the name of the facet to which it is related. This observation data will only be available through the Search Observations page, not the Browse pages.

For compliance-focused users, customers typically would want to keep the unused facet available so the compliance data is not lost. You can also remove the patterns as long as you keep the actual facet to maintain collected observations. Then only after the compliance data related to this old facet is no longer available, you can delete the facet without any data loss.

To delete a facet, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export.

3. Select the facet from the list of facets in the table on the page.
4. Click **Delete** to delete the facet. You will be prompted to confirm that you want to delete the facet.

Using Create Like to Create a New Facet

Facets that ship with the product or with a plug-in cannot be changed. If you want to enhance or modify the Oracle provided content, you must use the create-like functionality to make your own copy of the facet which can then subsequently be edited.

An important limitation to the Create Like function is that you cannot change the target type or entity type. The patterns contained in the facet may be dependent on target type or entity type. If you want to use Create Like and change these attributes, you should use Export to export the original facet, edit the name, target type, entity type in the XML, and then import as a new facet.

To use create like to create a new facet, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

3. Choose the facet from the facet table that you want to use as the basis for the new facet you want to create.
4. Click **Create Like**.

Cloud Control displays the Create Facet page. All the values that were applicable to the facet you want to clone are entered. Use the page to edit the values for the new facet and click **OK**.

It is important to understand that if the original base facet you used in the create like activity is changed, that change will not be reflected in the newly created facet. There is no relationship maintained when using Create Like.

5. For more information about using the Create Facet page, see [Creating and Editing Facets](#).

Importing and Exporting Facets

You can select facets and export or import them. All selected facets will be exported into one output file.

On import, if a facet of the same name/target type/entity type combination already exists, the import fails with an error that the facet already exists. The user must change the import file to remove the duplicate name and retry the import.

The combination of name, target type, and entity type define a unique facet. You can have the same name facet across different target types and entity types.

To export a facet, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.
Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export.
3. Select one or more facets from the list of facets on the Facet Library page that you want to export and then click **Export**.
4. On the Open dialog box, you can choose to open or save the facet xml file using an XML editor of your choice and then either edit or save the file to another location.

To import a facet, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.
Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import, and export.
3. Click **Import** and choose the facet XML file you want to import into the Facet Library.
4. Cloud Control imports all facets specified in the imported XML file. You can then edit the facet or use any other action on it as you would any other facet in the library.

Changing Base Facet Attributes Not Yet Used In a Rule

After a facet is in use in at least one rule (either as a monitoring facet or as a filter facet), you cannot change the facet name, target type, entity type, or target criteria of the facet since the rules that have been created are already bound to these attributes. The only attributes that can be changed are the facet patterns, parameters and description

fields. Although the rule is not dependent on the facet name, users have used them in their rules based on the name of the facet. Allowing the name of the facet to change after consumption will only lead to confusion of the rule authors when analyzing compliance results and observations of the rule authors.

If a facet is not currently in use but has been in use in the past, then it is treated the same as an in-use facet since the historic observation data will still be tied to the past facet.

You cannot make changes to the Oracle provided facets that ship with the Cloud Control product. If you want to use an Oracle provided facet with changes, you can perform a "Create Like" operation and then modify the newly created facet as needed.

To change base facet attributes, follow these steps:

1. From the **Enterprise** menu, select **Compliance**, then select **Library**.
2. Choose the **Real-time Monitoring Facets Library** tab.

Cloud Control displays the Facet Library page that lists all existing facets along with their target type, entity type, and other details about the facet. From this page you can perform administrative tasks such as create, create like, view, delete, import and export.

3. Choose the facet from which you want to create a new facet with modified attributes. Click **Create Like**.
4. Enter a new Facet Name and change whatever attributes to create a new facet based on the previous facet.

Examples

This section provides examples of using compliance. Examples include:

- [Creating Repository Rule Based on Custom Configuration Collections](#)
- [Creating Compliance Standard Agent-side and Manual Rules](#)
- [Suppressing Violations](#)
- [Clearing Violations](#)

Creating Repository Rule Based on Custom Configuration Collections

This example illustrates how a compliance rule can be created and run on a custom configuration which collects a sample configuration file (for this example, /tmp/foo.xml) for targets of type Host.

For this example, create a sample /tmp/foo.xml file with following contents:

```
<some_config>
  <prop foo="1"/>
  <prop bar="2"/>
</some_config>
```

The steps include how to:

- Create a custom configuration
- Create a custom-based repository rule

- Create a compliance standard
- Associate a target
- View results

To create a custom configuration:

1. From the Enterprise menu, select **Configuration**, then select **Configuration Extensions**.
2. From the Configuration Extensions page, click **Create**. The Create Configuration Extension page appears.

- a. Type the Name (for example, `compliance_ccs`), a description (optional), select Target Type (for this example, `Host`).
- b. In the Files & Commands section, type the Default Base Directory. [Use `/tmp` as the directory.]

This is an example. For a real target it should be the directory containing the target's configuration files.

Note: All files collected by custom configurations **MUST NOT** change on a daily basis, but should only change very rarely due to an explicit action by an administrator.

- c. Click **Add**.
 - In the Type column, select **File**.
 - In the File/Command column, type **foo.xml**. The Alias column is automatically filled in with `foo.xml`.
- Note:** You can use any file or files, not just xml and not just "foo.xml" expressions. Custom configuration supports many files and corresponding parsers.
- In the Parser column, select **XML Parser (default)**.

- d. Click **Save** located at the bottom of the page.
3. In the Custom Configurations page, highlight `compliance_ccs` and click **Deploy**. The Deployments page appears.

- a. Click **Add** to select targets on which CSS needs to be deployed.
- b. On the Search and Select: Targets page, highlight the host target where file `/tmp/foo.xml` was created and click **Select**.
- c. Click **Apply** on the Deployments page.

4. On the Submit Pending Deployment Actions popup, select **Yes**. This action will submit the deployment action.

On the Deployments page, click **Refresh Status** to refresh the status of the deployment until the Status column displays "Successfully deployed".

5. Now that deployment is submitted, click **Cancel** to exit the page. (**Note:** Clicking Save instead of Apply earlier, would have exited the page right after the submission of the deployment action.)

To create a custom-based repository rule based on custom configuration collection:

1. From the Enterprise menu, select **Compliance**, then select **Library**.
2. On the Compliance Library page, click the **Compliance Standard Rules** tab.

3. Click **Create**.
 - a. On the Create Rule popup, select **Repository Rule** and click **Continue**.
 - b. On the Create Rule: Repository Rule: Details page, type in the Rule name, for our example, **compliance_css_rule**.
 - c. For the Compliance Rule State, select Development, then select Minor Warning for the Severity. For Applicable To: select Host. Click Next located at the top-right of the page.
4. On the Create Rule: Repository Rule: Check Definition (Query) page, click **Model Query**. New Search Criteria page appears.
 - a. Select compliance_css (Parsed Data) from the Configuration Item menu under "Commonly Used Search Criteria".
 - b. Under the Host section and Parsed Data subsection, type **foo.xml** for Data Source contains. For the Attribute, select **is exactly** comparison operator and type **foo** to refer to the "foo" attribute in our sample file. (**Note:** % sign can also be used as a wild card character in these expressions for Data Source and Attribute.)
 - c. Click **Search** to see the rows returned for this filter. A table displays the data with value 1 for attribute foo in our file.
 - d. Click **OK**.
 - e. The Create Rule: Repository Rule: Check Definition (Query) displays again but this time the SQL Source appears.
 - f. Click **Next**. **Note:** In general, you could also update the query before proceeding, if needed.
5. The Create Rule: Repository Rule: Check Definition (Violation Condition) page displays.
 - a. Check all the columns as Key columns (VALUE, ATTR, CONTAINER, and DATA SOURCE NAME), except the INFO column.
 - b. In the Condition Type section of the page, select Simple Condition, and in the Column Name select VALUE and change the Comparison Operator to equal sign (=). In the Default Value column, type 1. Click **Next**.
6. In the Create Rule: Repository Rule: Test page, click the icon next to Target Name field. The Search and Select: Targets popup appears. Find the host where the custom configuration was deployed. Select it and click **Select**.
7. In the Create Rule: Repository Rule: Test page, click **Run Test**. When the test runs successfully, you get a confirmation stating that the Run Test - Completed Successfully.

You should see one violation after running the test because we specified value of "1" in step 5 above for violation condition and our sample file had value "1" for attribute foo. Click **Close**.
8. On the Create Rule: Repository Rule: Test page, click **Next**.
9. In the Create Rule: Repository Rule: Review page, ensure that all the information that you added is correct. Click **Finish**.

To create a compliance standard:

1. From the Enterprise menu, select **Compliance**, then select **Library**.

2. Click the Compliance Standards tab and click **Create**.
3. On the Create Compliance Standard popup, type **compliance_css_cs** in the Name field, select **Host** from Applicable To menu, and select **Repository** as the Standard Type. Click **Continue**.
4. The compliance standard page displays with the information regarding the compliance_css_cs compliance standard. Right-click on compliance_css_cs on the left side and select the **Add Rules...** option in the right-click menu.
5. On the Include Rule Reference popup, select compliance_css_rule. Click **OK**. Click **Save** to save the compliance_css_cs.
6. A confirmation message appears on the Compliance Library page stating that the compliance standard has been created. Click **OK**.

To associate targets:

1. Select the compliance_css_cs that was just created. Click **Associate Targets**.
2. On the "Target Association for Compliance Standard: compliance_css_cs" page, click **Add** to add targets.
3. On the Search and Select: Targets page, select a target where /tmp/foo.xml is present and click **Select**. Click **OK**.

You will then be prompted whether you want to Save the association or not. Click either Yes or No. You will then get an Informational message stating that the compliance standard has been submitted to the target for processing.

To view results:

1. From the Enterprise menu, select **Compliance**, then select **Results**.
On the Compliance Results page, select the compliance_css_cs compliance standard and click **Show Details** to view the details of the compliance standard created.
2. Click the **Violations** tab associated with the compliance_css_rule. The target is associated with one violation.
3. Click on the rule node in the tree to see the **Violation Events** tab, then click on this tab to see the violation details for the rule. Click on a violations row in the violations table, to view details of the violation.

Creating Compliance Standard Agent-side and Manual Rules

The purpose of this example is to create an agent-side compliance standard rule and a manual rule that test for DBMS privileged actions.

When creating an agent-side compliance standard rule, perform the following steps:

1. Create a configuration extension
2. Create the agent-side compliance rule
3. Create a manual rule
4. Create a compliance standard
5. Add the rules to the configuration standard
6. Associate the compliance standard to a target

Creating a Configuration Extension

Perform the following steps to create a configuration extension:

1. From the **Enterprise** menu, select **Configuration**, then select **Configuration Extensions**.
2. On the Configuration Extensions page, click **Create**.
3. Type a name for the extension, for example, DG0142 DBMS Privileged action audit. You will use this name on the Check Definition page.
4. Select **Database Instance** for the Target Type.
5. Click the SQL tab.
6. Click **Add** to add the first SQL statement.

- In the SQL field, type:

```
select distinct 'Unauthorized user '||owner||' owns application objects in
the database.' value
  from dba_objects where
owner not in ('ANONYMOUS','AURORA$JIS$UTILITY$', 'AURORA$ORB$UNAUTHENTICAT-
ED', 'CTXSYS','DBSNMP','DIP','DVF','DVSYS','EXFSYS','LBACSYS','MDDATA',
'MDSYS','MGMT_VIEW','ODM','ODM_MTR', 'OLAPSYS','ORDPLUGINS', 'ORDSYS',
'OSE$HTTP$ADMIN','OUTLN','PERFSTAT', 'PUBLIC','REPADMIN','RMAN','SI_IN-
FORMTN_SCHEMA',
'SYS','SYSMAN','SYSTEM','TRACESVR', 'TSMYSWK_TEST','WKPROXY','WKSYS',
'WKUSER','WMSYS','XDB','OWBSYS', 'SCOTT', 'ORACLE_OCM', 'ORDDATA',
'APEX_030200',
'OWBSYS_AUDIT', 'APPQOSSYS', 'FLOWS_FILES')
and owner not in (select grantee from dba_role_privs
  where granted_role='DBA')
```

- Type an alias, for example, DBMS application object ownership. This alias is useful when defining the rule on top of this configuration extension.
 - For the Parser, use Database Query Parser.
7. Click **Add** to add the second SQL statement.

- SQL

```
select distinct 'Application object owner account '||owner||' is not disa-
bled.' value
  from dba_objects, dba_users where
owner not in ('ANONYMOUS','AURORA$JIS$UTILITY$',
'AURORA$ORB$UNAUTHENTICATED','CTXSYS','DBSNMP','DIP','DVF',
'DVSYS','EXFSYS','LBACSYS','MDDATA','MDSYS','MGMT_VIEW','ODM',
'ODM_MTR','OLAPSYS','ORDPLUGINS','ORDSYS','OSE$HTTP$ADMIN',
'OUTLN','PERFSTAT','PUBLIC','REPADMIN','RMAN',
'SI_INFORMTN_SCHEMA','SYS','SYSMAN','SYSTEM','TRACESVR', 'TSMYSWK_TEST',
'WK_TEST','WKPROXY','WKSYS','WKUSER','WMSYS','XDB')
and owner in (select distinct owner from dba_objects where object_type <>
'SYNONYM')
and owner = username and upper(account_status) not like '%LOCKED%'
```

- Type an alias, for example, DBMS application object owner accounts.
 - For the Parser, use Database Query Parser.
8. Click **Save** then click **Yes** on Configuration box.

Figure 50-10 Completed Create Configuration Extension Page**Create Configuration Extension**

Use this page to create a new Configuration Extension or to edit an existing one. A Configuration Extension is defined for a given target type and can subsequently be deployed to targets of that type.

* Name: DG0142 DBMS Privileged action aud

Description:

* Target Type: Database Instance

Sample Target: [Select Target](#)

Name:

Host:

Files & Commands: **SQL**

Database Credentials: Monitoring Database Credentials [Add Monitoring Credentials](#)

JDBC URL for Sample Target: HOST:PORT:SID

[Add](#) [Delete](#) [Manage Parsers](#) [Parser Rules](#)

SQL	Alias	Parser
select distinct 'Unauthorized user ' owner ' owns application objects in the database.' value from dba_objects where owner not i	DBMS application object ownership	Database Query Parser
select distinct 'Application object owner account ' owner ' is not disabled.' value from dba_objects, dba_users where owner not i	DBMS application object owner accour	Database Query Parser

Creating an Agent-Side Compliance Standard Rule

To create an agent-side compliance rule:

1. From the Enterprise menu, select **Compliance**, then select **Library**.
2. On the Compliance Library, click **Compliance Standard Rules**.
3. Click **Create**. On the Create Rule pop-up, choose **Agent-side Rule**.
4. Click **Continue**.
5. On the Create Rule: Agent-side Rule: Details page provide the following information (see [Figure 50-11](#)):
 - a. Name: DBMS application object ownership
 - b. Compliance Rule State: Development
 - c. Severity: Critical
 - d. Applicable To: Database Instance
 - e. Description: Application objects should be owned by accounts authorized for ownership.
 - f. Rationale: Database object ownership implies full privileges to the owned object including the privilege to assign access to the owned objects to other subjects. Unmanaged or uncontrolled ownership of objects can lead to unauthorized object grants and alterations.
 - g. Click **Next**.

Figure 50-11 Completed Compliance Standard Rule Details Page

Create Rule: Agent-side Rule: Details Back Step 1 of 4 N

* Rule: DBMS application object ownership

Compliance Rule State: Development

Severity: Critical

* Applicable To: Database Instance

Target Property Filter

Description: Application objects should be owned by accounts authorized for ownership.

Rationale: Database object ownership implies full privileges to the owned object including the privilege to assign access to the owned objects to other subjects. Unmanaged or uncontrolled ownership of objects can lead to unauthorized object grants and alterations.

Recommendation:

Reference URL:

Keywords: + Add... - Remove

Rule Keywords

No data to display.

- On the Create Rule: Agent-side Rule: Check Definition Page search for the configuration extension and alias you defined earlier. See [Figure 50-12](#).

Note: The configuration extension name *and* the alias name are concatenated together to form the name in the Configuration Extension and Name field. For this example, the complete name is: DG0142 DBMS Privileged action audit-DBMS application object ownership.

Click **Next**.

Figure 50-12 Completed Compliance Standard Rule Check Definition Page

Create Rule: Agent-side Rule: Check Definition Back Step 2 of 4 N

Configuration Extension Details

Configuration Extension-Alias Name: DG0142 DBMS Privileged action aud

Violation Details

Compliant Message:

Non-Compliant Message:

- On the Create Rule: Agent-side Rule: Test Page, search for a target, and then click **Run Test**. A pop-up displays stating that the test is running. Click **Close** on the Confirmation pop-up. See [Figure 50-13](#).

Note: You can have test results that intentionally show violations. For example, if you are testing target type equal to host and you are evaluating a host target, then you will see violation results.

Click **Next**.

Figure 50-13 Completed Compliance Standard Rule Test Page

Create Rule: Agent-side Rule: Test Back Step 3 of 4 N

Select the target against which you want to test the rule.

Target Name: Oemrep_Database Run Test

☒ **TIP** Click the icon to select a target from the list of available targets.

8. On the Create Rule: Agent-side Rule: Review, ensure the information is as you intended. If not, click **Back** and make the necessary corrections. When the information is correct, click **Finish**. See [Figure 50-14](#).

Note: The compliance standard rule is not defined until you click **Finish**.

Tips

- Once the compliance standard rule has been created, it is not automatically evaluated. Consider adding the compliance standard rule to a compliance standard.
- Assign a corrective action to the rule after the rule has been created.
 - On the Compliance Standard Rules tab, highlight the rule you just created.
 - From the **Actions** menu, select **Assign Corrective Action**.
 - From the **Assign Creative Action** popup, select an existing corrective action and click **OK**.

Figure 50-14 Completed Compliance Standard Rule Review Page

Details Check Definition Test **Review**

Create Rule: Agent-side Rule: Review

Back Step 4 of 4 Finish Cancel

Rule	DBMS application object ownership
Applicable To	Database Instance
Configuration Extension	DGO142 DBMS Privileged action audit
Name	
Alias Name	DBMS application object ownership
Severity	Critical
Compliance Rule State	Development
Description	Application objects should be owned by accounts authorized for ownership.
Rationale	Database object ownership implies full privileges to the owned object including the privilege to assign access to the owned objects to other subjects. Unmanaged or uncontrolled ownership of objects can lead to unauthorized object grants and alterations.
Reference URL	Keyword
Compliant Message	Compliant Message
Non-Compliant Message	Non-Compliant Message

9. Repeat these steps for the second rule.

Note: The compliance standard rule is not defined until you click **Finish**.

Creating a Manual Rule

The purpose of creating this manual rule is to keep track of the checks that cannot be automated: ensuring that test plans and procedures have been followed prior to production.

To create a manual rule:

1. On the Compliance Library, click **Compliance Standard Rules**.
2. Click **Create**. On the Create Rule pop-up, choose **Manual Rule**.
3. Click **Continue**.
4. On the Create Manual Rule page, provide the following information (see [Figure 50-15](#)).
 - a. Name: DBMS testing plans and procedures
 - b. Compliance Rule State: Production
 - c. Severity: Warning
 - d. Applicable To: Database Instance

- e. Description: Plans and procedures for testing DBMS installations, upgrades, and patches should be defined and followed prior to production implementation.
- f. Rationale: Updates and patches to existing software have the intention of improving the security or enhancing or adding features to the product. However, it is unfortunately common that updates or patches can render production systems inoperable or even introduce serious vulnerabilities. Some updates also set security configurations back to unacceptable settings that do not meet security requirements. For these reasons, it is a good practice to test updates and patches offline before introducing them in a production environment.
- g. Recommendation: Develop, document and implement procedures for testing DBMS installations, upgrades and patches prior to deployment on production systems.
- h. Compliant Message: Plans and procedures for testing DBMS installations, upgrades and patches are defined and followed prior to production implementation.
- i. Non-Compliant Message: Plans and procedures for testing DBMS installations, upgrades and patches are not defined or followed prior to production implementation.
- j. Reference URL: <http://iase.disa.mil/stigs/index.html>
- k. Rule Keywords: Security
- l. Click **Finish**.

Figure 50-15 Completed Manual Rule Page

Create Manual Rule

* Rule: DBMS testing plans and procedures

Compliance Rule State: Production

Severity: Warning

* Applicable To: Database Instance

> Target Property Filter

Description: Plans and procedures for testing DBMS installations, upgrades, and patches should be defined and followed prior to production implementation.

Rationale: Updates and patches to existing software have the intention of improving the security or enhancing or adding features to the product. However, it is unfortunately common that updates or patches can render production systems inoperable or even introduce serious vulnerabilities. Some updates also set security configurations back to unacceptable settings that do not meet security requirements. For these reasons, it is a good practice to test updates and patches offline before introducing them in a production environment.

Recommendation: Develop, document and implement procedures for testing DBMS installations, upgrades and patches prior to deployment on production systems.

Compliant Message: Plans and procedures for testing DBMS installations, upgrades and patches are defined and followed prior to production implementation.

Non-Compliant Message: Plans and procedures for testing DBMS installations, upgrades and patches are not defined or followed prior to production implementation.

Reference URL: <http://iase.disa.mil/stigs/index.html>

Keywords

Rule Keywords

Security

Finish Cancel

Creating a Compliance Standard

To create a compliance standard, perform the following steps:

1. From the Enterprise menu, select **Compliance**, then select **Library**.
2. On the Compliance Library page, click **Compliance Standards**.
3. Click **Create**. On the Create Compliance Standard pop-up, provide the following (see Figure 50-16):

- Name: CS1 - DB Check
- Applicable To: Select Database Instance
- Author: SYSMAN
- Standard Type: Agent-side
- Click **Continue**

Figure 50-16 Completed Create Compliance Standard Pop-Up

Create Compliance Standard

Enter the details to create Compliance Standard.

Name	CS1 - DB Check	Applicable To	Database Instance
Author	SYSMAN	Standard Type	Agent-side

[Continue](#) [Cancel](#)

4. On the Compliance Standard: CS1 - DB Check page, right-click the standard in the navigation tree. Select **Add Rules**. On the Include Rule Reference, select DBMS application object ownership, DBMS application owner accounts, and DBMS testing plans and procedures. See [Figure 50-17](#). Click **OK**.

Figure 50-17 Compliance Standard Rules

Compliance Standard: CS1 - DB Check

Select a Compliance Standard node to see its details. Right click the node (or select the node and press Ctrl+Alt+M) to modify the hierarchy.

Properties

DBMS application object owner accounts (Compliance Standard Rule) [Save](#) [Cancel](#)

Name: DBMS application object owner accounts

Description:

Compliance: Development

Rule State:

Importance: Normal

Rule Check Definition

Configuration Extension: DG0142 DBMS Privileged action audit

Name:

Alias Name: DBMS application object owner accounts

Query: select distinct 'Application object owner account' || owner || ' is not disabled.' value from dba_objects, dba_users where owner not in ('ANONYMOUS','AURORA\$IIS\$UTILITY\$', 'AURORA\$ORB\$UNAUTHENTICATED', 'CTXSYS', 'DBSNMP', 'DIP', 'DIPF', 'DVSYS', 'ENFSYS', 'LBACSYS', 'MODDATA', 'MDSYS', 'MGMT_VIEW', 'ODM', 'ODM_MTR', 'OLAPSYS', 'ORDPLUGINS', 'ORDSYS', 'OSE\$HTTP\$ADMIN', 'OUTLN', 'PERFSTAT', 'PUBLIC', 'REPADMIN', 'RMAN', 'SL_INFORMTN_SCHEMA', 'SYS', 'SYSMAN', 'SYSTEM', 'TRACESVR', 'TSMYSYS', 'WK_TEST', 'WKPROXY', 'WKUSER', 'WMSYS', 'XDB') and owner in (select distinct owner from dba_objects where object_type <> 'SYNONYM') and owner = username and upper(account_status) not like '%LOCKED%'

5. Click **Save**.

Associating the Compliance Standard to a Target

To associate the compliance standard to a target, perform the following steps:

1. From the Enterprise menu, select **Compliance**, then select **Library**.
2. On the Compliance Library page, click **Compliance Standards**.

Figure 50-18 Compliance Standards Library Page

Compliance Library Page Refreshed Apr 23, 2014 9:21:24 AM PDT

Compliance Frameworks **Compliance Standards** Compliance Standard Rules Real-time Monitoring Facets

> Search

To perform an operation on a standard, highlight the row and select an operation. To delete multiple standards, select multiple rows and click Delete.

Actions View Create... Create Like... Show Details Edit... Delete Associate Targets... Override Target Type Settings...

Compliance Standard	Description	Compliance Standard State	Applicable To	Keywords	Author	Standard Type	Association Co
CS1 - DB Check		Development	Database Instance		SYSMAN	Agent-side	
All WLS v10 rules	All WLS v10 rules.	Production	Oracle WebLogic Domain	Security, Configuration	ORACLE	WebLogic Server S...	
All WLS v11 rules	All WLS v11 rules.	Production	Oracle WebLogic Domain	Security, Configuration	ORACLE	WebLogic Server S...	

- Highlight the newly created standard (CS1 - DB Check) and click the **Associate Targets** button.
- On the Target Association for Compliance Standard: CS1 - DB Check page click **Add**.
- Choose one or more targets, for example, Oemrep_Database. See [Figure 50-19](#).

Figure 50-19 Completed Target Association Page

Compliance Standard Target Association

Target Association for Compliance Standard: CS1 - DB Check OK C

This table lists the targets that are associated with the compliance standard selected in the Compliance Standard Library.

View Add... Remove Edit... Enable Disable Force Transfer

Target Name	Target Type	Evaluation Status	Customized	Transfer Status	Association Time Stamp
Oemrep_Database	Database Instance	Enabled			

- Click **Select**. Click **OK**.
- Click **Yes** to Save the Association.

Suppressing Violations

The purpose of this example is to suppress violations. We will suppress the violation that arose due to the manual rule defined in [Creating Compliance Standard Agent-side and Manual Rules](#).

Follow these steps:

- From the **Compliance** menu, select **Results**.
- In the Evaluation Results tab, locate the compliance standard named CS1 - DB Check. Notice that there is a violation against the standard.

Figure 50-20 CS1 - DB Check Compliance Standard in Evaluation Results Tab

Compliance Results Page Refreshed Apr 24, 2014 6:32:34 AM PDT

Compliance Frameworks **Compliance Standards** Target Compliance

Evaluation Results Errors

> Search

View Show Details Manage Violations

Compliance Standards	Applicable To	Compliance Standard State	Target Evaluations	Violations	Average Score (%)
CS1 - DB Check	Database Instance	Development	0 1 0	0 1 0	66

3. Select the compliance standard and click the **Manage Violations** tab.
4. On the Manage Violations page, ensure the **Unsuppressed Violations** tab is selected.
5. Select **DBMS testing plans and procedures**. See [Figure 50-21](#).

Figure 50-21 Manage Violations Page - Unsuppressed Violations

Manage Violations

Compliance Standard: CS1 - DB Check Return

Unsuppressed Violations | Suppressed Violations | Manual Rule Violations

This table lists all unsuppressed violations. Select rows in the table and click on Suppress Violations button in order to suppress the selected violations.

Search

View bd Suppress Violations

Rule	Target Name	Applicable To	Keywords	Severity	Recommendation
DBMS testing plans...	Oemrep_Database	Database Instance	Security	Warning	Develop, document and implement procedures for testing DBMS installations, upgrades and patches prior to deployment on production systems.

Rows Selected 1

Plans and procedures for testing DBMS installations, upgrades and patches are not defined or followed prior to production implementation.

Details Not Completed

6. To suppress the violation, click the **Suppress Violations** tab.
7. On the Violation Suppressed Confirmation popup, select **Suppress Violations Indefinitely**.
8. Once the violation is suppressed, it no longer appears on the Evaluation Results page. See [Figure 50-22](#).

Figure 50-22 Evaluation Results Page After Violation Is Suppressed

Compliance Results Page Refreshed Apr 23, 2014 11:23:51 AM PDT

Compliance Frameworks | **Compliance Standards** | Target Compliance

Evaluation Results | Errors

Search

View bd Show Details Manage Violations

Compliance Standards	Applicable To	Compliance Standard State	Target Evaluations	Violations	Average Score (%)
CS1 - DB Check	Database Instance	Development	0 0 1	0 0 0	100

9. To unsuppress the violation, use the **Suppressed Violations** tab as shown in [Figure 50-23](#). Select the rows and then click **Unsuppress Violations**.

Unsuppressing a violation causes the compliance score to be recomputed accounting for the violations that were unsuppressed.

Figure 50-23 Manage Violations Page Showing the Suppressed Violations Tab

Manage Violations

Compliance Standard: CS1 - DB Check Return

Unsuppressed Violations | **Suppressed Violations** | Manual Rule Violations

This table lists all suppressed violations. Select rows in the table and click on Unsuppress Violations button in order to unsuppress the selected violations.

Search

View bd Unsuppress Violations

Rule	Target Name	Applicable To	Keywords	Severity	Recommendation
DBMS testing plans...	Oemrep_Database	Database Instance	Security	Warning	Develop, document and implement procedures for testing DBMS installations, upgrades and patches prior to deployment on production systems.

Clearing Violations

Clearing of manual rule violations causes the violations to be cleared, and the compliance score to go up for the corresponding compliance standard or target. To clear violations, perform the following steps:

1. From the **Compliance** menu, select **Results**. Select the CS1 - DB Check compliance standard.
2. Click **Manage Violations**.
3. On the Manage Violations page, highlight the **DBMS testing plans and procedures rule**.
4. Click the **Manage Rule Violations** tab.
5. On the Manage Violations page, highlight the rule and click the **Manual Rule Violations** tab.

Figure 50-24 Clearing Manual Rule Violations

Manage Violations

Compliance Standard: CS1 - DB Check Return

Unsuppressed Violations Suppressed Violations **Manual Rule Violations**

This table lists all manual rule violations. Select rows in the table and click on Clear Violations button in order to clear the selected violations.

Search

View

Rule	Target Name	Applicable To	Keywords	Severity	Recommendation
DBMS testing plans...	Oemrep_Database	Database Instance	Security	Warning	Develop, document and implement procedures for testing DBMS installations, upgrades and patches prior to deployment ...

Rows Selected 1

Plans and procedures for testing DBMS installations, upgrades and patches are not defined or followed prior to production implementation.

Details Not Completed

6. Select the rows and then click **Clear Violations**. On the Clear Violations Confirmation pop-up, select either **Clear Violations Indefinitely** or **Clear Violations Until** and specify a date. For completeness, provide a reason for clearing the violation.

Managing Enterprise Data Governance

This chapter introduces Enterprise Data Governance and describes how to use the feature to protect sensitive data. The chapter includes the following sections:

- [Overview of Enterprise Data Governance](#)
- [Using Enterprise Data Governance](#)

Overview of Enterprise Data Governance

This section provides a brief overview of Enterprise Data Governance. The section covers the following topics:

- [About Enterprise Data Governance](#)
- [What Are Protection Policies?](#)
- [What Are Application Signatures?](#)

About Enterprise Data Governance

Enterprise Data Governance offers a comprehensive solution for identifying, securing, managing, and tracking sensitive data in the data center. The solution involves a two-pronged approach to provide this protection:

- Perform user-initiated and automatic discovery on a regular basis of databases that potentially contain sensitive data. This is metadata discovery, also referred to as a shallow scan, so-called because it looks only at metadata involving schema, table, and column name patterns.
- Perform user-initiated discovery of sensitive data in databases identified by the metadata discovery. This is data discovery, also referred to as a deep scan, so-called because it drills down in the actual data, looking for matches to user-supplied sensitive types and object-level protection details.

Enterprise Data Governance forms the first steps in the recommended workflow to mask sensitive data:

1. Discover databases that potentially contain sensitive data.
2. Aided by (but not limited to) the results of discovering database candidates, drill down to the data within the tables and columns of databases to further identify sensitive data.
3. Armed with the results of this discovery, flag columns as sensitive and identify them within the context of an Application Data Model (ADM).
4. Select these columns within an ADM and apply masking formats to protect the data in the testing environment.

What Are Protection Policies?

A Protection Policy defines a security mechanism for protecting a sensitive data object. It controls the way a sensitive data object is protected. After a policy is created for a sensitive object, it serves as a template that can be applied to all the sensitive data objects of a similar type and structure. This ensures that a sensitive data object is protected consistently no matter where it is present in the database cloud.

A Protection Policy maps to a security feature available in Oracle Database. Metadata discovery identifies databases that contain objects that are protected via one or more of the following database security features:

- **Transparent Data Encryption (TDE)**—A database feature that automatically encrypts data when it is written to the database and automatically decrypts data when accessed.
- **Data Redaction**—A database feature that protects data by presenting a masked version of the data to nonprivileged users. The masked version of the data preserves the format and referential integrity of the data, so any application that uses the data continues to work as expected.
- **Virtual Private Database (VPD)**—A database feature that enforces data access at the row and column level, using security conditions to protect the data.
- **Oracle Label Security (OLS)**—A database feature that provides data classification and control access using security labels.

Metadata discovery checks for each security feature listed. The scan does not, however, collect protection policy details, nor does it necessarily scan for all the policies. Any protection policy found is sufficient to flag the database as potentially sensitive. This strategy keeps the scan fast and lightweight.

What Are Application Signatures?

An application signature is a set of database objects such as schemas, tables, and views that uniquely identifies a specific application. A database that contains these objects is assumed to contain the application and is noted as a sensitive database candidate. Oracle supplies signatures for the following applications:

- Oracle E-Business Suite
- Oracle Fusion Applications
- Oracle PeopleSoft Enterprise

You can also create custom application signatures, see [Creating Custom Application Signatures](#).

Using Enterprise Data Governance

This section covers the following topics:

- [The Enterprise Data Governance Dashboard](#)
- [Working with Sensitive Database Discovery Results](#)
- [Working with Metadata Discovery Jobs](#)
- [Working with Data Discovery Jobs](#)

- [Creating Custom Application Signatures](#)

The Enterprise Data Governance Dashboard

Enterprise Data Governance provides the means to identify databases within the enterprise that potentially contain sensitive data, and then to evaluate the data within these candidates to determine if sensitive data exists.

The Enterprise Data Governance dashboard summarizes discovery activity and provides links to:

- Review the results of sensitive discovery jobs (see [Working with Sensitive Database Discovery Results](#)).
- Manage and review metadata discovery jobs (see [Working with Metadata Discovery Jobs](#)).
- Manage and review data discovery jobs (see [Working with Data Discovery Jobs](#)).
- Create application signatures (see [Creating Custom Application Signatures](#)).

You can also manage the Application Data Model (ADM) environment and sensitive column types from the dashboard. See Chapter 2, "Application Data Modeling," in the *Oracle Data Masking and Subsetting Guide* for information on these activities.

To navigate to the dashboard within the Cloud Control console, select **Databases** on the **Targets** menu, then select **Enterprise Data Governance** on the **Security** menu. Whenever you navigate away from the dashboard, use the **Enterprise Data Governance** bread crumb at the top to return.

Working with Sensitive Database Discovery Results

On the Sensitive Database Discovery Summary page, you can perform the following tasks:

- Review databases discovered to have sensitive data or considered to be sensitive data candidates.
- Create a metadata discovery job (see [Creating a Metadata Discovery Job](#)).
- Create a data discovery job (see [Creating a Data Discovery Job](#)).
- Click a number in a metadata column to see a pop-up list of items found. For example, click the number in the Data Protections column to see which data protections are in play for the database candidate.
- Click the database name itself to open the database instance home page.

Working with Metadata Discovery Jobs

On the Metadata Discovery Jobs page, you can perform the following tasks:

- Create a metadata discovery job (see [Creating a Metadata Discovery Job](#)).
- Manage automatic metadata discovery.
- Manage job results.

Since a metadata discovery job looks only at schema, table, and column name patterns but not at the data itself, there are no database credentials required to execute the job.

Creating a Metadata Discovery Job

Run a metadata discovery job to scan database metadata looking for candidates that potentially contain sensitive data.

Creating a metadata discovery job involves the following steps:

1. Click **Create Metadata Discovery Job**.
2. Set the criteria for sensitive column types, application signatures, and data protections.

For sensitive column type, select a row and click **View Search Criteria** to see applicable criteria such as pattern matching, regex formatting and Boolean condition.

When done, click **Next** to continue.

3. Select the targets on which you want to perform metadata discovery. First, select the target type, then click **Add** to select the targets within a given type. Note that you can include searches from the configuration search library as part of your target search criteria.

You cannot select targets of a different type. If you select targets of one type and then select targets of a different type, targets of the first selected type are deselected.

When done, click **Select** to close the selection dialog, then click **Next** to continue.

4. Schedule the job. Provide a meaningful name and description. Set other parameters as appropriate. Note that metadata discovery is a job you would typically want to repeat on a rotating schedule to be vigilant in monitoring your databases for sensitive data.

When done, click **Submit**.

5. A confirmation message appears at the top of the page. Click the link to view job details in the Jobs system. Refresh the Metadata Discovery Jobs page to see the completed job.

Managing Automatic Metadata Discovery

Automatic metadata discovery happens independent of user-initiated metadata discovery and ties directly to target discovery. By default, whenever a database is discovered as part of target discovery, the metadata discovery job runs on that database. You can disable this feature by choosing **Disable Metadata Discovery During Target Discovery** from the **Automatic Metadata Discovery** drop-down menu. You may want to disable the feature if you want more control over when the metadata discovery job is run and on which databases. When you disable the feature, the menu selection toggles to **Enable metadata discovery during target discovery** so you have the option of resuming automatic metadata discovery.

You can also choose to retain the feature but with a different set of criteria. Out-of-box criteria for automatic metadata discovery uses Oracle-defined sensitive column types, data protection policies, and application signatures, but you can change the default settings and add user-defined entities as well. Select **Edit Automatic Metadata Discovery Parameters** from the **Automatic Metadata Discovery** drop-down menu to edit the criteria.

Managing Metadata Discovery Results

The results of a metadata discovery job help you ascertain which databases actually contain sensitive data and the nature of the sensitivity.

Work with metadata discovery job results by doing the following:

1. Select a job in the top table to see the discovery results at the bottom.
2. Use the **Show** drop-down list to filter the display based on all databases evaluated or only those with or without sensitive data.
3. Click **View Discovery Results Detail** to see matching metadata based on specified criteria.
4. Click a number in a metadata column to see a pop-up list of items found. For example, click the number in the Data Protections column to see which data protections are in play for the database candidate.
5. Click the database name itself to open the database instance home page.

Working with Data Discovery Jobs

On the Data Discovery Jobs page, you can perform the following tasks:

- Create a data discovery job
- Manage job results

Creating a Data Discovery Job

Run a data discovery job to search for sensitive data within a database candidate identified by the metadata discovery job.

Creating a data discovery job involves the following steps:

1. Click **Create Data Discovery Job**.
2. Click the search icon to select the database candidate on which you want to perform data discovery. Note that you can include searches from the configuration search library as part of your target search criteria.

Click **Select** to close the selection dialog.

3. Set the criteria for sensitive column types, application signatures, and data protections.

For sensitive column type, select a given column row and click **View Search Criteria** to see applicable criteria such as pattern matching, regex formatting and Boolean condition. Set the number of rows you feel constitutes an adequate sample size. Indicate whether to scan empty tables.

The data discovery job ignores empty tables on the basis that data is what makes a column sensitive. You may, however, want to include empty tables in the discovery search based on other factors such as column name and comment patterns. While an empty table is defined as a table without data values, the metadata discovery job might report some nonempty tables as empty, if the statistics collection job has yet to run.

Click **Next**.

4. Specify schema and table parameters (those to include or exclude). Use pattern matching to scope the searches. Alternatively, you can opt to include all of either or both entities.
Click **Next**.
5. Schedule the job. Specify a meaningful name and description. Provide credentials to access the database. Set the job schedule.
Click **Submit**.
6. A confirmation message appears at the top of the page. Click the link to view job details in the Jobs system. Refresh the Data Discovery Jobs page to see the completed job.

Managing Data Discovery Results

Use the results of data discovery to identify sensitive columns and associate the database with an Application Data Model.

Work with data discovery job results by doing the following:

1. Click the database name link in the job row to open the database instance home page; click the job status link to open the job summary page in the Jobs system.
2. Optionally associate a database with either a new or existing ADM. Select a data discovery job row, then click **Assign Application Data Model** and choose the appropriate option.
3. Select a job in the top table to see the discovery results at the bottom. Review job results by clicking the job criteria tabs. Expand tab contents as necessary to drill down to the details.
4. Click the **Sensitive Data Columns** tab to see the origin and nature of the data in the sensitive columns. As noted, if there is an ADM assigned, you can interactively set the sensitivity status by selecting a row and choosing a status from the **Set Sensitive Status** drop-down menu.

Use the information in the table to inform your decision to declare a column sensitive. For example, the sample data and columns matching the criteria both in name and as a percentage of data are strong indicators of the column's sensitivity.

If there is no ADM assigned to the data discovery job, sensitivity status is disabled, and the relevant schema is displayed in place of an application.

5. Click the **Application Signatures** tab to see database objects that uniquely identify the application.
6. Click the **Objects with Data Protection Policies** tab to see the specific objects the job discovered that are protected by supported protection policies.

Set sensitive column status on the discovered objects:

- a. Click **Select Sensitive Columns**.
- b. Provide credentials to log in to the database discovered by the job.
- c. Click the **List Columns** button to display all the columns in the table covered by the protection policy.
- d. Set status to sensitive and select an associated sensitive column type for those columns you consider sensitive within the application.
- e. Click **OK** when done to confirm your selections.

The selected columns are identified as sensitive within the assigned ADM.

If there is no ADM assigned to the data discovery job, the sensitive status feature is disabled, and the relevant schema is displayed in place of an application.

Creating Custom Application Signatures

Customize application signatures to facilitate sensitive data discovery within your business enterprise.

Creating a custom application signature involves the following steps:

1. Open the Application Signature link from the Enterprise Data Governance dashboard.

Click **Create**. The editor page opens.

2. Specify a name and optional description.
3. Click **Add** and select from the available objects to include in the signature. The name provided for any of these object types can be specified explicitly or with a pattern (for example, HR%).

- Schema—schema name is required
- Table—schema name is optional; table name is required
- View—schema name is optional; view name is required

Click **OK**. The object appears in the table.

4. Repeat Step 3 to include additional objects in the signature. Remember that all signature objects must be found in the database for there to be a match.
5. When done, click **OK** to complete the signature definition.

The editor window closes and the signature appears in the table on the Application Signature page. The signature can now be used as search criteria for metadata discovery and data discovery jobs.

Managing Database Schema Changes

This chapter introduces database change management solution in the following sections:

- [Overview of Change Management for Databases](#)
- [Using Schema Baselines](#)
- [Using Schema Comparisons](#)
- [Using Schema Synchronizations](#)
- [Using Change Plans](#)
- [Using Database Data Comparison](#)

Overview of Change Management for Databases

To manage the lifecycle of enterprise applications, an organization will need to maintain multiple copies of an application database for various purposes such as development, staging, production, and testing. Each of these databases must adhere to different processes. For example, for production databases, it is essential to ensure adherence to proper production control procedures. It is vital that administrators have the tools to detect unauthorized changes, such as an index being dropped without the requisite change approvals. In such cases, monitoring changes to production databases day over day or week over week becomes vital.

Database compliance, that is, ensuring that all databases meet the gold standard configuration, is another important aspect of life cycle management. Compliance with organizational standards or best practices ensures database efficiency, maintenance, and ease of operation.

On development databases, developers make changes that the database administrator needs to consolidate and propagate to staging or test databases. The goal is to identify the changes made to development and then make the same changes to staging or test databases taking into account any other changes already in production database.

Typically, most applications will get upgraded over time. Also, most applications are customized by the business user to suit their needs. Application customizations are usually dependent on database objects or PL/SQL modules supplied by the application vendor. The application vendor supplies the upgrade scripts and the customer has very little transparency about the impact of the upgrade procedure on their customizations. When customers test upgrade databases, they can capture a baseline of the application schema before and after the upgrade. A comparison of the before and after baselines will tell the user what modules were changed by the application. This gives them a better idea about how their customizations will be impacted as a result of upgrading their application.

The following are core capabilities of Change Management that allow developers and database administrators to manage changes in database environments:

- Schema Baseline—A point in time of the definition of the database and its associated database objects.
- Schema Comparison—A complete list of differences between a baseline or a database and another baseline or a database.
- Schema Synchronization—The process of promoting changes from a database definition capture in a baseline or from a database to a target database.
- Schema Change Plans—A means of deploying specific changes from a development environment to one or more target databases.
- Data Comparison—A list of differences in row data between two databases.

For database versions 9.x and above, the user logged into the database target through Cloud Control must have `SELECT ANY DICTIONARY` privilege and `SELECT_CATALOG_ROLE` role for capturing or comparing databases. To perform schema synchronization, the user logging in to the destination database must have the `SYSDBA` privilege. To create or delete change plans, Cloud Control users need the Manage Change Plans resource privilege, `EM_ALL_OPERATOR` privilege, `VIEW` and `CONNECT` privilege for the targets, and Create resource privilege for the job system and Create new Named Credentials resource privilege. Users can also be granted View and Edit privileges on specific change plans.

When submitting a data comparison job, the user whose credentials are specified for the reference and candidate databases must have `SELECT` privilege on reference and candidate objects respectively. Additionally, the users needs these privileges: `SELECT ANY DICTIONARY`, `SELECT_CATALOG_ROLE`, and `CREATE VIEW`. When comparing objects with LOB type columns included, the users need to be granted `EXECUTE` privilege on `SYS.DBMS_CRYPTO` package, since cryptographic hash value of the columns will be compared instead of actual column values. And in case you specify the comparison to be performed as of a time stamp or system change number (SCN), the users must also be granted `FLASHBACK` privilege directly on the reference and candidate objects in their respective databases.

Further, the user whose credentials are specified as reference database credentials must be a DBA with `EXECUTE` privilege on `DBMS_COMPARISON` program and in case the reference database is not the same as candidate database, the `CREATE DATABASE LINK` privilege as well.

Database link, comparison definitions, and views may be created in the reference database by the data comparison job. Views may be created in the candidate database. These objects created during the comparison processing will be dropped when the comparison is deleted, unless you specify the option to skip dropping them at the time of deletion.

Data comparison cannot be performed connecting to a remote candidate database as user `SYS` since `SYS` credentials cannot be used over database links.

Using Schema Baselines

A schema baseline contains a set of database definitions captured at a certain point in time. Baselines are stored in the Cloud Control repository as XML documents.

Each baseline must be assigned a unique name. A good practice to name baselines is to match it on the scope of the database objects being captured in the baseline, for example, Financial 11.5.10 or HR Benefits or PO Check Print. A baseline can have a series of versions that have been captured at different points in time. Creating multiple

versions of a baseline allows you to track changes to the definitions of a set of database objects over time. You can compare two versions of the same baseline to see the changes that have occurred between them.

When creating a baseline, you also create a corresponding baseline scope specification, which describes the names and the types of database objects and schemas from which they should be captured. When you have created the baseline definition, you can then capture the first version of the baseline by submitting an Cloud Control job. At a later time, or at regular intervals, you can capture additional versions of the same baseline. Each baseline version records the metadata definitions as they exist at the time the version is captured.

Change management schema baselines are retained in the system until you delete them. When you delete a baseline, it is deleted from the system permanently. Delete operation cannot be undone. However, if a baseline may be needed in future, you can first export the baseline to a dump file (created on the repository database server host) and then delete the baseline. Baseline can then be imported back from the file at a later time if needed.

Overview of Scope Specification

A scope specification identifies the database objects to be captured in a baseline. (Scope specifications also identify objects to process in schema comparisons and synchronizations.) After you have specified the scope of a baseline, you cannot change the scope specification. This restriction ensures that all versions of the baseline are captured using the same set of rules, which means that differences between versions result from changes in the database, not scope specification changes. To capture schema objects using a different scope specification, you must create a new baseline.

Baseline scope specifications take three forms.

- You can specify schemas and object types to capture. For example, you can capture all Tables, Indexes and Views in schemas APPL1 and APPL2. This form of scope specification is appropriate when there is a well-defined set of schemas that contain your application objects. In addition to schema objects, you can also capture non-schema objects (such as Users, Roles and Tablespace) and privilege grants to Users and Roles.
- You can specify schemas to exclude, and object types. This form of scope specification captures objects that are contained in all schemas other than those you specify. For example, you can capture all object types in schemas other than SYSTEM and SYS. This form of scope specification is appropriate when you want to capture all database objects, with the exception of objects contained in Oracle-provided schemas. As with the first form of scope specification, you can also capture non-schema objects and privilege grants.
- Finally, you can capture individual schema objects by specifying the type, schema and name of each object. This form of scope specification is appropriate when you want to capture a few specific objects, rather than all objects contained within one or more schemas. While capturing individual schema objects, you can also capture non-schema objects and privilege grants.

If you include a non-schema object type, such as User or Role, in a scope specification, all objects of that type are captured. There is no way to capture individual non-schema objects.

About Capturing a Schema Baseline Version

As the final step of defining a baseline, you specify when to capture the first version of the baseline. You can capture the first version immediately, or at a later time (for example, when the database is not being used in active development work). You can also indicate that additional versions of the baseline should be captured at regular intervals without further intervention on your part.

You can also capture a new baseline version at any time by selecting the baseline and specifying "Recapture Now."

Baselines processed after the initial version generally complete substantially faster than the initial version. Only those objects that have changed are captured in the new version. This also means that the storage required for additional baseline versions is only slightly larger than the storage used by the initial version, assuming a small percentage of objects has changed.

About Working With A Schema Baseline Version

Within a single schema baseline version, you can examine individual object attributes, generate DDL for individual objects, and generate DDL for all the objects in the baseline version. You cannot modify object definitions captured in baseline versions, since they are intended to represent the state of objects at a particular point in time.

- Viewing a baseline object displays the object's attributes graphically.
- Selecting a baseline object and specifying "Generate DDL" displays the DDL used to create the object.
- Selecting a baseline version and specifying "Generate DDL" generates the DDL for all objects in the baseline version. While an effort is made to create objects in the correct order (for example, creating tables before indexes), the resulting DDL cannot necessarily be executed on a database to create the objects in the baseline version. For example, if you capture all the schema objects contained in schema APPL1, then try to execute the baseline version DDL on a database that does not contain User APPL1, the generated DDL will fail to execute.

Baseline versions are also used with other Database Lifecycle Management Pack applications, such as Compare and Synchronize. You can compare a baseline version to a database (or to another baseline version). You can also use a baseline version as the source of object definitions in Synchronize, allowing you to re-create the definitions in another database.

About Working With Multiple Schema Baseline Versions

When a baseline contains more than one version, you can examine the differences between the versions.

- To see what has changed between a version and the version that precedes it, select the version and specify "View Changes Since Previous Version." The display shows which objects have changed since the previous version, which have been added or removed, and which are unchanged. Selecting an object that has changed displays the differences between the object in the two versions.
- To see how an individual object has changed over all the versions of the baseline, select the object and specify "View Version History." The display identifies the ver-

sions in which the object was initially captured, modified, or dropped. From this display, you can compare the definitions of the object in any two baseline versions.

Exporting and Importing Schema Baselines

You can use the export/import baseline functionality for the following:

- Transferring baselines between two Cloud Control sites with different repositories.
- Offline storage of baselines. Baselines can be exported to files, deleted, and then imported back from files.

You can select a schema baseline or a version and then export it to a file. The system uses Data Pump for export and import. The dump files and log files are located in the Cloud Control repository database server host. They can be located in directories set up on NFS file systems, including file systems on NAS devices that are supported by Oracle.

Creating Directory Objects for Export and Import

To export a schema baseline version from the repository to an export file or import schema baselines from an import file on the repository database server, select the directory object in the repository database for the export or import and specify a name for the export or import file.

To create a new directory object for export or import, do the following:

1. Log in to the repository database as a user with CREATE ANY DIRECTORY privilege or the DBA role.
2. Create a directory object as the alias for a directory on the repository database server's file system where the baselines are to be exported or where the import dump file is stored.
3. Grant READ and WRITE privileges on the directory object to SYSMAN.

The newly created directory will be available for selection by Cloud Control administrators for export and import of schema baselines. Data pump log files from the export and import operations are also written to the same directory.

During import, new values can be set for name, owner, and source database. Super administrators can set another administrator as the owner at the time of import.

The export operation does not export job information associated with a baseline. On import, the job status will hence be unknown.

For non-super administrators, the following applies:

- Non-super administrators can export their own baselines. They can also export a version of baseline owned by another administrator, provided they have the privilege to view the version and see the list of schema objects in that version.
- At the time of import, a non-super administrator must become the owner of the baseline being imported. A non-super administrator cannot set another administrator as the owner. If the baseline in the import dump file was owned by another administrator, its new owner is set to the logged-in non-super administrator at the time of import.

- View privileges granted on the baseline to non-super administrators are lost during import and cannot be re-granted after the import, since there is no associated job information.

Using Schema Comparisons

A schema comparison identifies differences in database object definitions between a baseline or a database and a baseline or a database, or two schemas within a single database/baseline.

A comparison specification is defined by left and right sources, scope, and owner. The scope specification describes the names and types of database object definitions to be included in the comparison and the schemas that contain these object definitions.

Comparisons identify differences in any attribute value between objects of any type. Use comparisons to create multiple versions of a comparison. Each version has a unique version number and a comparison date. Use these versions to associate comparisons of database/schemas made over time.

Comparisons show differences between definitions in the original baseline for your application and those in your current database. After creating a new comparison version, it identifies the differences between the original definitions at the start of the development cycle, and those same definitions at the current time.

Use another comparison specification to compare definitions from your most recent baseline with those in your previous baseline. With each newly created version of this comparison using the comparison specification, that comparison version identifies the differences in the definitions since the previous baseline.

Defining Schema Comparisons

A schema comparison definition consists of left and right sources for metadata definitions, the scope specification, an optional schema map, and comparison options. Once created, a schema comparison definition cannot be modified. This ensures that each version of the schema comparison reflects changes in the databases being compared and not in the comparison's definition.

Schema Comparison Sources

Each comparison has a left source and a right source. The comparison matches object definitions from the left and right sources. A source can be a database or a baseline version.

- When the source is a database, object definitions are taken directly from the database at the time the comparison version is created.
- When the source is a baseline, object definitions are taken from the specified baseline version. Using a baseline version allows you to compare a database (or another baseline version) to the database as it existed at a previous point in time. For example, a baseline version might represent a stable point in the application development cycle, or a previous release of the application.

For baseline sources, there are various ways to specify the version to be used.

- If you want a specific baseline version to be used in all versions of the comparison, specify the baseline version number. This is appropriate for comparing a well-defined previous state of the database, such as a release, to its current state.

- You can also request that the latest or next-to-latest version be used in the comparison. If you specify "Latest," you can also request that the baseline version be captured before the comparison takes place. This option allows you to capture a baseline and compare it to the other source in a single operation. For example, every night, you can capture a baseline version of the current state of a development database and compare it to the previous night's baseline, or to a fixed baseline representing a stable point in development.

Scope Specification

The scope specification for a schema comparison identifies the objects to compare in the left and right sources. Creating a comparison scope specification is the same as creating a baseline scope specification, described in the "Schema Baselines" section. As with baselines, you can specify object types and schemas to compare, or individual objects to compare.

Schema Map

Normally, schema objects in one source are compared to objects in the same schema in the other source. For example, table APPL1.T1 in the left source is compared to APPL1.T1 in the right source.

However, there may be cases where you want to compare objects in one schema to corresponding objects in a different schema. For example, assume that there are two schemas, DEV1 and DEV2, which contain the same set of objects. Different application developers work in DEV1 and DEV2. You can use the optional schema map feature to allow you to compare objects in DEV1 to objects with the same type and name in DEV2.

To add entries to the schema map, expand the "Mapped Objects" section of the comparison "Objects" page. You can create one or more pairs of mapped schemas. Each pair designates a left-side schema and a corresponding right-side schema.

When using a schema map, you can compare objects within a single database or baseline version. In the example above, DEV1 and DEV2 can be in the same database. You specify that database as both the left and right source, and supply the schema map to compare objects in DEV1 to those in DEV2.

Comparison Options

You can select several options to determine how objects are compared. These options allow you to disregard differences that are not significant. The options include the following:

- **Ignore Tablespace and Ignore Physical Attributes**— These two options allow you to compare stored objects without regard to the tablespaces in which they are stored or the settings of their storage-related attributes. This is useful when you are comparing objects in databases having different size and storage configurations, and you are not interested in differences related to the storage of the objects.
- **Match Constraints By Definition or By Name**— If you are more interested in the definitions of table constraints – for example, the columns involved in primary or unique constraints – choose "Match Constraints By Definition." This causes constraints with the same definitions to match; their names appear as differences (unless you also choose "Ignore Name Differences"). If the names of constraints are

meaningful, choose "Match Constraints By Name." With this choice, constraints with the same names will match and their definitions will appear as differences.

- **Partitioned Objects: Ignore Partitioning** — Choose this option to ignore partitioning in tables and indexes.
- **Partitioned Objects: Ignore High Values** — Tables that are otherwise the same might have different partition high values in different environments. Choose this option to ignore differences in high values.
- **Logical SQL Compare**— Choose this option to ignore meaningless formatting differences in source objects such packages, package bodies, procedures and functions and to ignore white space differences in comments.
- **Compare Statistics**— Choose this option to compare optimizer statistics for tables and indexes.
- **Ignore Table Column Position**— Choose this option if tables that differ only in column position should be considered equal.

Creating Comparison Versions

When you have finished defining the comparison, you specify when to create the first comparison version. You can create the first version immediately, or at a later time. You can also schedule new comparison versions at regular intervals.

In addition to scheduling comparison versions, you can create a new comparison version at any time by selecting the comparison and specifying "Repeat Now."

Comparisons processed after the initial version generally complete substantially faster than the initial comparison. Only those objects that have changed on the left or right side are compared in the new version. This also means that the storage required for additional comparison versions is only slightly larger than the storage used by the initial version, assuming a small percentage of objects has changed.

About Working with Schema Comparison Versions

A schema comparison version records the results of comparing the left and right sources in accordance with the scope specification. Objects in a comparison version have one of four states:

- **Left Only** – The object is present only in the left source.
- **Right Only** – The object is present only in the right source.
- **Identical** – The object is present in both left and right sources, and is the same.
- **Not Identical** – The object is present in both left and right sources, and is different.

The page lists all versions of a comparison and shows the number of objects in each state within each version. On the Comparison version page, you can see the objects in each state individually. Objects that are "Not Identical" can be selected to view the differences, and to generate DDL for the left and right definitions.

You can take two further actions to record information about objects in a comparison version:

- You can add a comment to an object. For example, the comment might explain why two objects are different.

- You can ignore the object. Ignoring the object removes it from lists of comparison version objects. You might ignore an object that is different if you decide that the difference is not important.

Using Schema Synchronizations

A schema synchronization synchronizes differences in database object definitions between two databases or a baseline and a database. The basic action of a database synchronization is to create or modify selected object definitions in a database to match those in another database or in a baseline.

Synchronizations are generated using synchronization specifications. For synchronizations, the scope specification does not include the names of individual objects. It can only specify the types, and the schemas to be included or excluded. You can additionally supply a prefix to limit the objects selected to those whose names start with that prefix.

Schema synchronizations synchronize differences in any attribute value between objects of any supported type. Use synchronization specifications to create multiple versions of a synchronization. Each version has a unique version number and a synchronization date. Use these versions to associate synchronizations of database/schemas made over time.

About Defining Schema Synchronizations

Source and Destination

The synchronization source provides the object definitions (and optionally, the data) from which the destination database is synchronized. A synchronization source may be a database, or a baseline version. If the source is a baseline version, it is not possible to propagate data to the destination, since a baseline does not capture data.

The synchronization destination must always be a database. The purpose of synchronization is to create or modify object definitions in the destination to match those in the source.

The options for specifying which version of a source baseline to use are similar to those used with schema comparisons. You can specify a fixed baseline version, or "Latest" or "Latest-1". If you specify "Latest," you can also request that the baseline version be captured first before synchronizing the destination from it.

When defining a baseline to be used as the source for synchronization, it is important that the baseline contain all the objects to be synchronized. For this reason, the baseline's scope specification should include at least all the schemas and object types that will be synchronized. The baseline should also include User and Role objects, along with privilege grants. A baseline to database synchronization is recommended in environments where changes are expected to be applied to the source database frequently thus necessitating the need for a point-in-time snapshot of the database, i.e. a baseline as the source of the synchronization.

Scope Specification

Defining a scope specification for a schema synchronization is similar to defining a scope specification for a schema baseline or comparison. However, there are restrictions on what you can include in the scope specification.

- You cannot specify individual objects for synchronization. You must specify object types and either schemas to include or schemas to exclude.
- Certain schemas, such as SYS and SYSTEM, cannot be synchronized.
- You cannot directly include User and Role objects for synchronization. (However, Users and Roles are automatically included as needed during the synchronization process.)
- Oracle recommends that the following object types be selected as a group: Table, Index, Cluster, Materialized View, and Materialized View Log.

The scope specification for a synchronization should be carefully tailored to the application. Do not include more schemas than are needed. For example, if you are synchronizing a module of a large application, include only those schemas that make up the module. Do not attempt to synchronize a large application (or the entire database) at one time.

Schema Map

The definition and use of the schema map is the same in schema synchronizations as in schema comparisons. When you use a schema map, object definitions in each schema map are synchronized to the mapped schema in the destination, rather than to the schema with the same name. In addition, schema-qualified references (other than those contained in PL/SQL blocks and view queries) are changed according to the schema map.

For example, assume the schema map has two entries, as follows:

- DEV1A -> DEV2A
- DEV1B -> DEV2B

Table DEV1A.T1 has an index, DEV1A.T1_IDX. It also has a foreign key constraint that refers to DEV1B.T2. Synchronize will create objects as follows:

- Table DEV2B.T2
- Table DEV2A.T1, with a foreign key reference to table DEV2B.T2
- Index DEV2A.T1_IDX, on table DEV2A.T1

Synchronization Options

Schema synchronization options are similar to the options you can specify with schema comparisons. In synchronization, the options perform two functions:

- During initial comparison of source and destination objects, the options determine whether differences are considered meaningful. For example, if the "Ignore Tablespace" option is selected, tablespace differences are ignored. If two tables are identical except for their tablespaces, no modification to the destination table will occur.
- When generating the script that creates objects at the destination, some options control the content of the script. For example, if "Ignore Tablespace" is selected, no TABLESPACE clauses are generated.

In addition to the options provided with schema comparison, the following options are specific to Synchronize:

- "Preserve Data In Destination" and "Copy Data From Source"—These two options control whether table data is copied from the source database to the destination

database. (The option is not available if the source is a baseline.) By default, Synchronize preserves data in destination tables. Choosing "Copy Data From Source" causes Synchronize to replace the destination data with data from the source table.

- "Drop Destination-Only Schema Objects"—Choosing this option directs Synchronize to drop schema objects that are present in the destination but not the source. For example, if the source contains table DEV1.T1 and the destination contains DEV1.T1 and DEV1.T2, Synchronize will drop DEV1.T2 if this option is chosen. This action applies only to schema objects that are within the scope specification. By default, Synchronize does not drop destination-only objects. Synchronize never drops non-schema objects.

Synchronization Mode

The next step in defining a synchronization is to choose the synchronization mode. There are two options:

- Unattended synchronization mode carries out the entire synchronization process in one step, ending with execution of the synchronization script. However, if an error is detected that makes it impossible to generate a correct script, the process will terminate without attempting to execute the script.
- Interactive synchronization mode pauses after initial comparison of the source and destination objects, and again after generation of the synchronization script. Interactive mode gives you a chance to examine the results of comparison and script generation, and to take appropriate action before proceeding with the next step.

Creating Synchronization Versions

When you have finished defining the synchronization, you specify when to create the first synchronization version. You can create the first version immediately, or at a later time. You can also schedule new synchronization versions at regular intervals.

Depending on the synchronization mode you select, the synchronization may run to completion (unattended mode), or pause after initial object comparison (interactive mode). In the latter case, you run each subsequent phase of the synchronization in a new job.

When using interactive mode, the destination database should not be modified from the time the objects are initially compared until the synchronization script has executed. Otherwise, the script may encounter problems. For example, assume the source has a table that the destination does not have. Object comparison notes the source-only table, and the generated script includes statements to create the table. However, after object comparison but before script execution, you manually create the table at the destination database. The script will fail when it attempts to create the table because the table already exists.

See [About Working with Schema Comparison Versions](#).

Creating a Synchronization Definition from a Comparison

You can use a schema comparison as the starting point for synchronization. Select the comparison, then choose **Synchronize**. This creates a new synchronization with the following initial information from the comparison:

- Source and destination, from the comparison's left and right sources, respectively. This means that you cannot create a synchronization from a comparison whose right source is a baseline.
- Scope specification.

 Note:

Some comparison scope specification options are not available in a synchronization. For example, you cannot synchronize individual schema objects, User objects, or Role objects.

- Comparison options.

Working with Schema Synchronization Versions

Each synchronization version represents an attempt to modify the destination objects selected by the scope specification to match the corresponding source objects. (It is "an attempt" because the process may not complete, for various reasons.) This section describes how a schema synchronization version is processed, and how you can monitor and control the process.

About the Schema Synchronization Cycle

There are three steps involved in processing a synchronization version. As noted previously, you can combine these steps into one (by choosing "Unattended Mode") or run each step separately (by choosing "Interactive Mode"). In either case, all three steps must be carried out when processing a successful synchronization version.

The following sections detail each of the three steps and describe what you can do following each step, when operating in interactive mode.

Object Comparison Step

The first step is to compare objects in the source to corresponding objects in the destination. Only those objects selected by the scope specification are compared. At the end of this step, the synchronization version has recorded all the objects. Each object is in one of the following states:

- Source Only
- Destination Only
- Identical
- Not Identical

In interactive mode, you can view the objects that are in each state, or all objects at once. For objects that are not identical, you can view the differences between the objects. At this stage, you can anticipate what will happen to each destination object:

- Source-only objects will be created in the destination.
- Destination-only objects will be unaffected, unless you chose the "Drop Destination-Only Schema Objects" option, in which case they will be dropped from the destination.

- Identical objects will be unaffected. However, if you chose the "Copy Data From Source" option, the data in tables that are identical will be replaced with data from the source.
- Non-identical objects will be modified to match the source object. Depending on the differences, the modification may be done with ALTER statements or by dropping and re-creating the object.

Before proceeding to script generation, you can exclude objects from the synchronization. For example, if you notice a source-only view that you do not want to create at the destination, you can exclude it now.

Script Generation Step

During script generation, Synchronize uses the results of the comparison step to create the PL/SQL script that will create and modify objects at the destination so that it matches the source. As part of script generation, several activities take place:

- Dependency analysis assures that the destination database provides a suitable environment for each object and that objects are created and modified in the correct order.
- Change analysis determines whether an object can be modified with ALTER statements or if it must be dropped and re-created.
- Messages are placed in the impact report for the synchronization version. The messages provide information about the synchronization process, and may report one or more error conditions that make it impossible to generate a usable script.
- The DDL statements needed to carry out the synchronization are generated and recorded in the synchronization version.

Dependency analysis examines each object to determine its relationships with other objects. It makes sure these other objects exist (or will be created) in the destination. Some examples of these relationships are:

- A schema object depends on the User object that owns it. For example, table DEV1.T1 depends on user DEV1.
- An index depends on the table that it is on. For example, index DEV1.T1_IDX depends on table DEV1.T1.
- A table that has a foreign key constraint depends on the table to which the constraint refers.
- A source object such as a package body depends on other source objects, tables, views, and so on.
- A stored object depends on the tablespace in which it is stored, unless you choose "Ignore Tablespace."

The relationships established during dependency analysis are used later in the script generation process to make sure script statements are in the correct order.

Dependency analysis may determine that a required object does not exist in the destination database. For example, a schema object's owner may not exist in the destination database. Or, a table may have a foreign key constraint on another table that is in a different schema. There are several possible outcomes.

- If the required object is in the source and is selected by the scope specification, Synchronize creates the object. For example, if DEV1.T1 has a foreign key constraint that refers to DEV2.T2, and both DEV1 and DEV2 are in the scope specification, Synchronize creates DEV2.T2.

- If the required object is a user or role, and the object is available in the source, Synchronize automatically includes the user or role and creates it at the destination. This occurs even though User and Role objects are not part of the Synchronize scope specification.
- If a required schema object is in the source but is not selected by the scope specification, Synchronize does not automatically include the object; instead, it places an Error-level message in the impact report. This restriction prevents uncontrolled synchronization of objects outside the scope specification. It is for this reason that scope specifications should include all the schemas that make up the application or module.
- If the source is a baseline version, it may not include the required object. For example, a baseline might not capture Users and Roles. Synchronize cannot look for objects outside the baseline version, so it places an Error-level message in the impact report. This is why it is important to include Users, Roles, and privilege grants in any baseline that will be used for synchronization.

At the end of the script generation step, Synchronize has added the impact report and the script to the synchronization version. In interactive mode, you can examine the script and impact report before proceeding to script execution.

The impact report contains messages about situations that were encountered during script generation. There are three types of messages:

- Informational messages require no action on your part. They simply report something of interest. For example, when script generation automatically includes a User object, it adds an informational message to the impact report.
- Warning messages report a situation that requires your attention, but that may not require action. For example, if Synchronize is unable to determine if a reference in a source object can be resolved, it adds a warning message to the impact report. You need to verify that situations reported in warning messages will not prevent script execution.
- Error messages indicate a situation that will prevent script execution if not corrected. For example, if Synchronize is unable to locate a required dependency object, it adds an error message to the impact report. Depending on the message, you may be required to create a new synchronization. For example, if the dependency object is not in the synchronization scope, or if the source is a baseline that does not contain the dependency object, you will need to create a new synchronization with an expanded scope or a different source baseline. In other cases, you can resolve the situation by excluding one or more objects from the synchronization and regenerating the script.

The script display contains the statements of the generated script in the order they will be executed. You can examine the script if you have any concerns about its correctness. The display allows you to locate statements that are associated with a particular object or object type.

Following script generation, you can continue to script execution unless an error was encountered during the script generation step. In this case the impact report will contain one or more Error-level messages detailing the problem and solution. In some cases, you may be able to solve the problem by selecting "Regenerate Script," excluding an object from the synchronization, and regenerating the script.

There may be cases where you need to create a new version of the synchronization in order to correct the problem. For example, if you need to modify the definition of an object in the source or destination or add an object in the destination, you will need to

create a new version. This allows the new or modified object to be detected during the comparison step. In this case, the old version becomes "abandoned" since you cannot continue to script generation.

Script Execution Step

Following successful script generation, the script is ready to execute. In unattended mode, the script executes as soon as script generation completes. In interactive mode, you proceed to script execution as soon as you have reviewed the impact report and the script.

The script executes in the Cloud Control job system. Once script execution is complete, you can view the execution log. If the script fails to execute successfully, you may be able to correct the problem and re-start the script execution from the point of failure. For example, if the script fails due to lack of space in a tablespace, you can increase the size of the tablespace, then re-start the script.

Creating Additional Synchronization Versions

Following processing of the initial synchronization version, you can create additional versions of the synchronization. Select the synchronization, then choose "Synchronize Again." Note that you cannot choose a different source or destination, or modify the scope specification or synchronization options, when creating a new synchronization version. However, you can choose a different mode (Unattended or Interactive) when starting a new synchronization version.

Creating additional synchronization versions allows you to propagate incremental changes in the source to the destination. For example, you can refresh a test database from the latest changes made to a development system.

Synchronizations processed after the initial version generally complete substantially faster than the initial synchronization.

Using Change Plans

Change Plans are a new feature of the Cloud Control Database Lifecycle Management Pack. Change Plans complement and extend the capabilities of existing Change Management components by allowing users to select and package metadata changes for deployment to multiple databases. Change Plans support database application development methodologies that are not adequately supported by existing Database Lifecycle Management Pack tools such as Schema Synchronizations.

Change Plans are flexible enough to support a variety of development methodologies, yet powerful enough to automate many database administration tasks previously carried out with custom scripts. These tasks include:

- Deploying project-specific development changes from a shared development database to one or more destination databases such as integration, test, or production staging.
- Deploying development changes from a stand-alone project development database to an integration database that collects changes from multiple development databases.
- Upgrading common modules in development databases from a central integration database.

Change Plans are tightly integrated with the other tools in the Database Lifecycle Management Pack. Specifically:

- Change Plan change requests that create objects can get the object definitions from Change Management Schema Baselines.
- Change requests that modify objects can use the contents of an object in a Change Management Schema Comparison to specify the change.
- Change Plans complement Change Management Schema Synchronizations, allowing for finer control of changes and “change-only” change requests.

About Working with Change Plans

The first phase of using a change plan to create or modify object definitions is to plan and define the changes that you want to make. For example, you may want to make one or more changes to an existing object definition in one or more databases. Or, you may want to reproduce one or more object definitions from one schema or database in another schema or database.

Figure 52-1 Steps in a Change Plan




Figure 52-1 shows the steps in a change plan. A change plan is a named container for change requests. You can define change requests to reproduce or modify object definitions at a destination database. A destination database is a database where you want to apply the change requests in a change plan. After you finish planning and defining the changes, evaluate the impact of the changes that you want to make.

To evaluate the impact of the change requests at a particular database, generate a script and an impact report for a change plan and that destination database. The impact report explains the changes that will be made by the script when it executes at the destination database. It also describes any change requests that cannot be applied at the destination database.

To implement the change requests in a change plan at a destination database, execute the script at the destination database.

Creating a Change Plan

This section explains the different methods of creating change plans.

You can create change plans through any of the following ways:

- [Creating and Applying a Change Plan From a Schema Comparison](#)
- [Using External Clients to Create and Access Change Plans in Cloud Control](#)

Creating and Applying a Change Plan From a Schema Comparison

This section explains how to create a change plan from a schema comparison.

Prerequisites for Creating a Change Plan

Following are the prerequisites:


- Ensure that the Application Developer (AD) is an Cloud Control user who has the following privileges:
 - Connect Target privilege to the development and production-staging databases targets or Connect Any Target privilege
 - DBA privileges to the development database
 - Create Privileges for Job System (Resource Privilege)
 - Create new Named Credential (Resource Privilege)
 - Edit Resource Privilege on the change plans
 - Execute Command Anywhere (Target Privilege)
 - EM_ALL_OPERATOR privilege
- Ensure that the Database Administrator (DBA) is an Cloud Control user who has the following privileges:
 - Connect Target privilege to the development and production-staging databases targets or Connect Any Target privilege
 - DBA privileges to the development database
 - Create Privileges for Job System (Resource Privilege)
 - Create new Named Credential (Resource Privilege)
 - Manage Change Plans (Resource Privilege)
 - Execute Command Anywhere (Target Privilege)
 - EM_ALL_OPERATOR privilege
- It is recommended that the development and destination databases are identical at the start of the development work. For example, they may both be at the current production version, or both updated to a common interim development version.
- The Application Developer would have made changes in the development database. After creating a change plan, the application developer can create and update change items in the change plan through external clients such as SQL Developer. For more information, see [Using External Clients to Create and Access Change Plans in Cloud Control](#).

Creating a Change Plan


Follow these steps to create a change plan:

1. Log in to Cloud Control as a database administrator (DBA).
2. Identify the schemas that contain application objects.
3. Use Metadata Baselines wizard to define a baseline that includes the schemas of interest. Schedule a job to capture the first version of the baseline.
4. Save the baseline.
5. Use Schema Comparisons wizard to define a comparison between the baseline version and the development database.

6. Schedule a job to create the first version of the comparison and save the comparison.
7. In the Schema Change Plans page, click **Create**.
8. Specify a **Name** and **Description** for the change plan and click **OK** to save the change plan.



9. In the Change Items page, click **Create From Comparison**.
10. In the Create Change Items from Schema Comparison page, select the Comparison Version created earlier, specify the development database as the Change To side and the production-staging database as the Change From side in the Conversion Assignment and click **OK**.



11. In the Create Change Items from Schema Comparison: Select Differences page, select:
 - All Differences in the Schema Comparison to add all differences in the comparison to the change plan
 - Specific Differences in the Schema Comparison to select the differences in the comparison you want to add to the change plan. Select the differences.

Click **Finish**.

12. Submit request to apply the Change Plan on the destination database.

Applying a Change Plan

Follow these steps to apply a change plan:

1. Log in to Cloud Control as a database administrator (DBA).
2. In the DBA role, examine the Change Plan, evaluating its suitability for application to the proposed database. Remove individual Change requests, if required.

3. From the Schema Change Plans page, select **Create Synchronization from Change Plan**.
4. Specify the details in the Schema Synchronization wizard with the source as the Change Plan instance created earlier. For information about using the Schema Synchronization wizard, see *Synchronizing with Production Staging*. By default, the synchronization created from change works in the interactive mode.
5. Schedule script generation.
6. Check Impact Report and schedule script execution.
7. Check completed script execution job for errors. If the change plan job failed, do the following:
 - If the failure is due to a condition noted in an impact report error warning, perform the suggested user action.
 - If the failure is due to a condition in the source or destination database that can be fixed manually, fix the problem and perform the operation again.
 - If the failure is in the script execution phase, view the script output in the job details. If the problem can be resolved by actions such as issuing missing grants, fix the problem manually in the database and then click **Retry Script Execution**.
8. Fix the errors and submit the change plan creation job again.

Using External Clients to Create and Access Change Plans in Cloud Control

Cloud Control provides support for external clients such as SQL Developer to create and access change plans. You can use these applications to connect to the Cloud Control repository and create change plans and add and update change items in them.

Client users are of two types:

- Users who can create and access all change plans
- Users who can access (view and possibly edit) specific change plans

Following are the steps:

1. Configure the repository database listener to allow access by a trusted client. It is recommended that you make the repository database inaccessible to login from non-trusted clients. For information about configuring the listener, see *Oracle® Database Net Services Administrator's Guide 11g Release 2 (11.2)*.
2. Set up an Cloud Control administrator for use by an external client.

The following section describes how to set up administrators for change plans, for access from Cloud Control and from external clients.

Setting Up Cloud Control Administrator For Change Plans

Follow these steps:

1. Log in to Cloud Control as a super administrator.
2. From the **Setup** menu, click **Security** and then select **Administrators**.
3. In the Administrators page, click **Create**.

4. In the Create Administrator: Properties page, specify the Name and Password for the user. This creates a database user with the specified name and password, as well as creating the Cloud Control administrator. Click **Next**.
5. In the Create Administrator: Roles page, click **Next**.
6. In the Create Administrator: Target Privileges page, click **Next**.
7. In the Create Administrator: Resource Privileges page, select Change Plan Security Class and click the **Manage Privilege Grants** icon.
8. In the Create Administrator: Manage Privileges page, do the following:
 - If you want to create an administrator who has all access to all change plans, select **Manage Change Plans** in the Resource Type Privileges section.
 - If you want to create an administrator who has specific access to one or more change plans, click **Add** in the Resource Privileges section. In the list of change plans that have been created already, select one or more and click **Select**. The selected plans are added to the Resource Privileges section. By default, the administrator is granted View Change Plan privilege; you can edit this to grant Edit Change Plan privilege.
9. Click **Continue**.
10. In the Create Administrator: Review page, click **Finish** to create the new administrator.
11. For an external client to be able to access change plans using any of these privilege types, follow these steps:
 - a. Log in to the repository database as a user with DBA privileges.
 - b. Grant the **CHANGE_PLAN_USER** database role to the database user corresponding to the new administrator (through Schema->Users in Enterprise Manager Administrator, or in SQL Plus).

Submitting Schema Change Plans From SQL Developer Interface

To enable developers to submit their schema changes to Enterprise Manager Schema Change Plans through SQL Developer interface, perform the following manual configuration steps:

1. Ensure that the repository administrator has configured the repository database to accept remote database connection from SQL developer. You can do this by configuring the repository listener process.
2. Create a local administrator account on the OMS.
3. Provide the repository user of the local OMS account privileges to be a change plan user by running the following SQL commands on the repository database as user SYS:

```
grant CHANGE_PLAN_USER to PUBLIC;
```

or:

```
grant CHANGE_PLAN_USER to <repos_user>;
```

4. Edit the OMS users resource privileges to give the user access to edit the change plans.

Using Database Data Comparison

A data comparison operation compares data in a set of database objects in a candidate database with those in a reference database. To compare objects residing in the same database, select that database as both the reference and the candidate. You can create a comparison specifying which objects are to be compared and submit a Cloud Control job to compare them immediately or at a later time. On job completion, select the data comparison and view results. The results will be purged when you delete the comparison.

Cloud Control data comparison uses DBMS_COMPARISON package for comparison. It can compare the following types of database objects:

- Tables
- Single-table views
- Materialized views
- Synonyms for tables, single-table views, and materialized views

Database objects of different types can be compared at different databases. For example, a table at one database and a materialized view at another database can be compared.

Requirements for Database Data Comparisons

For data comparison, you will need to meet the requirements explained in this section.

The database character sets must be the same for the databases that contain the database objects being compared.

For index column, the number, timestamp, and interval columns datatypes are as follows:

- Number columns are of the following datatypes: NUMBER, FLOAT, BINARY_FLOAT, and BINARY_DOUBLE.
- Timestamp columns are of the following datatypes: TIMESTAMP, TIMESTAMP WITH TIME ZONE, and TIMESTAMP WITH LOCAL TIME ZONE.
- Interval columns are of the following datatypes: INTERVAL YEAR TO MONTH and INTERVAL DAY TO SECOND.

The database objects must have one of the following types of indexes:

- A single-column index on a number, timestamp, interval, DATE, VARCHAR2, or CHAR datatype column.
- A composite index that only includes number, timestamp, interval, DATE, VARCHAR2, or CHAR columns. Each column in the composite index must either have a NOT NULL constraint or must be part of the primary key.

If the database objects do not have one of these types of indexes, then the EM data comparison does not support the database objects. For example, if the database objects only have a single index on an NVARCHAR2 column, then the data comparison does not support them. Or, if the database objects have only one index, and it is a composite index that includes a NUMBER column and an NCHAR column, then the data comparison does not support them.

The index columns in a comparison must uniquely identify every row involved in a comparison. The following constraints satisfy this requirement:

- A primary key constraint.
- A unique constraint on one or more non-NULL columns.

If you specify an index, then make sure the columns in the index meet these requirements for comparison.

Data Comparison feature in Cloud Control can compare data in columns of the following datatypes:

- VARCHAR2
- NVARCHAR2
- NUMBER
- FLOAT
- DATE
- BINARY_FLOAT
- BINARY_DOUBLE
- TIMESTAMP
- TIMESTAMP WITH TIME ZONE
- TIMESTAMP WITH LOCAL TIME ZONE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- RAW
- CHAR
- NCHAR

If a column with datatype `TIMESTAMP WITH LOCAL TIME ZONE` is compared, then the two databases must use the same time zone. Also, if a column with datatype `NVARCHAR2` or `NCHAR` is compared, then the two databases must use the same national character set.

Data comparison feature cannot compare data in columns of the following datatypes:

- LONG
- LONG RAW
- ROWID
- UROWID
- CLOB
- NCLOB
- BLOB
- BFILE
- User-defined types (including object types, REFs, varrays, and nested tables)
- Oracle-supplied types (including any types, XML types, spatial types, and media types)

You can compare database objects that contain unsupported columns by excluding the unsupported columns when providing comparison specification. Edit the comparison item and include only the supported columns in the Columns To Include list of column names.

Since data comparison cannot compare LOB column values directly, their cryptographic hashes will instead be used for comparison. If you include LOB type columns to be compared, make sure that the database users connecting to the reference and candidate databases have EXECUTE privilege on SYS.DBMS_CRYPTO package. For more information about DBMS_COMPARISON, see *Database PL/SQL Packages and Types Reference* book for the database version of your reference database.

 Note:

A Data Comparison job may fail with the error "ORA-28759: failure to open file error."

This failure occurs when data comparison tries to get data from the candidate database into the reference database over a database link in the reference database for comparing them.

The database server (candidate/source database) requires the use of TCPS protocol for connections, but the client (reference/destination database) does not have a valid wallet location. Connection over the database link fails since no wallet was specified on the client side.

This problem can be fixed by specifying a valid WALLET_LOCATION entry in sqlnet.ora file (which is by default located in the \$ORACLE_HOME/network/admin directory). The following wallet location must be specified at the reference database:

```
WALLET_LOCATION = (SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/net/slc05puy/scratch/dbwallets/s wal-
lets)))
```

Comparing Database Data and Viewing Results

The following procedure enables you to specify which pairs of objects you want to compare in the reference and candidate databases, submit a job to process your choices, then view the differences after the job successfully completes.

1. From the main Data Comparisons page, click **Create**. The Create Data Comparison page appears.
2. Provide the required input:
 - a. If you want to compare objects residing in two databases, select one database as the Reference and the other as the Candidate.
 - The Reference database always executes the comparison, so it must be version 11g or later. The Candidate database must be version 10g or later.
 - Be advised that the Reference database carries an additional processing load and requires some space to store the row IDs of differing rows (not the entire rows themselves). If you compare data between a production

system and a test system, it might be appropriate to process and store the results on the test system.

- b. Click **OK** when you have finished. The Data Comparison Specification page appears.

 Tip:

It is recommended that you define the comparison specification once and run it many times.

3. Open the **Actions** menu, then select **Add Object Pair** or **Add Multiple Objects**. If you select Object Pair, continue with the following sub-steps. If you select Multiple Objects, go to the next step.

- Adding an object pair consists of selecting one object from the reference database and one object from the candidate database. You can compare dissimilar object types, if desired, such as a table in the reference database and a materialized view in the candidate database.

- a. Specify the reference and candidate objects. The reference database can be the same as the candidate database. In this case, the objects are from the same database.
- b. Select one or more columns in the reference or candidate databases for comparison. The columns included must be common to both objects.
- c. Optionally select an index to be used for comparison. Columns in the comparison index must uniquely identify every row involved in a comparison. An index used for a primary key constraint or a unique constraint on one or more non-NULL columns satisfies this requirement. The comparison can use the specified index only if you select all of the columns in the list of Columns To Include.
You can select a composite index if you want to add multiple index columns.
- d. Specify an optional Where Condition per pair of objects being compared.
- e. Either specify or let the system compute the maximum number of buckets and minimum rows per bucket.
- f. Specify the point in time you want to compare data.
 - The System Change Number (SCN) is a sequential counter that uniquely identifies a precise moment in the database. This is the most accurate way to identify a moment in time. Whenever you commit a transaction, Oracle records a new SCN. You can obtain SCNs from the alert log.

- g. When you have finished the configuration, click **OK**.

The Data Comparison Specification page reappears, showing your selected objects in the list.

- h. Click **OK**, then go to Step 5.

4. Open the **Actions** menu, then select **Add Multiple Objects**.

- Adding multiple objects enables you to conveniently perform a bulk inclusion of multiple objects from the reference database into the specification. You can search and select multiple objects, such as many tables and views, from the reference database list of values, and then edit each item as needed.
- a. Specify the schema name, one or more object types, then click **Search**.

The table populates with object names.

- b.** Select the objects you want to compare, then click **OK**.

The Data Comparison Specification page reappears, showing your selected objects in the list.

5. Select your comparison name from the list, open the **Actions** menu, then select **Submit Comparison Job**. For information about privileges required for user credentials for the reference and candidate databases, see [Overview of Change Management for Databases](#).

6. Provide the required credentials in the page, schedule the job, then click **OK**.

The Data Comparisons page reappears and displays the following confirmation message:

"The job was submitted successfully. Click the link in the Job Status column to view job status."

After the Job Status column shows Succeeded, go to the next step.

7. Select your comparison name from the list, open the **Actions** menu, then select **View Results**. The Data Comparison Results page appears.

- Look for rows in the Result column with the `!=` symbol, indicating that there are differences between reference row and candidate row data.

- Data comparison attempts to compare all tables. If there is an error, you can see the error message by selecting the **Messages** tab. An error message is indicated with an X instead of the = or **!=** symbol.
- You can see the SQL statements that are running to perform the comparison by clicking the **Executed Statements** tab.

9. Select a dissimilar Reference/Candidate row, then click **View Row Differences** to see a detailed, indexed list of reference-only, candidate-only, and non-identical changed rows on the Row Data Differences page.

- The Row Source column indicates the origin of each row of data as a whole. Furthermore, data in a row differing between reference and candidate are displayed in contrasting colors, indicating whether the source of the data is the reference or candidate database.
- The comparison is shown based on a key column (depending on a chosen unique index). If the key column value is different, the row appears as a candidate or reference-only row. If other columns are different, the row appears as a non-identical row.

Schema Mapping

By default, a reference object will be compared with a candidate object in the same-named schema as the reference schema. Using schema mapping, you can optionally compare objects in a reference schema with objects in a different candidate schema. Any schema can only be mapped once. Provide reference and candidate schema names for mapping under the Schema Mapping section of the Data Comparison Specification page. Default candidate schema will then be picked from schema mapping you specified.

You may further override the candidate schema of individual item by editing the item, clicking the Override button next to the Candidate Object field, and explicitly specifying the candidate object belonging to any schema. For such items whose candidate objects are overridden in this way, schema mapping will be ignored.

Usage of Buckets

A bucket is a range of rows in a database object that is being compared. Buckets improve performance by splitting the database object into ranges and comparing the ranges independently. Every comparison divides the rows being compared into an appropriate number of buckets. The number of buckets used depends on the size of the database object and is always less than the maximum number of buckets specified for the comparison by the maximum number of buckets specified.

When a bucket is compared, the following results are possible:

- No differences are found —
The comparison proceeds to the next bucket.
- Differences are found —
The comparison can split the bucket into smaller buckets and compare each smaller bucket. When differences are found in a smaller bucket, the bucket is split into still smaller buckets. This process continues until the minimum number of rows allowed in a bucket is reached, whereupon a comparison reports whether there are differences in the bucket and identifies each row difference.

You can adjust the maximum number of buckets and minimum rows per bucket to achieve the best performance when comparing a particular database object.

The comparison program uses the `ORA_HASH` function on the specified columns in all the rows in a bucket to compute a hash value for the bucket. If the hash values for two corresponding buckets match, the contents of the buckets are assumed to match. The `ORA_HASH` function efficiently compares buckets, because row values are not transferred between databases. Instead, only the hash value is transferred.



Note:

If an index column for a comparison is a `VARCHAR2` or `CHAR` column, the number of buckets might exceed the value specified for the maximum number of buckets.

Additional Setup for Real-time Monitoring

This section describes Oracle Enterprise Manager Cloud Control's (Cloud Control) Compliance features include the ability to monitor certain elements of your targets in real time to watch for configuration changes or actions that may result in configuration changes.

These features include Operating System level file change monitoring, process starts and stops, Operating System user logins and logouts, Oracle database changes and more.

The real-time monitoring for these features takes place from the Cloud Control agent. Some of these monitoring capabilities require specific setup steps depending on the type of monitoring you will do and what Operating System is being monitored.

This chapter outlines the specific requirements and pre-requisites that exist to use the Compliance Real-time Monitoring features. For details on how to use Real-time monitoring from Cloud Control, see the chapter Compliance Management in this document. This chapter covers the following topics:

- [Overview of Real-Time Monitoring](#)
- [Overview of Resource Consumption Considerations](#)
- [Configuring Monitoring Credentials](#)
- [Preparing To Monitor Linux Hosts](#)
- [Preparing To Monitor Windows Hosts](#)
- [Preparing To Monitor Solaris Hosts](#)
- [Preparing to Monitor AIX Hosts](#)
- [Preparing To Monitor the Oracle Database](#)
- [Setting Up Change Request Management Integration](#)
- [Overview of the Repository Views Related to Real-time Monitoring Features](#)
- [Modifying Data Retention Periods](#)
- [Real-time Monitoring Supported Platforms](#)

Overview of Real-Time Monitoring

Real-time monitoring is configured through the Cloud Control Server. Users with the EM_COMPLIANCE_DESIGNER role create Compliance Standard Rules that are of type "Real-time Monitoring Rule." These rules are then associated with Compliance Standards and these standards are subsequently associated with one or more targets.

After the Compliance Standard to target association is complete, the set of monitoring rules are sent to the agent to enable real-time monitoring. All monitoring for Real-time monitoring occurs on the agents and all observed action data is sent from the agent to the Cloud Control server for reporting and data management.

Overview of Resource Consumption Considerations

The Real-time monitoring features are built into the Cloud Control agent. There are some specific resource considerations if you use the Real-time monitoring features. The following sections describe issues you should consider when using Real-time monitoring features.

OS File Monitoring Archiving

An optional setting when monitoring for file changes in real time is to make an archive copy of the file on the agent. When monitoring first begins, a copy of the file at that time is made and stored into a private directory in the `ORACLE_HOME` directory of the agent. Then, any subsequent changes to that file will result in additional copies of the file being archived in that same directory. This feature allows you to later perform a file diff from the user interface or to issue a job to roll back a file to a previous version.

This feature however will use disk space to make copies of the file. Care should be taken to ensure that this feature is only enabled for files that are critical. During rule creation, the user can specify how many copies of a file to save. The default is five historic versions. This can also be adjusted to tune potential resource consumption.

Select the checkbox in the Ignore Events Prior to Rule field to ignore all previous Oracle database change events when the Oracle database monitoring module runs the first time.

OS File Read Monitoring

The Operating System File level monitoring can monitor many types of changes to files, but can also monitor reads to files. If you have a Rule to monitor a facet that has file patterns that are read frequently, this may result in a very large number of observations. You can reduce the number of observations by ensuring that your Rule includes a filter on either time or a user that you want to ensure does not read the file.

For instance, monitoring the `/etc/passwd` file for reads for All users will result in many observations being created. However, if you only monitor the `/etc/passwd` file by a specific user, you can create a user filter for this specific user during rule creation. You will then only receive an observation when that specific user attempts to read the file.

Creating Facets That Have Very Broad Coverage

It is important to remember that facets are created to specify files that are very important to monitor for security/compliance purposes. For instance, monitoring all modifies to a log file that change every few seconds will result in reporting many file changes making it harder for you to identify the critical file changes you care about. Instead, in this case, it may be appropriate to create a rule to monitor the log file for all changes, but filter only when the log change is made by a non-application user. This would only capture the log file change if a regular user attempted to change or tamper with the log rather than when the log is simply being updated by an application.

Cloud Control Repository Sizing

Database sizing considerations for Real-time monitoring depend on several factors. The most important factor is the number of observations expected in a month. The second factor is the number of months data will be retained in the repository. Repository retention rates are explained in the Enterprise Manager Administration Guide.

In general, each observation consumes roughly 1.5KB of space in the database. This is a guideline and this number can vary depending on many factors for each installation.

For example, if a customer expected a total of 10 million Real-time observations per month across all targets and wanted to retain the data for 12 months, then the database size required for this would be roughly 180GB.

10,000,000 Observations x 12 Months x 1500 Bytes = 180,000,000,000 Bytes

This size represents Real-time monitoring data only and does not include database storage needs for other areas of Cloud Control.

The number of observations to expect per month can vary from environment to environment and can also depend on what types of monitoring are configured. You may be required to tune the expected size over time after Rules and Facets have been enabled for some time and configured to fit the organizations requirements. You can easily find your observations usage over a month by selecting **Compliance** from the **Enterprise** menu, then choosing **Browse By Systems UI Report** from the **Real-time Observations** page to select your systems and see the related counts of observations for each system over a period of time.

Configuring Monitoring Credentials

Many of the real-time monitoring capabilities require monitoring credentials that maintain the ability to launch monitoring programs with root privileges. These processes that Real-time monitoring uses begin with the prefix *nmxc*. Low-level monitoring uses operating system APIs that are not available to regular users.

Before starting to use the Real-time monitoring features on a target host for the first time, the following settings must be configured from the Enterprise Manager Console.

1. Ensure that the agent's *root.sh* script is run after agent installation.

After installing the agent, the *root.sh* script must be run as the root user. This script must be run before configuring the rest of these credential steps.

2. Configure Privilege Delegation.

Privilege Delegation settings are found from the **Setup** menu by choosing **Security**, then **Privilege Delegation**. On this page you can either set privilege delegation for each host manually or you can create a Privilege Delegation Setting Template.

Privilege delegation for each host that will have real-time monitoring must have SUDO setting enabled with the appropriate SUDO command filled in (for example, */usr/local/bin/sudo*).

3. Configure Monitoring Credentials.

Monitoring Credential settings are found from the Setup menu. Choose **Security** then **Monitoring Credentials**. From this page, select the Host target type and click **Manage Monitoring Credentials**.

For each entry with the credential “Host Credentials For Real-time Configuration Change Monitoring”, select the entry and click **Set Credentials**. You will be asked for a credential set to use. Ensure you also add “root” to the Run As entry. If “Run As” is not visible, then the privilege delegation was not set properly in the previous step.

To set monitoring credentials in bulk on multiple hosts at once, you can use EMCLI. For more information on using EMCLI to set monitoring credentials, see the section, *Managing Credentials Using EMCLI* in the Security chapter of *Oracle Enterprise Manager Administration*. Likewise, for more information about configuring monitoring credentials in Cloud Control, the Security chapter of *Oracle Enterprise Manager Administration*.

Preparing To Monitor Linux Hosts

The following sections describe how to prepare Linux hosts for monitoring.

OS File Monitoring

Before using Real-time file monitoring for Linux, a loadable kernel module must be installed on the host. This loadable kernel module provides you with the most efficient way of monitoring the host. This loadable kernel module is referred to as the File Audit Module, or Audit Module for short.

Acquiring the Kernel Module

The kernel audit module is available from <http://oss.oracle.com/projects/fileauditmodule>. There are two ways to get the file audit kernel module:

1. **Prebuilt .ko files** for which Oracle has already prebuilt, you can use this in your environment. You can look for the Prebuilt kernel modules under the **Downloads** link. To find the matching prebuilt version, run the `uname -r` command on the host being monitored and compare that version to the version of the prebuilt modules. The complete version string must match perfectly. For 32-bit machines, the post-fix of the .ko file name will be .ko. For 64-bit machines, the post-fix of the .ko file name will be .k64.ko.
2. **Build your own kernel module.** To build your own kernel module, you can download the following RPM from the **Downloads** link:

Fileauditmodule-emversion-revision-noarch.rpm

You should always retrieve the latest revision available at the time you are installing this module. The emversion field must match the version of Cloud Control agent and server you are using.

Install this RPM on the host you want to monitor as root. The installation of this RPM depends on the kernel-devel package matching your running kernel also existing on the host. This kernel-devel package comes with the same media as the Linux installers.

In addition to installing this package, you must ensure that the version of gcc available on your host matches the version with which the kernel was built. To do this, view the `/proc/version` file to see what gcc version the kernel was built with and

then run the command `gcc -v` to see what version of gcc is being used. These two versions should match.

Also check that the file `/boot/System.map-{version}` exists where `{version}` must match the kernel version you see when you run the `uname -r` command. This file contains system symbols that are required to decode the kernel symbols we are monitoring for real-time changes. Without this file, real-time file monitoring will not function. This file is standard on all default Linux installations.

After installing this package and checking prerequisites successfully, go to the directory where the package contents were installed (defaults to `/opt/fileauditmodule`) and run the following script:

```
compmo.sh
```

This will build the kernel module file (`.ko`, `.k64`, or `.o` extension depending on the OS version) and place it in the `/opt/fileauditmodule` directory.

If the audit module file is not created, check the `make.log` and `build.log` files for any errors in building the module.

If all of your hosts have the exact same kernel version as shown using the command `uname -r`, then you only need to compile the module on one machine. You can then copy the `.ko`, `.k64`, or `.o` file to the other servers without having to build on that specific host.

Deploying the Kernel Module

Once you have either the prebuilt `.ko` file or a `.ko` file that exists from building it from the source RPM, the `.ko` file must be located in the proper directory. The default location for this file is in the `bin` folder under the agent home directory. You can also place the file in any location on the host and change the `nmxc.properties` file under the `AGENT_INST/sysman/config` directory of the agent home. The property `nmxcf.kernel_module_dir` specifies the absolute path to the `.ko` directory.

Install Kernel Module Job

In addition to manually placing the `.KO` file on the agent, there is a Cloud Control job named *Real-time Monitoring Kernel Module Installation*. This job is configured with a list of Linux hosts on which you can install the kernel module. It will search in a directory locally on the Cloud Control server disk for prebuilt `.ko` files or the source RPM file. If it finds a matching prebuilt `.ko` file, it will send this to the matching agents; otherwise it will send the RPM to the agent and install and compile it resulting in a new `.KO` file.

Prior to using this job, files from OSS.Oracle.com must be manually retrieved by the user and placed into the `%ORACLE_HOME%/gccompliance/fileauditmodule/resources/linux` directory. This directory already exists on the server with a README file indicating this is the location to place these files. The files that must be placed here are either prebuilt `.KO` files or the source RPM file. If you have built your own `.KO` files in your environment, you can also place those `.KO` files into this directory on the server and deploy it to other hosts in your environment.

Special Considerations for Enterprise Linux 5 and Greater

For Enterprise Linux 5 and greater, the kernel audit module is not required. The monitoring will use the built-in audit subsystem if a kernel module is not detected at startup time. However, the functionality of the audit subsystem is not as robust as the capability that the kernel audit module can provide.

You will lose the functionality that provides the granularity of what type of change there has been to a file, whether it was a create action or a modify action. Without the kernel

module, all changes to a file will appear as a modify action. Additionally, monitoring a directory that does not exist yet or a directory that may exist now and gets removed later may be disrupted since the underlying Linux audit subsystem does not handle these cases.

It is recommended that you use the kernel audit module even with the newer versions of Linux, if possible.

Debugging Kernel Module Or Other File Monitoring Issues

You may detect a problem with the kernel module in a few different ways:

1. You may have noticed that you do not receive real-time file changes on the Enterprise Manager console for file changes that you know should occur.
2. In the Compliance Standard Target Associations or Real-time Observations page on the user interface, you may see an agent warning indicating a kernel module problem.
3. When examining the *nmxcf.log* file under *AGENT_INST/sysman/logs*, you may see errors indicating that the kernel module could not be loaded or used for various reasons.

If you encounter any of these issues, most likely there was a problem with compiling or inserting the Linux kernel module at run time.

You can confirm whether the auditmodule was loaded properly by running the following command.

```
grep -i auditmodule /proc/modules
```

If you do not receive any output, then the auditmodule is not loaded and the agent will not perform real time file monitoring.

If the audit module file was generated properly and it does not show up in the module list above, you can attempt to manually load the module to see if there are any errors. Use the following command where you replace {audit module file name} with the entire name of the .ko file that was created from *compmmod.sh*:

```
insmod {audit module file name}
```

If you experience no errors during this command, you can check the module list again by using the *grep* command above. If the audit module now appears, then the file monitoring capability should work. An agent restart is necessary; however there still may be a problem with the file monitoring process finding the .ko file which you will experience again next time your host is rebooted.

One additional step to debug any issues with the file monitoring process is to try to run it manually. To do this, follow these steps:

1. Get the process ID of the agent TMMMain process:

```
ps -eaf |grep TMMMain
```
2. Execute the nmxcf process using the following command replacing the values in {} with the proper path elements or the process ID from the previous command:

```
sudo {agent_home}/core/{agent_version}/bin/nmxcf -e {agent_home}/  
agent_inst/ -m {agent_home}/agent_inst/sysman/emd/state/fetch-  
let_state/CCCDDataFetchlet -w {process id of TMMMain}
```

Running the `nmxcf` process this way will not work in the long term since it will not start up again when the agent is restarted, but this can help in trying to debug any issues as to why the process cannot start.

If the module still is not able to load and if you need to contact Oracle support about the issue, please be sure to include the following information with your support ticket:

- Output of the command: `uname -a`
- Output of the command: `grep -i auditmodule /proc/modules`
- Output of the command: `rpm -q -a |grep -i kernel-devel`
- The `make.log` and `build.log` files from the `/opt/fileauditmodule` directory where you ran `compmod.sh` if you built your own `.ko` file
- The files `AGENT_INST/sysman/log/nmxc*.log`
- Any warnings or errors you received when trying to start `nmxcf` manually.

This information will help Oracle Support to determine if the real time file monitoring audit module of the agent can be built on your environment.

 Note:

Be careful when using the Linux OS command `rmmod` which is used to unload a kernel module. If the `nmxcf` binary is running and you use `rmmod`, there is a chance that a kernel panic can arise by trying to unload a kernel module in use. The use of `rmmod` in Linux should be done carefully no matter which module you are unloading.

Preparing To Monitor Windows Hosts

The Real-time monitoring features support Windows 2003 and 2008 Server along with Windows XP. The Real-time monitoring modules for Windows rely on various capabilities of the operating system to collect all of the information on actions. One part of this is to capture the user that made changes from the Windows Event Log. If you do not configure Windows to capture users that make changes, the agent will not capture this information. However it will still capture that a change occurred and when it occurred.

To configure the event log to work with real time monitoring, perform the following steps:

1. From Windows Explorer, select the directory that is being monitored by a Real-time Monitoring Rule, right-click and select **Properties**.
2. Go to the Security tab.
3. Click **Advanced**.
4. Select the Auditing tab.
5. Click **Add**. (In Microsoft XP, double-click the **Auditing Entries** window).
6. Select the Name **Everyone**, then click **OK**. You can also choose specific users if you are only monitoring for changes by specific users in Configuration Change Console rules. The rules filter the results by user as well, so even if you enable

audit for everyone, only users that you want to monitor changes of in your rules will be captured.

7. Select the following options (Successful and/or Failed) from the Access window. For Windows XP and Windows 2003:
 - Create Files/Write Data
 - Create Folders/Append Data
 - Delete Files Subfolders and Files
 - DeleteFor Windows 2008 and Windows 7:
 - Create Files/Write Data
 - Create Folders/Append Data
 - Write Attributes
 - Write Extended Attributes
 - Delete Files Subfolders and Files
 - Delete
 - Change Permissions
 - Take Ownership
8. Click **OK** to exit.
9. Repeat steps 1 through 7 for all other monitored directories and/or files.
10. From the **Start** menu, select **Settings**, then **Control Panel**, then **Administrative Tools**, then **Local Security Policy**, then **Local Policies**, then **Audit Policy**. Double-click and turn on the following policies (Success and/or Failure):
 - Audit account logon events
 - Audit logon events
 - Audit object access
11. Close the Local Security Settings screen.
12. From the **Start** menu, select **Settings**, then **Control Panel**, then **Administrative Tools**, and finally **Event Viewer**.
13. Select **System Log**, then click **Action** from the menu bar and select **Properties**.
14. From the System Log Properties panel, on the General tab, set the Maximum log size to at least 5120 KB (5 megabytes) and select **Overwrite Events as Needed**. Note that the log size depends on the number of events generated in the system during a two-minute reporting interval. The log size must be large enough to accommodate those events. If you extend the monitoring time for file events because you expect the change rate to be lower, you need to ensure that the audit log in Windows is large enough to capture the events.
15. Click **Apply** then **OK** to exit.

If Windows auditing is not configured properly, you will see warnings on the Compliance Standard Target Association page on the Cloud Control user interface. This is the same page where you associated your Real-time Monitoring compliance Standards to your targets.

Verifying Auditing Is Configured Properly

To verify that the host records login and logout events, follow these steps:

1. Log out of the host and then log back into the host.
2. From **Start**, select **Settings**, then **Control Panel**, then **Administrative Tools**, and finally **Event Viewer**.
3. Select **Security Log** and choose **Filter** from the **View** menu. Select **Security for the Event Source** and **Logon/Logoff** for the Category fields.
4. Click **OK**.

The Event Viewer should have the activity recorded as Event 528.

SubinacL External Requirements

As mentioned earlier, the agent will send warnings to the server when audit settings are not set properly. It, however, can only do this if the windows feature SUBINACL is installed. If this feature is not installed on the host, a warning will be sent to the server saying that the agent cannot detect whether audit settings are correct. This warning will be visible from the Compliance Standard Associate Targets page.

You can specify the absolute path to the directory that contain subinacL by setting the following property in the `AGENT_INST/sysman/config/nmxc.properties` file:

```
nmxcf.subinacL_dir=
```

SubInACL is available for download from Microsoft's Web site.

Preparing To Monitor Solaris Hosts

Real-time monitoring on Solaris systems utilizes the Solaris audit system which is part of the Solaris Basic Security Model (BSM). BSM auditing allows system administrators to monitor events and to detect user account logins and logouts as well as file changes.

Verify that BSM auditing is enabled by running the following command with root privilege:

```
/usr/sbin/auditconfig -getcond
```

You should see the following output:

```
audit condition = auditing
```

If the output is different from the above, it means the BSM auditing needs to be enabled through different methods in different Solaris releases.

Enabling BSM Auditing

You can enable BSM auditing using the steps below for each of the following environments.

Enabling BSM Auditing Using Solaris Versions 9 and 10

To enable BSM auditing, you can use the following command with root privilege:

```
/etc/security/bsmconv
```

See the Solaris BSM Auditing manuals for additional details on setting up BSM auditing.

If auditing is already enabled on the server, simply verify that the audit system configuration matches the configurations detailed below.

The audit file can be configured to include specific events. The */etc/security/audit_control* file controls which events will be included in the audit file. This section summarizes the configuration; for further details, refer to the Sun Product Online Documentation site.

For monitoring entity types OS FILE (file changes) and OS USER (user logins/logouts), the flags line in the file */etc/security/audit_control* should be set as follows:

```
flags: +fw,+fc,+fd,+fm,+fr,+lo
```

This configuration enables success/fail auditing for file writes (fw), file creates (fc), file deletes (fd), file attribute modifies (fm), file reads (fr) and login/logout events (lo); where '+' means to only log successful events.

If you are interested in logging the failed events as well, remove the "+" sign before each event in the flag.



Note:

Installing BSM on an existing host has the requirement that the host is rebooted.

Auditing Users: The *audit_user* file controls which users are being audited. The settings in this file are for specific users and override the settings in the *audit_control* file, which applies to all users.

Audit Logs and Disk Space: The *audit_control* file uses entries to control where the audit logs are stored and the maximum amount of disk space used by the audit system. The minimum requirement for file monitoring is approximately 10 minutes worth of data stored on the hard drive or the configured reporting interval time.

Enabling BSM Auditing Using Solaris 11

Auditing is enabled by default on Solaris 11, but only user login/logout events are monitored by default. For monitoring both the OS File change events and OS USER logins/logout events, you can execute the following command with root privilege:

```
/usr/sbin/auditconfig -setflags fw,fd,fc,fm,fr,lo
```

The configuration flags have the same meaning as defined in the last section.

 Note:

This configuration will not affect the existing sessions in which users already log into the host, so you must terminate all the existing sessions and then re-login or simply reboot the machine to ensure this change takes effect.

As the *bsmconv* command has been removed on Solaris 11, you can use the following command to enable the auditing feature, if needed:

```
audit -s
```

Managing Audit Log Files

Cloud Control Real-time Monitoring only reads the audit logs; it does not delete the logs. This might flood the system with log files and prevent it from logging additional events. To manage and delete old audit events while maintaining minimum monitoring requirements, follow these steps:

1. The auditing policy can be set to automatically drop new events (keeping only a count of the dropped events) rather than suspending all processes by running the following command:

```
# auditconfig -setpolicy cnt
```
2. Run the following command to force the audit daemon to close the current audit log file and use a new log file:

```
/usr/sbin/audit -n
```
3. Run the following command to merge all existing closed auditing log files into a single file with an extension of *.trash* and then delete the files:

```
/usr/sbin/auditreduce -D trash
```
4. Create a cron job to periodically run the commands in Step 2 and 3 above. The frequency at which these two commands are run can be adjusted based on the anticipated event volume and the amount of disk space allocated to auditing. The only requirement is that the time between the `audit -s` command and the `auditreduce -D trash` command is at least 15 minutes or twice the reporting interval if that is changed.

Preparing to Monitor AIX Hosts

Real-time monitoring on AIX systems utilizes the underlying AIX audit subsystem provided by the OS. IBM AIX 5.3 and 6.1 are the only currently supported versions.

Installation Prerequisite for AIX 5.3

Before using Real-time monitoring on AIX 5.3 hosts, ensure that you are using AIX 5.3 5300-08 service pack or higher. This maintenance package is available from IBM.

Administering AIX Auditing

The AIX auditing subsystem allows an administrator to record security-relevant information, such as User Logins, Logouts, and file changes, for analysis against existing security policies and detection of security violations.

Setting up auditing involves modification of the existing auditing configuration files. To set up auditing, follow these steps:

1. Log into the AIX machine as the root user.
2. Open a terminal window and change directory to `/etc/security/audit`
3. Open the config file in `vi`.
4. Locate the following sections, and update or add the listed values:

```
start:
    binmode = off
    streammode = on

stream:
    cmds = /etc/security/audit/ccstream

classes:
...
    filewatch = PROC_Create,PROC_Delete,FILE_Open,FILE_Write,FILE_Close,FILE_
Link,FILE_Unlink,FILE_Rename,FILE_Own-
er,FILE_Mode,FILE_Fchmod,FILE_Fchown,FS_Chdir,FS_Fchdir,FS_Chroot,FS_Mkdir,FS_Rmd
ir,FILE_Symlink,FILE_Dupfd,FILE_Mknod,FILE_Utimes

users:
    root = filewatch
    default = filewatch
```

Note:

- In this case, default refers to all users that are not root. Further note that the last line of the config file should be a blank line.
- Each parameter (binmode, streammode, filewatch, root, and default) must have a tab in front of them. You can verify that the audit system has used all variables properly by using the *audit query* command. Make sure the filewatch property appears in the output.

5. Save your modifications and exit `vi`.
6. In the same directory (`/etc/security/audit/`) open the file `streamcmds` in `vi`.
7. Clear all text from the file. The default configuration for this file is not necessary, as the File Monitoring agent module (`nmxcf` process) will operate as a direct audit reader. Clearing the file helps to reduce CPU usage and improve overall auditing performance.
8. Save the file and exit `vi`.

9. At the terminal prompt, enter the following command to initialize Auditing at system startup:

```
mkitab "audit:2:once:/usr/sbin/audit start"
```
10. At the prompt, restart audit using the command `/usr/sbin/audit shutdown` and `/usr/sbin/audit start` or directly reboot the host to make the auditing effective.
11. At the prompt, use the command `audit query` to view the configuration the audit system is using. Ensure that the properties are set correctly and that the required settings for filewatch are set.

Verifying AIX System Log Files for the OS User Monitoring Module

The OS User monitoring module relies on the following system log files:

- `/etc/security/failedlogin`
- `/var/adm/wtmp`
- `/var/adm/sulog`

Be sure the log files exist before running the OS User monitoring module on an AIX host. If any of the log files is missing, refer to the AIX System documentation for more information about how to generate it.

Preparing To Monitor the Oracle Database

This section describes the steps involved in setting up auditing within an Oracle database. If you are going to monitor an Oracle database with any of the Oracle entity types, you will need to perform these steps before events will be captured.

Before configuring auditing it is suggested you review the Auditing Database Use section of the *Oracle Database Administrator's Guide*. This document provides an overview of Oracle's auditing functionality, as well as basic concepts and guidelines for auditing configurations. Note that this document does not cover all details of configuring and fine tuning the Oracle audit system. Instead, this document serves as an example of the basic steps involved to configure the Oracle audit system to enable Real-time monitoring through Real-time monitoring rules.

Setting Auditing User Privileges

When you create a Real-time Monitoring Compliance Standard Rule to monitor an Oracle instance target, the agent will read the audit trail to perform its monitoring.

Real-time monitoring for Oracle entity types requires the audit trail to be stored in the database as opposed to a file. To verify if a setting is correct, follow these steps:

1. In Cloud Control, go to the target home page for the Oracle Database target for which you want to enable Real-time Monitoring.
2. From the **Administration** menu, select **Initialization Parameters**.
3. Log in to the database as a sys user, connecting as SYSDBA.
4. Find the parameter `audit_trail` and ensure it is set to DB. If not, this parameter needs to be changed in the Oracle Database.
5. This change will require a restart of the database.

Specifying Audit Options

Through SQL plus, an Oracle DBA can use audit and noaudit statements to configure audit options for the database. The audit statement allows you to set audit options at three levels:

Table 53-1 Audit Options Table

Level	Effect
Statement	Audits specific SQL statements or groups of statements that affect a particular type of database object. For example, AUDIT TABLE audits the CREATE TABLE, TRUNCATE TABLE, COMMENT ON TABLE, and DELETE [FROM] TABLE statements.
Privilege	Audits SQL statements that are executed under the umbrella of a specified system privilege. For Example, AUDIT CREATE ANY TRIGGER audits statements issued using the CREATE ANY TRIGGER system privilege.
Object	Audits specific statements on specific objects, such as ALTER TABLE on the employee table

To use the audit statement to set statement and privilege auditing options a DBA must be assigned AUDIT SYSTEM privileges. To use the audit statement to set object audit options, the DBA must own the object to be audited or be assigned the AUDIT ANY privilege within Oracle. Privilege assignments are covered in the following section.

Audit statements that set statement and privilege audit options can also include a BY clause to supply a list of specific users or application proxies to audit, and thus limit the scope of the statement and privilege audit options.

Some examples of audit statements can be seen below. Feel free to use these as a basis for the audit settings you specify within your database. Once all audit settings are in place you can create application policies, using the Oracle (SQL Trace) agent module with which to monitor the Oracle database instance.

Statement Audit Options (User sessions)

The following statement audits user sessions of users Bill and Lori.

```
AUDIT SESSION BY scott, lori;
```

Privilege Audit Options

The following statement audits all successful and unsuccessful uses of the DELETE ANY TABLE system privilege:

```
AUDIT DELETE ANY TABLE BY ACCESS WHENEVER NOT SUCCESSFUL;
```

Object Audit Options

The following statement audits all successful SELECT, INSERT, and DELETE statements on the dept table owned by user jward:

```
AUDIT SELECT, INSERT, DELETE ON jward.dept BY ACCESS WHENEVER SUCCESSFUL;
```

Example Oracle Audit Monitor Configurations

The following command audits all basic statements. Extra statements are not audited.

Audit all by access;

The following statement audits all extra statements:

```
audit ALTER SEQUENCE, ALTER TABLE, DELETE TABLE, EXECUTE PROCEDURE, GRANT  
DIRECTORY, GRANT PROCEDURE, GRANT SEQUENCE, GRANT TABLE, GRANT TYPE, IN-  
SERT TABLE, LOCK TABLE, UPDATE TABLE by access;
```

The following command displays audit settings for statements:

```
SELECT * FROM DBA_STMT_AUDIT_OPTS;
```

Once you have specified your audit configuration you can then create real-time monitoring rules from the Cloud Control Server that uses the Oracle Database entity types.

Setting Up Change Request Management Integration

This section explains how to install and configure integration with a Change Management Server and to be able to determine whether changes that occur are authorized automatically.

BMC Remedy Action Request System 7.1 Integration

Remedy ARS 7.1 is a supported Change Management system for automatic reconciliation of observations. The following steps outline how to setup Remedy and also configure the integration with Cloud Control.

Installing and Customizing Remedy ARS

Follow these steps to install and customize Remedy ARS 7.1.

1. Install Remedy ARS 7.1. Ensure the following components are all installed and properly licensed:

ARS 7.1.00 Patch 011

Midtier 7.1.00 Patch 011

Flashboard Server 7.0.03

Assignment Engine 7.1

Asset Management 7.0.03*

CMDB 2.1.00 Patch 4

CMDB Extension Loader

Approval Server 7.1

Change Management Server 7.0.03 Patch 008*

Problem Management Server 7.0.03*

Incident Management Server 7.0.0*3

User Client

Administrator Client

These packages all come with the IT Service Management Pack. Oracle provides example customizations for the Remedy under ITSM 7.0.03 Patch 008 environ-

ment. For different versions, the customizations may need to be adjusted to account for changes in the version of Remedy.

2. Install the Cloud Control EMCLI_Client on the same host on which Remedy is installed. This will need to be able to communicate to your Cloud Control Server.
 - a. Log in to the Enterprise Manager console.
 - b. Choose **Setup**, then select **Command Line Interface** from the **My Preferences** menu.
 - c. Click **Download the EM CLI kit to your workstation** and download the jar to your Remedy server.
 - d. Follow the steps given on the page to install the EMCLI client on the Remedy server.

3. Get the latest version of the Change Request Management connector self-update package. Also acquire the latest version of the example Remedy ARS customizations for Cloud Control version 12c.

These definition files provide a guideline of customizations that must be made in your environment for the integration. These customization files assume a fresh install of Remedy ARS. When integrating with a production instance of Remedy, care should be taken to make sure these customizations are compatible with any previous customizations that have been made to the Remedy instance.

- ActiveLinks_Customization.def
- Forms_Customization.def
- Menus_Customization.def
- Webservices_Customization.def

To get these definition files, in the Enterprise Manager Self Update user interface, export the connector. The definition files are inside this connector package.

4. Install the four definition files (.DEF) files in the running Remedy environment by completing these steps:
 - a. Log into the Remedy Administrator tool.
 - b. Select the **Remedy** instance from the hierarchy on the left.
 - c. From the **Tools** menu, select **Import Definitions**, then select **From Definition File...**
 - d. Select the definition file to import from the list above.
 - e. Check the box labeled **Replace Objects on the Destination Server**.
 - f. Choose the drop down option **Replace With New Type**.
 - g. Click **Import**.
 - h. You should not encounter any errors during this process. At the end of import there should be an Import Complete message.
 - i. When done, repeat for the rest of the customization files.
5. Customize Web Services.
 - a. Log into Remedy Administrator tool.
 - b. Select **Webservices**, then select the webservice **EMCCC_GetCR**. Right click, then select **Open**.

- c. Select the **WSDL** tab.
 - d. In the input on top, modify the `midtier_server` and `servername` values in the **WSDL Handler URL**.
 - e. If `midtier` is on localhost, you can enter localhost right after `http://`.
 - f. If the `midtier` uses port 80, you can omit the port, otherwise include the port after the server name.
 - g. For the `servername` after "`public/`", enter the name of the Remedy server.
 - h. Click **View**.
 - i. You should see an XML representing the webservice WSDL file for the web-service.
 - j. If you see an error, check the `midtier_server` name, port, or `servername`. Also, you can try adding/removing the domain part of the `servername`. Another possible issue occurs when the `midtier` password set in Remedy's System > General > Serverinfo > Connection Settings may not be set correctly. Be sure to check this also if the WSDL XML is not returned.
 - k. If you see the XML content after clicking View, then close this window and save the changes.
 - l. Repeat all above steps with the webservices `EMCCC_PublishCSData` and `EMCCC_UpdateChangeRequest`.
6. Customize Active Links.
- a. Log in to Remedy Administrator tool.
 - b. Select active links and then select the active link **EMCCCC_ApprovedCR**. Right click, then select **Open**.
 - c. Click the **If Action** tab.
 - d. Click the Current Action **Run Process** at the end of the list of actions.
 - e. In the Command Line field, change the path to `emcli.bat` to match that of where you installed the `emcli` on the local host.
7. Create a user in Remedy that will be used for creating requests that will be used for automatic observation reconciliation:
- a. Log in to BMC Remedy User Client as an administrative user.
 - b. Click **Application Administration Console** on the User Client Home Page.
 - c. Click **Create** for each step 1 through 4 in this wizard.
 - d. When adding the person, add the support group under the Support Groups tab.
 - e. Under the Support Groups Tab, select sub tab **Support Group Functional Roles**.
 - f. Add Support groups with functional role of *Infrastructure Change Management*. Without this, you will not be able to create change requests as the Infrastructure Change Manager fields support group will not have values.
 - g. Go to AR System Administrator Console.
 - h. On the left side bar, select **Application**, then **Users/Groups/Roles**, then **Select Users**.
 - i. This will load the user search page. Click **Search** at the top right.

- j. Double-click the newly created user above to bring up the user form.
- k. Click the down arrow next to "Group List" field and select **Infrastructure Change Master**.
- l. Repeat the previous step and add the following Groups to this user as well.
Infrastructure Change Submit
Infrastructure Change User
Infrastructure Change Viewer
- m. Save the changes to this user by clicking the **Save** button in the upper right hand corner of the window.

Adding the Connector to Cloud Control

Follow these steps to add the connector to Cloud Control.

1. Add the Change Management Connector to Cloud Control.
 - a. Log into Enterprise Manager as an Administrative user that has privileges to create a connector.
 - b. From the **Setup** menu, select **Provisioning and Patching**, then choose **Software Library**.
 - c. Click **Actions**, then select **Administration**.
 - d. Click **Add**.
 - e. Provide a name, such as "self update swlib".
 - f. Provide a location where the swlib files will be located on the Cloud Control server. This can be anywhere, but must be a path that the Cloud Control user can access. You must put the full absolute path in this input.
 - g. This process will take several minutes to complete.
 - h. Locate the connector self-update package file.

The connectors jar can be downloaded from the Cloud Control store to EM@Customer using the Self Update console, and can be exported to any local directory using the export functionality of Self Update.
 - i. Run: `emcli import_update -file=<full path>/connector.zip -omslocal` (where *connector.zip* is an example name of the self update package)
 - j. If you have errors with the previous step, make sure the user you run emcli as has permissions to access this directory and file. Also, be sure you are using absolute path for the *-file* switch.
 - k. When successful, you will receive the following message:

Operation completed successfully. Update has been uploaded to Cloud Control. Please use the Self Update Home to manage this update.
 - l. Log into the Enterprise Manager console.
 - m. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
 - n. Find the type "Management Connector" and click the link "1" under "Downloaded Updates" for this entry.
 - o. Select the Connector from the table and click **Apply**.

2. Create a Change Management Connector instance.
 - a. Log in into Enterprise Manager console.
 - b. From the **Setup** menu, select **Extensibility**, then select **Management Connectors**.
 - c. Select "Remedy Change Management Connector" from the drop-down after "Create Connector", then click **Go**.
 - d. Provide a name and description for the connector. This name is used to choose the connector when creating a Real-time Monitoring Compliance Standard Rule.
 - e. After returning to the management connector listing page, select the newly added row, then click **Configure**.
 - f. Under the Web Service End Points label, change the [servername] and [port] to match that of your Remedy instance Web Services. The values you put here will be similar to what you configured in the Web Services step earlier in these instructions.
 - g. Enter the Remedy username and password you are using for the connector integration.
 - h. Enter the locale ('en', for example).
 - i. Enter the time zone offset of the remedy server from UTC, ('-08:00', for example).
 - j. Enter the Change ID to use as a test. This should be a valid Change Request ID currently existing in Remedy that is used to test the connectivity between Cloud Control and Remedy.

Using Automatic Reconciliation Rules

Once Remedy is customized and the Cloud Control connector is configured, to utilize the automatic reconciliation features you need to create Real-time Monitoring Rules that are configured to use automatic reconciliation. Use the following steps:

1. Create a Real-time monitoring Rule:
 - a. Follow the normal steps to create a Real-time monitoring Rule.
 - b. On the Settings page, choose **Authorized Observations Automatically** using Change Request Management System. This configures Cloud Control to use this change request from Remedy for reconciliation of Real-time Observations that are detected.
 - c. Select the connector from the drop-down.
 - d. Click to annotate change requests with authorized observations check box.
 - e. Continue to save the rule after this. The Real-time Monitoring Rule can be used like any other Real-time Monitoring rule. The integration with a new Change Management server will not begin until at least one Real-time Monitoring Standard with a rule using Automatic Reconciliation is associated to a target. Create a Compliance Standard, add this rule to the Compliance Standard, and associate this compliance standard to one or more targets.

The configuration of rules is discussed in more detail in the [Managing Compliance](#) section.

Creating Change Requests for Upcoming Changes

Now that integration is set up and Real-time monitoring rules have been created, Change Requests can be created by Remedy users in the Remedy interface. These Change Requests will be compared to observations that occur to automatically determine if these observations are from actions that were authorized by change requests or not.

To make this correlation, some new fields that have been added to the Change Request form must be filled out by the change request filer. Not all fields are required; correlation only occurs on the fields that are present in the Change Request.

For instance, the following fields have been added to the Change Request form under the Oracle Enterprise Manager Integration tab:

- **Connector:** Choose the Cloud Control connector this Change Request will use to integrate with Cloud Control.
- **Hostname:** the hostname(s) this change request is for. These are the hosts that this change request is specifying someone needs to make changes to. An empty value in this field indicates that all hosts will be correlated to this change request.
- **Target User List:** the user name(s) this change request is for based on target users. These are the target users you expect to log in to the target to make a change. An empty value in this field means that all users on the target will be correlated to this change request.
- **Target Type:** the target type this change request is against. An empty value in this field means that any target type will be correlated to this change request.
- **Target:** The target this change request is specifically for. An empty value in this field indicates that any target will be correlated to this change request.
- **Facet:** The facet this change request is specifically for. An empty value in this field indicates that all facets on the above target type and target will be correlated for this change request.

When creating a change request that you want to use to authorize changes detected by Real-time monitoring rules, follow these steps in addition to whatever requirements your organization implements for creation of Change Requests:

1. Under the Dates tab of the Change Request form, fill out the Scheduled Start date and Scheduled End Date. These are the date ranges the request is valid for reconciliation. If an action occurs outside this time, it is marked as unauthorized by the Real-time Monitoring feature.
2. Select the Oracle Enterprise Manager **Integration** tab.
3. Select the Cloud Control connector from the drop-down list.
4. Optionally select values for the five reconciliation criteria as described above: Hostname, Target User List, Target type, Target and Facet. The last three -- Target Type, Target, and Facet -- will be Choice lists based on content in Real-time Monitoring Rules that have been created in Cloud Control that belong to Compliance Standards which are associated to targets. You can add multiple values separated by commas.

**Note:**

This form can be customized in Remedy to look differently. The example form elements from the customizations loaded earlier are only examples.

5. Change the auditable status to True. This configures Remedy to allow Cloud Control to use this change request for reconciliation of Real-time Observations that are detected.
6. Save the change request.
7. A popup displays, notifying you that active links will send the content to Cloud Control. You will see a DOS command window open and then close.

Overview of Reconciliation Functionality

After creating a change request that references a target and/or facet that is being monitored by Real-time Monitoring rules, any observations that happen against that rule will be correlated to all open and matching change requests.

When the observation arrives at the Cloud Control server, all open change requests that were active (based on Scheduled Start/Stop time) and have matching correlation criteria from the Cloud Control Integration tab will be evaluated. If any change request exists that matches the criteria of the observation, this observation will be marked with an "authorized" audit status. If the annotation check box was checked in the Rule configuration, details of these authorized observations will be put into a table in the Enterprise Manager Integration tab of the Remedy Change Request.

If no open change requests can be correlated to the observation and the rule was configured to use automatic reconciliation, then this observation is set to an Unauthorized audit status. The Observation bundle to which this observation belonged will be in violation and results in a Cloud Control event being created. This event can further be used through creation of a Cloud Control Event Rule.

An observations audit status can be seen whenever looking at observation details either by selecting Compliance, then Real-time Observations, then Observation Search, or either of the Browse By screens. A user with the proper role can also override the audit status for individual observations from these pages.

Any bundles that are in violation because they contain unauthorized observations will be reflected as violations in the Compliance Results page. These violations cause the compliance score skew lower. If these violations are cleared, the score becomes higher; however, the history of these audit status changes will be retained for the given observation.

Overview of the Repository Views Related to Real-time Monitoring Features

The following views exist to allow access to Real-time Monitoring data.

View: mgmt\$ccc_all_observations

Description: This view returns all observations that have occurred. Any query against this view should ensure that filtering is done on appropriate fields with *action_time* being the first to take advantage of partitions.

Fields:

Field	Description
OBSERVATION_ID	Unique ID given to the observation when detected by the agent
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
ENTITY_TYPE	Entity type of the entity that had an action against it
ACTION	Action that was observed
ACTION_TIME	Time the action occurred
USER_TYPE	Type of user that performed the action (for example, OS user versus DB user)
USER_PERFORMING_ACTION	Name of the user that performed the action
ORIGINAL_USER_NAME	Previous user name in the case of a SU/SUDO action (only applicable to some entity types)
AFFECTED_ENTITY_NAME	Name of the entity that was affected by this action (file name, and so on)
AFFECTED_ENTITY_PREVIOUS_NAME	Name of the entity prior to the action. For instance for file rename actions, this would be the old file name.
SOURCE_HOST_IP	Source IP of a connection when an action comes from another host (only applicable to some entity types)
ACTION_PROCESS_ID	PID of the process that performed the action (only applicable to some entity types)
ACTION_PROCESS_NAME	Name of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_ID	PID of the parent process of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_NAME	Name of the parent process of the process that performed the action (only applicable to some entity types)
ENTITY_PREVIOUS_VALUE	Previous value of the entity (only applicable to some entity types)
ENTITY_NEW_VALUE	New value of the entity (only applicable to some entity types)
FILE_ENTITY_PREVIOUS_MD5_HASH	Previous MD5 hash value of the entity (only applicable to some entity types)
FILE_ENTITY_NEW_MD5_HASH	New MD5 hash value of the entity (only applicable to some entity types)
AUDIT_STATUS	Current audit status of the observation (unaudited, authorized, unauthorized, and so on)

Field	Description
AUDIT_STATUS_SET_DATE	Date the most recent audit status was set
AUDIT_STATUS_SET_BY_USER	User who set the most recent audit status

View: mgmt\$ccc_all_obs_bundles

Description: This view returns a summary of all observations bundles. Any query against this view should ensure that filtering is done on appropriate fields with *bundle_start_time* being the first to take advantage of partitions.

Fields:

Field	Description
BUNDLE_ID	Bundle to which this observation belongs based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
RULE_NAME	Name of the Real-time Monitoring Compliance Standard Rule
ENTITY_TYPE	Entity type of the entity that had an action against it
USER_PERFORMING_ACTION	Name of the user that performed the action
BUNDLE_IN_VIOLATION	Boolean value if the bundle currently is in violation. This means at least one observation in the bundle is unauthorized. True indicates the bundle is in violation.
BUNDLE_START_TIME	Date of the first observation in this bundle
BUNDLE_CLOSE_TIME	Date when this bundle was closed
BUNDLE_CLOSE_REASON	Explanation of why this bundle was closed
DISTINCT_OBS_COUNT	Total number of observations in this bundle
AUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently authorized
UNAUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently unauthorized
UNAUTH_CLEARED_OBS_COUNT	Number of observations in this bundle that are currently cleared (that were at one point unauthorized)
UNAUDITED_OBS_COUNT	Number of observations in this bundle that are currently unaudited. They have not been evaluated manually or with Change Management integration to determine audit status.

View: mgmt\$ccc_all_violations

Description: This view returns all real-time monitoring violations caused by an observation bundle having at least one unauthorized observation in it.

Fields:

Field	Description
ROOT_CS_ID	Root Compliance Standard GUID. This is used for internal representation of the violation context.
RQS_ID	Runtime compliance standard GUID. This is used for internal representation of the violation context.
RULE_ID	Rule GUID. Internal ID of the rule having a violation.
TARGET_ID	Target GUID. Internal ID of the target having a violation.
ROOT_TARGET_ID	Root Target GUID. Internal ID of target hierarchy.
RULE_TYPE	Type of rule (Repository, Weblogic Server Signature, Real-time Monitoring)
SEVERITY	Severity Level of the rule (Info, Warning, Critical)
BUNDLE_ID	Internal ID of the Observation Bundle that is in violation. This observation bundle has one or more unauthorized observations in it
BUNDLE_START_TIME	Time the Observation Bundle started
BUNDLE_CLOSE_TIME	Time the Observation Bundle closed
TARGET_TYPE	Target Type of the Observation Bundle and all observations inside that bundle.
ENTITY_TYPE	Entity Type of the Observation Bundle and all observations inside that bundle.
USER_NAME	User name that performed the actions in this bundle
AUTHORIZED_OBS_COUNT	Number of Authorized observations in the observation bundle involved in this violation.
UNAUTHORIZED_OBS_COUNT	Number of Unauthorized observations in the observation bundle involved in this violation.
UNAUDITED_OBS_COUNT	Number of unaudited observations in the observation bundle involved in this violation.
RULE_NAME	Rule Name this violation is against.
COMPLIANCE_STANDARD_NAME	Compliance Standard Name this violation is against.
TARGET	Target Name this violation is against.

View: mgmt\$compliant_targets

Description: This view returns all evaluation and violation details for all targets. This is the same data that is shown in the Compliance Summary dashboard regions for targets.

Fields:

Field	Description
TARGET_ID	Internal representation of the Target

Field	Description
TARGET_NAME	Name of the Target
TARGET_TYPE	Target Type of the Target
TARGET_TYPE_INAME	Internal representation of the Target Type
CRIT_EVALS	Number of Critical-level Evaluations
WARN_EVALS	Number of Warning-level Evaluations
COMPLIANT_EVALS	Number of Compliant Evaluations
CRIT_VIOLATIONS	Number of Critical-level Violations
WARN_VIOLATIONS	Number of Warning-level Violations
MWARN_VIOLATIONS	Number of Minor Warning-level Violations
COMPLIANCE_SCORE	Current Compliance Score for the target

View: mgmt\$compliance_summary

Description: This view returns all evaluation and violation details for Compliance Standards and Frameworks. This is the same data that is shown in the Compliance Summary dashboard regions for Standards and Frameworks.

Fields:

Field	Description
ELEMENT_NAME	Display name of the Compliance Standard or Compliance Framework
ELEMENT_ID	Internal ID of the compliance standard or compliance framework
FRAMEWORK_ID	Internal ID of the Compliance Framework
CRIT_EVALS	Number of Critical-level Evaluations
WARN_EVALS	Number of Warning-level Evaluations
COMPLIANT_EVALS	Number of Compliant Evaluations
CRIT_VIOLATIONS	Number of Critical-level Violations
WARN_VIOLATIONS	Number of Warning-level Violations
MWARN_VIOLATIONS	Number of Minor Warning-level Violations
COMPLIANCE_SCORE	Current compliance score for the standard or framework
NON_COMPLIANT_SCORE	Current non-compliant score for the standard or framework
ELEMENT_TYPE	Type of element (1=Compliance Standard, 4=Compliance Framework)
AUTHOR	Author of the standard or framework
VERSION	Version of the standard or framework
ELEMENT_INAME	Internal representation of the standard or framework

View: mgmt\$compliance_trend

Description: This view returns the last 31 days compliance trend information for compliance frameworks and standards. This is the same data that is shown in the Compliance Summary dashboard trend regions for Standards and Frameworks.

Fields:

Field	Description
ELEMENT_ID	Internal ID representation of the standard or framework
FRAMEWORK_ID	Internal ID representation of the compliance framework
ELEMENT_NAME	Display name of the Compliance Standard or Compliance Framework
ELEMENT_INAME	Internal representation of the standard or framework
AVG_COMPLIANCE_SCORE	Average compliance score over last 31 days
DAILY_AVG_VIOLATIONS	Average number of violations per day over last 31 days
SNAPSHOT_TS	The snapshot timestamp
TOTAL_EVALS	Total evaluations over last 31 days
ELEMENT_TYPE	Type of element (1=Compliance Standard, 4=Compliance Framework)

Modifying Data Retention Periods

Real-time Monitoring features use partitioning and data retention configuration.

The following are the tables along with their default retention periods. When changing any retention periods, all tables related to Real-time monitoring must be changed to the same value to ensure that data is consistent across various features.



Note:

For more information about modifying data retention values, see the chapter "Maintaining and Troubleshooting the Management Repository" in the book *Oracle Enterprise Manager Administration*.

Table Name	Default Retention Period	Description
EM_CCC_WATCHDOG_ALERTS	366 Days	This table stores warnings from the agents when we detect that monitoring was not active.
EM_CCC_HISTORY_JOBEXEC	366 Days	This table stores history of all Enterprise Manager Jobs that are run as part of the Real-time Monitoring functionality.

Table Name	Default Retention Period	Description
EM_CCC_OBSERVATION	366 Days	This table stores each individual observation of a user action (for example, each file change, login/logout, process start/stop, each database object change, and so on).
EM_CCC_OBSGROUP	366 Days	This table stores information about how a single observation is related to a bundle based on the bundle settings set in the Real-time Monitoring Rule's user interface.
EM_CCC_OBS_GROUP_MAP	366 Days	This table stores the relationship between each single observation bundle and the target, rule, and standard that was monitoring for that observed action.
EM_CCC_HISTORY_OBS_STATUS	366 Days	This table stores the state change history for audit status (unaudited, unauthorized, authorized) for each observation.
EM_CCC_HA_OBS	366 Days	This table stores analytic summaries of counts of observations by hour and other attributes for reporting.
EM_CCC_HA_OBSGROUP	366 Days	This table stores analytic summaries of counts of observations bundles by hour and other attributes for reporting.
EM_CCC_FILEOBS_DIFF	366 Days	This table stores past file comparison for OS File based observations.
EM_CCC_AUTHOBS_CR_MAP	366 Days	This table stores the mapping between Change Management Request System change requests that were used to authorize an observation.
EM_CCC_CMPUBACTION	366	This table stores requests to publish data from EM server to an integrated Change Management Server using the connector.
EM_CCC_CMPUBACTION_DETAIL	366	This table stores additional details for requests to publish data from EM server to an integrated Change Management Server using the connector.

Real-time Monitoring Supported Platforms

The following tables display the various platforms that support Real-time monitoring. For all tables, an X indicates support for the listed action and NS indicates "Not Supported".

The following Operating System platform combinations are not supported at this time:

- Microsoft Windows -- IA64
- Any Linux -- IA64, PA-RISC, POWER

OS User Monitoring

The following table displays the platforms that support OS User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-2 OS User Monitoring

Actions to Monitor	Oracle/Redhat Linux					Windows					
	V4		V5		V6	XP		2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit
Telnet Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Telnet Logout (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Telnet Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Logout (Successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SSH Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
Console Login (successful)	X	X	X	X	X	X	X	X	X	X	X
Console Logout (successful)	X	X	X	X	X	X	X	X	X	X	X
Console Login (failed)	X	X	X	X	X	X	X	X	X	X	X
FTP Login (successful)	NS	NS	NS	X	X	NS	NS	NS	NS	NS	NS
FTP Logout (successful)	NS	NS	NS	X	X	NS	NS	NS	NS	NS	NS
FTP Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Login (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Logout (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SU Login (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SUDO (successful)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
SUDO (failed)	X	X	X	X	X	NS	NS	NS	NS	NS	NS
RDP Login (Successful)	NS	NS	NS	NS	NS	X	X	X	X	X	X
RDP Logout (Successful)	NS	NS	NS	NS	NS	X	X	X	X	X	X
RDP Login (failed)	NS	NS	NS	NS	NS	X	X	X	X	X	X

Table 53-3 OS User Monitoring

Actions to Monitor	SUSE Linux			Solaris				AIX			
	V10	V11		V9		V10		V11		V 5.3	V 6.1
	X86 32 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sparc	X86 64 bit	Sparc	X86 64	Sparc	POWER	POWER
Telnet Login (successful)	X	X	X	X	X	X	X	X	X	X	X
Telnet Logout (successful)	X	X	X	X	X	X	X	X	X	X	X
Telnet Login (failed)	X	X	X	X	X	X	X	X	X	X	X
SSH Login (successful)	X	X	X	X	X	X	X	X	X	X	X
SSH Logout (Successful)	X	X	X	X	X	X	X	X	X	X	X
SSH Login (failed)	X	X	X	X	X	X	X	X	X	X	X
Console Login (successful)	NS	X	X	X	X	X	X	X	X	NS	NS
Console Logout (successful)	NS	X	X	X	X	X	X	X	X	NS	NS
Console Login (failed)	NS	X	X	X	X	X	X	X	X	NS	NS
FTP Login (successful)	X	NS	NS	X	X	X	X	X	X	X	X
FTP Logout (successful)	NS	NS	NS	X	X	X	X	X	X	X	X
FTP Login (failed)	X	X	X	X	X	X	X	X	X	X	X
SU Login (successful)	X	X	X	X	X	X	X	X	X	X	X
SU Logout (successful)	NS	X	X	NS	NS	NS	NS	X	X	NS	NS
SU Login (failed)	X	X	X	X	X	X	X	X	X	X	X
SUDO (successful)	X	X	X	NS	NS	NS	NS	NS	NS	NS	NS
SUDO (failed)	X	X	X	NS	NS	NS	NS	NS	NS	NS	NS
RDP Login (Successful)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
RDP Logout (Successful)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
RDP Login (failed)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS

OS Process Monitoring

The following table displays the platforms that support OS User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-4 OS Process Monitoring

Actions to Monitor	Oracle/Redhat Linux			Windows					Solaris						
	V4	V5	V6	XP	2003 Server		2008 Server (R1 and R2)			V9	V10		V11		
	X86 32 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	X86 64 bit	SPARC	X86 64 bit	SPARC	X86 64 bit	SPARC
Process Start (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Process Stop (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Table 53-5 OS Process Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
Process Start (successful)	X	X	X	X	X
Process Stop (successful)	X	X	X	X	X

OS File Monitoring

For Linux v5, there are two possible ways monitoring can occur. Some actions to monitor below will work only on one or the other method. The two methods are to use the Loadable Kernel Module. Actions that are detectable ONLY with this method are annotated with "(KO)". The other option is to not use the loadable kernel module, which will result in using the Linux built-in audited method. The actions that can only be monitored using this method are annotated with "(non-KO)". The actions that have no annotation other than the check mark can be monitored using either approach.



Note:

Monitoring remote file systems on Unix-based platforms is not supported. Likewise, monitoring remote file systems on Windows platforms is also not supported.

When restoring a file from the Recycle Bin on the Microsoft Windows operating system, capturing the user that made the change is not available since that feature is not available from the Operating System.

When using the audited monitoring method on Linux operating systems, not the Oracle kernel audit module method, directory creations are reported as file creation. Additionally, file create activity will be reported as a file modification instead of create. These

are limitations of using the audited method of monitoring. If you use the Oracle kernel audit module approach for OS file monitoring on Linux, these limitations will not exist.

An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-6 OS File Monitoring

Actions to Monitor	Linux				Windows						Solaris						
	V4		V5		V6		XP		2003 Server		2008 Server (R1 and R2)		V9		V10		V11
	X86 32 bit	X86 32 bit	X86 64 bit	X86 64 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	X86 32 bit	Sp arc	X86 64 bit	Sp arc	X86 64 bit
File Read (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Delete (Successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Rename (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Create (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Content Modified (successful)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Modified without content change	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Modified (failed)	NS	X (No n-KO)	NS	X (No n-KO)	X	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
File Permission Change (successful)	NS	X (non-KO)	X (non-KO)	X (KO)	X	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X
File Ownership Change (successful)	NS	X (non-KO)	X (non-KO)	X (KO)	X	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X
File content modified (successful) Archive File	NS	X (non-KO)	X (non-KO)	X	X	X	X	X	X	X	X	X	X	X	X	X	X
File Read (failed)	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X
File Delete (failed)	NS	X (No n-KO)	X (No n-KO)	NS	NS	NS	NS	NS	NS	NS	NS	NS	X	X	X	X	X

Table 53-6 (Cont.) OS File Monitoring

Actions to Monitor	Linux				Windows						Solaris							
	V4		V5		V6		XP		2003 Serv-er		2008 Serv-er (R1 and R2)		V9		V10		V11	
	X86 32 bit	X86 32 bit	X86 64 bit	X8 6 32 bit	X86 64 bit	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 64 bit	Sp arc	X8 6 64 bit	Spa rc	X8 6 64 bit	Spa rc	
File Rename (failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X	
		(No n-KO)	(No n-KO)	(no n-KO)	(non - KO)													
File Create (failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X	
		(non - KO)	(no n-KO)	(no n-KO)	(non - KO)													
File Permis-sion Change (Failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X	
File Owner-ship Change (failed)	NS	X	X	X	X	NS	NS	NS	NS	NS	NS	X	X	X	X	X	X	

Table 53-7 OS File Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
File Read (successful)	X	X (KO)	X (KO)	X	X
File Delete (Successful)	X	X (KO)	X (KO)	X	X
File Rename (successful)	X	X	X	X	X
File Create (successful)	X	X	X	X	X
File Content Modified (successful)	X	X	X	X	X
File Modified without content change (successful)	X	X	X	X	X
File Modified (failed)	NS	NS	NS	X	X
File Permission Change (successful)	X	X (KO)	X	X	X
File Ownership Change (successful)	X	X (KO)	X	X	X
File content modified (successful) Archive File	X	X	X	X	X
File Read (failed)	NS	NS	NS	X	X
File Delete (failed)	NS	NS	NS	X	X
File Rename (failed)	NS	X (Non-KO)	X (Non-KO)	X	X

Table 53-7 (Cont.) OS File Monitoring (continued)

Actions to Monitor	SUSE Linux			AIX	
	V10	V11		V5.3	V6.1
	X86 32 bit	X86 32 bit	X86 64 bit	POWER	POWER
File Create (failed)	NS	NS	X (Non-KO)	X	X
File Permission Change (Failed)	NS	X (Non-KO)	X (Non-KO)	X	X
File Ownership Change (failed)	NS	X (Non-KO)	X (Non-KO)	X	X

OS Windows Registry Monitoring

The following table displays the platforms that support OS Windows Registry Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-8 OS Windows Registry Monitoring

Actions to Monitor	Windows					
	XP		2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit	X86 32 bit	X86 64 bit
Create Key (successful)	X	NS	X	X	X	X
Delete Key (successful)	X	NS	X	X	X	X
Create Value (successful)	X	NS	X	X	X	X
Modify Value (successful)	X	NS	X	X	X	X
Delete Value (successful)	X	NS	X	X	X	X
Create Key (failed)	X	NS	X	NS	NS	NS
Create Value (failed)	X	NS	X	NS	NS	NS
Modify Value (failed)	X	NS	X	NS	NS	NS
Delete value (failed)	X	NS	X	X	X	X

OS Windows Active Directory User Monitoring

The following table displays the platforms that support OS Windows Active Directory User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-9 OS Windows Active Directory User Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
User Create (successful)	X	X	X	X
User Delete (successful)	X	X	X	X

Table 53-9 (Cont.) OS Windows Active Directory User Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
User Attribute Modify (successful)	X	X	X	X

OS Windows Active Directory Computer Monitoring

The following table displays the platforms that support OS Windows Active Directory Computer Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-10 OS Windows Active Directory Computer Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
Computer Create (successful)	X	X	X	X
Computer Delete (successful)	X	X	X	X
Computer Attribute Modify (successful)	X	X	X	X

OS Windows Active Directory Group Monitoring

The following table displays the platforms that support OS Windows Active Directory Group Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-11 OS Windows Active Directory Group Monitoring

Actions to Monitor	Windows			
	2003 Server		2008 Server (R1 and R2)	
	X86 32 bit	X86 64 Bit	X86 32 bit	X86 64 bit
Group Create (successful)	X	X	X	X
Group Delete (successful)	X	X	X	X
Group Attribute Modify (successful)	X	X	X	X
Group Member Add (successful)	X	X	X	X
Group Member Delete (successful)	X	X	X	X

Oracle Database Table Monitoring

The following table displays the platforms that support Oracle Database Table Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-12 Oracle Database Table Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Insert (successful)	X	X	X	X
Select (successful)	X	X	X	X
Update (successful)	X	X	X	X
Delete (successful)	X	X	X	X
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Truncate (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Comment (successful)	X	X	X	X
Rename (successful)	X	X	X	X
Lock (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X
Flashback (successful)		X	X	X

Oracle Database View Monitoring

The following table displays the platforms that support Oracle Database View Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-13 Oracle Database View Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Insert (successful)	X	X	X	X
Select (successful)	X	X	X	X
Update (successful)	X	X	X	X
Delete (successful)	X	X	X	X
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Comment (successful)	X	X	X	X

Table 53-13 (Cont.) Oracle Database View Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Rename (successful)	X	X	X	X
Lock (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X
Flashback (successful)		X	X	X

Oracle Database Materialized View Monitoring

The following table displays the platforms that support Oracle Database Materialized View Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-14 Oracle Database Materialized View Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Insert (successful)	X	X	X	X
Select (successful)	X	X	X	X
Update (successful)	X	X	X	X
Delete (successful)	X	X	X	X
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Comment (successful)	X	X	X	X
Lock (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Index Monitoring

The following table displays the platforms that support Oracle Database Index Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-15 Oracle Database Index Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Analyze (successful)	NS	X	X	X

Oracle Database Sequence Monitoring

The following table displays the platforms that support Oracle Database Sequence Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-16 Oracle Database Sequence Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Select (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Procedure Monitoring

The following table displays the platforms that support Oracle Database Procedure Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-17 Oracle Database Procedure Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Execute (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X

Table 53-17 (Cont.) Oracle Database Procedure Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Function Monitoring

The following table displays the platforms that support Oracle Database Function Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-18 Oracle Database Function Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Execute (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Package Monitoring

The following table displays the platforms that support Oracle Database Package Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-19 Oracle Database Package Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Execute (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Library Monitoring

The following table displays the platforms that support Oracle Database Library Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-20 Oracle Database Library Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Execute (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X

Oracle Database Trigger Monitoring

The following table displays the platforms that support Oracle Database Trigger Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-21 Oracle Database Trigger Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X

Oracle Database Tablespace Monitoring

The following table displays the platforms that support Oracle Database Tablespace Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-22 Oracle Database Tablespace Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X

Oracle Database Cluster Monitoring

The following table displays the platforms that support Oracle Database Cluster Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-23 Oracle Database Cluster Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Truncate (successful)	X	X	X	X

Oracle Database Link Monitoring

The following table displays the platforms that support Oracle Database Link Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-24 Oracle Database Link Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X

Oracle Database Dimension Monitoring

The following table displays the platforms that support Oracle Database Dimension Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-25 Oracle Database Dimension Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X

Oracle Database Profile Monitoring

The following table displays the platforms that support Oracle Database Profile Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-26 Oracle Database Profile Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X

Oracle Database Public Link Monitoring

The following table displays the platforms that support Oracle Database Public Link Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-27 Oracle Database Public Link Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X

Oracle Database Public Synonym Monitoring

The following table displays the platforms that support Oracle Database Public Synonym Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-28 Oracle Database Public Synonym Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X

Oracle Database Synonym Monitoring

The following table displays the platforms that support Oracle Database Synonym Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-29 Oracle Database Synonym Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Drop (successful)	X	X	X	X

Oracle Database Type Monitoring

The following table displays the platforms that support Oracle Database Type Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-30 Oracle Database Type Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Create Type Body (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Drop Type Body (successful)	X	X	X	X
Grant (successful)	X	X	X	X
Revoke (successful)	X	X	X	X
Audit (successful)	X	X	X	X
NOAUDIT usage	X	X	X	X

Oracle Database Role Monitoring

The following table displays the platforms that support Oracle Database Role Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-31 Oracle Database Role Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Drop Type Body (successful)	X	X	X	X
Set (successful)	X	X	X	X

Oracle Database User Monitoring

The following table displays the platforms that support Oracle Database User Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-32 Oracle Database User Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
Create (successful)	X	X	X	X
Logon (successful)	X	X	X	X
Drop (successful)	X	X	X	X
Alter (successful)	X	X	X	X
Logoff	X	X	X	X
Grant Role (successful)	X	X	X	X
Revoke Role (successful)	X	X	X	X
System Grant (successful)	X	X	X	X
System Revoke (successful)	X	X	X	X

Oracle Database SQL Query Statement Monitoring

The following table displays the platforms that support Oracle Database SQL Query Statement Monitoring. An X indicates support for the listed action and NS indicates "Not Supported".

Table 53-33 Oracle Database SQL Query Statement Monitoring

Actions to Monitor	Oracle Database			
	9i	10g	11g	12g
SQL Query Output Changed	X	X	X	X

Overview of Change Activity Planner

Change Activity Planner (CAP) enables you to plan, manage, and monitor operations within your data center. These operations involve dependencies and coordination across teams and business owners, as well as multiple processes. Operations can include rollout of security patches every quarter, building new servers to meet a business demand, migration or consolidation of data centers, and rolling out compliance standards across an environment.

Using CAP, you can:

- Plan change activity, including setting start and end dates; and creating, assigning and tracking task status.
- Manage large numbers of tasks and targets, using automated task assignment and completion support.
- Use a dashboard where you can monitor your plans for potential delays and quickly evaluate overall plan status.
- Have task owners manage their tasks using a task-based dashboard showing task priorities and schedules.

This chapter covers the following:

- [Before Getting Started](#)
- [Creating a Change Activity Plan](#)
- [Operations on Change Activity Plans](#)
- [Managing a Change Activity Plan](#)
- [Viewing My Tasks](#)
- [Example of Using Change Activity Planner](#)

Before Getting Started

Before you start using Change Activity Planner, ensure you have the necessary privileges and familiarize yourself with the CAP terminology.

Change Activity Planner Roles and Privileges

For security purposes, Change Activity Planner provides the following roles:

- EM_CAP_ADMINISTRATOR
- EM_CAP_USER

The EM_CAP_ADMINISTRATOR role grants the CREATE_JOB, CREATE_CAP_PLAN, and BASIC_CAP_ACCESS privileges. To create, create-like, edit, and delete plans, you must have been granted the EM_CAP_ADMINISTRATOR role. The EM_CAP_USER role grants the BASIC_CAP_ACCESS privilege only. Without

the basic access privilege enabled, you will be unable to access Change Activity Plan and My Tasks features.

Change Activity Planner security model adheres to owner-based management. Object level privileges are currently not supported.

Change Activity Planner Terminology

The major concepts in Change Activity Planner are plans, task definitions, task groups, and tasks. This section explains these concepts in detail.

- [Plan](#)
- [Task Definition](#)
- [Task Group](#)
- [Task](#)

Plan

Plans introduce new changes into an organization, and specifies the start and end date for the required changes. It identifies the set of changes that are required, as well as the targets where the changes are needed.

To create plans you must have been granted the EM_CAP_ADMINISTRATOR role. (See 'Privileges and Roles' in the [Oracle Enterprise Manager Cloud Control Administrator's Guide](#) for additional information.) Once a plan is created, it is possible to monitor the progress of the plan. A plan is considered complete once all tasks that make up that plan are closed.

The following terms are used when working with plans.

Plan Term	Description
Activation in Progress	Plan is in the process of being activated. Tasks are being created. Plan cannot be edited.
Activated	Plan definition is complete and the plan has been activated. Activation means that tasks have been created and you can start managing the plan's progress.
Overdue Plan	An Active Plan which still has open tasks but has passed the Plan's scheduled end date.
Due Within One Week	Active Plan which is due within seven days.

The various statuses of a plan are:

Plan Status	Description
Active	Plan for which tasks have been created, and which is already in progress. A plan is also considered active if the plan has been activated, and CAP is in the process of creating tasks. While CAP is creating tasks, the plan's progress will be displayed as 'activation in progress'.
Completed	Plan is complete. All tasks are completed or canceled.

Plan Status	Description
Deactivated	Plan was activated but then canceled. All tasks associated with the plan are canceled. The plan can still be viewed in the Change Activity Plans page, and operations such as Create Like, Manage, and Delete are available. You can view the Manage Plan page for Deactivated Plans.
Definition in Progress	Plan has been partially defined and saved for later use. Edit the plan to continue or complete its definition, and to activate it.
Failed in Activation	Error occurred when the plan was being activated. The Comments and Audit Trail or the log files will have more details on the error.
Pending Activation	Plan has been scheduled to be activated at a future date.

Task Definition

The description of one step of the work to be accomplished to meet the plan's goal.

The action described in the task definition can be associated with a target type, or the target type can optionally be set to **None** which means that no target will be associated with the task. This is usually the case when the task defines a manual action, for example 'Get management approval to proceed'. In this case, the task owner will obtain management approval, acknowledge the approval by marking the task as completed, and then proceed with the follow on tasks.

If the task definition has a target type specified, then it can be associated with one or more targets. When a plan is activated, a task will be created for each task definition / target combination. If no target type is specified, then a single task will be created.

A task is given a start and scheduled end date and is assigned to an owner.

- Task Type: Specify the type of task to generate for users.
 - Custom (or Manual)

Allows you to specify the Action to associate with this task definition. Custom task types also support two modes of verification:

 - * Manual

Task owners will have to update the task status.
 - * Use evaluation results from the compliance standard rule to update the status. (Only available for task definitions with a target type specified.)

A compliance standard rule can be associated with the task definition. When the tasks are generated, and as work is done on the target, this rule is evaluated and the results are used to determine if the task is still open, or if the task can be automatically closed.
 - Patch Plan

This choice integrates with the My Oracle Support Patches & Updates support. A patch template is associated with the task definition, and when you assign your tasks, you can use the template to create and deploy a patch plan.

Both Patch Template and Compliance Rule require a target type selection on which to base the evaluation.
 - Job

Enables the plan creator to create tasks based on jobs. In turn the task owner submits a job for the task or associates the task with an existing job execution. Use either the My Tasks page or the Manage Plan Tasks tab for submitting jobs and associating tasks.

A library job must be associated with the task. Only those library jobs whose target type matches the task's target type will be available for selection.

There are a number of ways to specify when a task is completed.

- * Task is marked as complete when the job completes successfully. This is the default. Once a day CAP updates the job-based tasks based on the status of their associated job executions.

To immediately update the status of a job-based task, select the task in the My Tasks page or the Manage Plan Tasks tab.

- * Task is manually marked as complete by the task owner.
- * Task is marked as complete based on the evaluation of a compliance rule.

Note: When selecting a job for use with a CAP task definition, ensure that task owners have the appropriate access to the job and any appropriate job runs or executions. You can change the access to a job from the Job Library page. When creating a job, or editing the job later, select the Access tab. Add access for the appropriate users. Failure to grant the proper access will prevent task owners from viewing the job associated with their task.

— Deployment Procedure

Enables the plan creator to associate a configured deployment procedure with a task definition. A popup displays the available configured deployment procedures.

Note that only configured deployment procedures will be available for selection. To create a configured deployment procedure, select a procedure from the Procedure Library page. Click the **Launch** button, fill in the desired attributes, and then save the configuration.

The task owner can associate a CAP task with a procedure run which has been launched from the Procedure Library.

There are a number of ways to specify when a task is completed.

- * Task is marked as complete when the deployment procedure completes successfully. This is the default.
- * Task is manually marked as complete by the task owner.
- * Task is marked as complete based on the evaluation of a compliance rule.

Note: When selecting a deployment procedure for use with a CAP task definition, ensure that task owners have the appropriate access to the procedure definition and any procedure runs for the selected procedure. You can edit the access for procedure definitions from the Procedure Library page. You can edit the access for procedure runs from the Procedure Activity page. After launching a procedure, Edit Permissions for the procedure execution associated with the CAP task, and grant access to the appropriate administrators. Failure to grant the proper access will prevent task owners from viewing the procedure associated with their task.

- Task Status: The task definitions and task groups all have an associated status, indicated by an icon. The status is determined by the status of the tasks that make up the task definition or task group.

Example: If a task definition has four associated Tasks: Task 1 is Unacknowledged, Task 2 is Complete, Task 3 is In Progress and Task 4 is Canceled, then the task definition status will be In Progress.

- **Dependent Task (or Dependency):** Task that cannot be started until another task is complete.

When creating a task definition or task group, you can define a *single* dependency on another task definition or task group. When the plan is activated, tasks are generated. Task dependencies are determined based on the dependencies defined during plan creation. If a task is dependent on another task, it will be flagged as waiting until that task is complete. If a task is dependent on a task group, it will be flagged as waiting until all the tasks in the task group are complete.

Task Group

A folder containing one or more task definitions and can be used to represent a phase in the plan. When defining a plan, targets can be assigned to task groups. All task definitions within the task group will inherit the target assignments. You can set the task group to be dependent on another task definition or group, providing you a quick way to establish dependencies among sets of task definitions.

The task group is in progress if at least one task is in progress. The task group is completed when all tasks in the group are closed.

Task

When a plan is activated, tasks are created based on task definitions. If the task definition has targets, a task will be created for each licensed target.

The following terms are used when working with tasks:

Task Term	Description
Unassigned	Task that does not have an owner.
Open	Task with a status of Unacknowledged, Acknowledged, or In Progress.
Closed	Task with a status of Completed, Canceled, or Deactivated.
Waiting	Indicates the task is dependent on another task, and the task it depends on has not been closed. If a task is waiting, it is indicated in the Waiting column of the Tasks table. To see details, click on the task. The Tracking section will show information on the task's dependency.
Require Attention	The system checks tasks for certain conditions, and if they are found, the task is flagged as needing attention. The system checks for the following conditions: • Is Overdue - Indicates the task is still open and its scheduled end date has passed. • Is Due Within One Week - Indicates the task is due within seven days. • End Date Beyond Plan End Date - Indicates the task's scheduled end date is after the plan's scheduled end date. • Is Unassigned - Indicates the task does not have an owner. • Was Completed Out of Order - Indicates the task was set to Completed, while it was still waiting.

The task status shows where the task is in its lifecycle. The task statuses include:

Task Status	Description
Unacknowledged	When a plan is activated, all tasks start as acknowledged.
Acknowledged	Indicates a user has seen the task but has not yet started work on it.
In Progress	Indicates work has started on the task.
Completed	Indicates work on the task is complete. Tasks with automatic verification will enter this status when the system verifies the expected changes were made. For manual tasks, users will have to set the task status to Complete once they finish their work. Once a task is completed, the details become read only. Note: Completing a task cannot be undone.
Canceled	Indicates the work specified in the task is no longer needed, and will not be done. When a task is canceled, the details become read only. Note: Canceling a task cannot be undone.
Deactivated	If a plan is deactivated, all associated tasks are marked deactivated.

Creating a Change Activity Plan

When you create a plan, you specify the tasks required to complete the plan. When you define tasks, you:

- Associate targets with the tasks.
- Define dependencies between tasks. By setting dependencies, you provide a warning signal that a dependent task should not be started until the task it is depending on is completed or canceled either automatically or when marked as completed by the task owner. This provides the ability for task owners to detect when their tasks should be performed.
- Specify a task hierarchy that contains nested groups of tasks with related dependencies. Using task groups, you can organize your tasks and structure the flow of how your plan should be processed.
- Either auto-assign tasks on plan activation or manually assign tasks on plan activation.
- Specify custom instructions for completing the task or select a patch template that should be applied to the targets.
- Specify whether the system will automatically detect that the task is complete or whether the owner of the task will manually close the task.

To create a Change Activity Plan:

1. From the **Enterprise** menu select **Configuration**, then select **Change Activity Plans**.
2. On the Change Activity Plans page, click **Create**. This activates the Create Plan Wizard.
3. On the Create Plan: Properties page, provide general information:
 - Name
 - Target type - A plan's target type determines the types of targets that can be associated with the plan's tasks.

- **Priority** - When managing a plan, priority may be used to indicate the order of importance of ensuring task ownership is assigned and completion dates are met.
- **Description** - Provide pertinent information.
- **External Links** - You can also add Links to documentation pertinent to the plan. For example, a link could be one document that explains a specific patch set, a specific internal procedure or policy, the set of applications impacted by the task, and so on. Provide a name for the link and type the URL.

Click **Next**.

4. In the Create menu on the Create Plan: Task Definitions page, select either **Task Group** or **Task Definition** to create a hierarchy of task actions to perform in this plan.

See [Creating a Task Definition](#) and [Creating a Task Group](#) for detailed information.

Click **Next**.

5. On the Create Plan: Schedule page, you can:
 - Activate the plan immediately or at a time in the future.
 - Determine the duration or date by which the plan needs to be completed. Specify either the exact date by which the plan needs to be completed or the completion date relative to when the plan is activated.
 - Assign tasks automatically or manually. If you choose to manually assign tasks, assign the tasks after the plan has been activated. If you choose to assign tasks automatically, tasks will be assigned to target owners when the plan is activated. See [Assigning Tasks to the Group Owner](#) for information on how to create a Task Owner.



Note:

You can leave this Activate page blank until you are ready to activate the plan.

Click **Next**.

6. On the Create Plan: Review page, ensure all the information you have entered is as you intended. If updates are necessary, click the **Back** button and make the necessary changes.

To save your in-progress plan definition, go to the end of the flow (the Review step), and click the **Save and Exit** button.



Note:

The information on this page is not saved until you click **Save and Exit** or **Activate Plan** on the **Create Plan: Review** page.

Creating a Task Definition

When creating a task definition, provide the following information:

1. Type the name for the task. Task group and task definition names should be unique within a plan

Note: If you select an existing task, you can create a task before, after, in parallel, or inside the selected task. If you select **Before** or **After**, it specifies the order in which the tasks should be performed, and sets the task dependencies accordingly. If you decide that you do not like the order, you can change the order by using **Move** or **Set Dependency**.

2. Select the target type the task definition is associated with.

Use the **NONE** option when a task definition is not based on a target.

3. Select the task type. The options are:

- **Custom (or Manual)**

A custom task is one that you define. Custom tasks may include tasks you do manually, typically outside Enterprise Manager, but want to account for as part of the plan.

Provide enough information in the **Action** field for the task owner to complete the task, or add a URL to the Links section with instructions for the task.

In addition, specify how to verify that the task is completed.

- Select **Task owner will update the status** to indicate that the task owner will manually mark the task as completed when appropriate. Tasks can be marked as completed from the Manage Plan page or from the My Tasks page.
- Select **Use evaluation results from the compliance standard rule to update the status** to indicate that the task should be automatically marked as completed when a target is compliant with the selected compliance rule. Click the magnifying glass to access the **Search and Select Compliance Standard Rule** dialog box to select the compliance rule from the list of rules appropriate for the task target type.

- **Patch Plan**

When you select this option, use the Patch Template dialog box to select a patch template which specifies the list of patches that should be applied to the targets for the task.

The tasks for each target in this task definition will be closed automatically when the system detects that the patches have been deployed to the target.

You can add more instructions in the **Action** field, or add a URL in the Links section that leads to more information for the task owner.

- **Job**

This option enables the plan creator to create tasks based on jobs. Click the magnifying glass to access the **Search and Select Job** dialog box to select the job to be executed.

The task owner submits a job for the task or associates the task with an existing job execution. Use either the **My Tasks** page or the **Manage Plan Tasks** tab for submitting jobs and associating tasks.

A library job must be associated with the task. Only those library jobs whose target type matches the task's target type will be available for selection.

There are a number of ways to specify when a task is completed.

- Task is marked as complete when the job completes successfully. This is the default.

Once a day CAP updates the job-based tasks based on the status of their associated job executions. To immediately update the status of a job-based task, select it in the **My Tasks** page or the **Manage Plan Tasks** tab.
- Task is manually marked as complete by the task owner.
- Task is marked as complete based on the evaluation of a compliance rule.

Note: When selecting a job for use with a CAP task, ensure that task owners have the appropriate access to the job and any appropriate job runs or executions. You can change the access to a job from the Job Library page. When creating a job, or editing the job later, select the Access tab. Add access for the appropriate users. Failure to grant the proper access will prevent task owners from viewing the job associated with their task.

- **Deployment Procedure**

Enables the plan creator to associate a deployment procedure with a task definition. Click the magnifying glass to access the Search and Select Deployment Procedure dialog box and select the appropriate deployment procedure.

Note: Only configured deployment procedures will be available for selection. To create a configured deployment procedure, select a procedure from the Procedure Library page. Click the **Launch** button, fill in desired attributes, and then save the configuration.

The task owner can associate a CAP task with a procedure run which has been launched from the Procedure Library. Ensure that the task owners have permission to view the necessary procedure runs.

There are a number of ways to specify when a task is completed.

- Task is marked as complete when the deployment procedure completes successfully. This is the default.
- Task is manually marked as complete by the task owner.
- Task is marked as complete based on the evaluation of a compliance rule.

Note: When selecting a deployment procedure for use with a CAP task, ensure that task owners have the appropriate access to the procedure definition and any procedure runs for the selected procedure. You can edit the access for procedure definitions from the Procedure Library page. You can edit the access for procedure executions from the Procedure Activity page. After launching a procedure, Edit Permissions for the procedure execution associated with the CAP task, and grant permission to the appropriate administrators. Failure to grant the proper access will prevent task owners from viewing the procedure associated with their task.

4. Provide the action for this task definition, for example, schedule a blackout, back-up a database, and so on.

5. Select the verification method. The methods displayed are dependent on the type of task. Methods include:
 - **Task owner will manually update the status**
 - **Use status from job execution to update the status**
 - **Use status from procedure run to update the status**
 - **Use evaluation results from the compliance standard rule to update the status.** Select the compliance standard rule.

Patch plan deployments are automatically verified by the system.

6. Click **Add** to add links to documentation pertinent to the task definition. Provide a name for the link and type the URL.
7. Click **Next**.

Note: The information on this page is not saved until you click **Save and Exit** or **Activate Plan** on the **Create Plan: Review** page.

Setting Dependencies

When you add task definitions to the plan, using the Add After or Add Before options automatically set a task dependency to specify the order in which the tasks should be completed. If you want to change the order, use the Move option to move a single task before or after another task.

You can manually specify a dependency between task definitions to indicate whether a task definition should be completed before another one starts. Use **Set Dependency** to move a task and all the tasks after it to be performed after another task. To set a dependency, select a task definition and click **Set Dependency** located in the toolbar. Specify the task definition that should be completed before the selected one starts. This does not prevent the task owner from completing a task definition out of order, but the task owner will be able to see the dependency and will be warned that the task definition should wait until the dependency is completed.

Dependencies can only be created on task definitions that are defined at the same level of the task definition tree, and that appear above the selected task definition.

Moving Task Definitions

Task definitions can be moved to a different location in the tree by selecting a task definition and clicking the **Move** button. Note that this does not affect the order of task definition execution, unless by moving the task definition a dependency must be removed. Moving a task definition in or out of a task group can affect the targets added to the task definition if the target is inherited from the task group.

Note: The information on this page is not saved until you click **Save and Exit** on the **Create Plan: Review** page.

Adding Targets

Click **Add Targets** to select the targets associated with the task. To activate the plan, there must be at least one target associated with each task except for tasks with None as the target type.

If the target type for the plan is a system, tasks will only be created for targets in the system appropriate to the task definition.

Note the following:

- You can only add targets that correspond to the target type for the plan. If the target type for the plan is a system, tasks will only be created for targets appropriate to the task definition.

For example, if the plan target type is a Database System and the task definition is for a Listener, tasks will be created for any Listener targets associated with the Database Systems you select.

- Assigning targets to task groups

If you assign targets to a task group, the targets will also be added to all the tasks within the group.

- If a task specifies a target type, at least one target must be associated with it before the plan can be activated.
- To see which targets are added to a task or task group, select the task definition and click the **Edit Targets** button, or the target count in the table.

The Edit Targets dialog box will show the targets added directly to the selected task, as well as the targets inherited from enclosing task groups.

Creating a Task Group

Use task groups to organize tasks. Task groups are useful for assigning targets to tasks. When targets are added to a task group, the targets will also be added to the tasks inside the group.

Provide the following information:

1. Type the name for the task group. Task group and task definition names should be unique within a plan.

Note: If you select an existing task group, you can create a task group (or task definition) before, after, in parallel, or inside the selected task group. If you select **Before** or **After**, it specifies the order in which the task groups (or task definitions) should be performed, and sets the task group (task definition) dependencies accordingly. If you decide that you do not like the order, you can change the order by using **Move** or **Set Dependency**.

2. Provide a description of the task group, for example, what is the purpose of the group. A task group is a folder used to group related task definitions.
3. Click **OK**.

Note: The information on this page is not saved until you click **Save and Exit** on the **Create Plan: Review** page.

Operations on Change Activity Plans

You can perform the following on plans:

- [Creating a Plan Like Another Plan](#)
- [Editing a Plan](#)
- [Deleting a Plan](#)
- [Deactivating a Plan](#)
- [Exporting Plans](#)

- [Printing Plans](#)
- [Changing the Owner of a Plan](#)

Creating a Plan Like Another Plan

When you have a plan that meets your requirements, you can create a plan like another plan. This saves you time and money.

1. On the Change Activity Plans page, highlight the plan which you want to mimic. Click **Create Like**. In the dialog box, type the name of the new plan. Click **OK**.

This activates the Create Plan Wizard. Follow the steps outlined in [Creating a Change Activity Plan](#). The following steps alert you to changes that have consequences when creating a plan like another plan.

2. On the Create Plan: Properties page, you can change the name of the plan. You can also change the target type but if you do, all the task definitions will be removed!

Click **Next**.

3. On the Create Plan: Tasks page, you can add, edit, and remove task definitions and task groups. You can also move task definitions and task groups within the plan.

When moving or copying task definitions and task groups, dependencies if set will be retained. It is strongly recommended that all task names be unique. This ensures less confusion when managing tasks after the plan has been activated. During the copy operation, a new name with an updated index is provided to propagate unique naming. The name can be changed using edit.

Click **Next**.

4. On the Create Plan: Targets page, you can add and edit targets associated with a task definition. On the resulting dialog box, select the targets associated with the task definition. There must be at least one target associated with each task definition. Click **OK** after you have made the changes. Click **Next**.

5. On the Create Plan: Activate page, you can activate the schedule.

Note: You can leave this Activate page blank until you are ready to activate the plan.

Click **Next**.

6. On the Review page, ensure all the information you have entered is as you intended. If updates are necessary, click the **Back** button and make the necessary changes.

You also have the option to **Save and Exit** if you are not ready to activate the plan.

Click **Activate** if you are ready to activate the plan. Click **Refresh** on the Change Activity Plans page to see that the plan is activated. You can only manage a plan after it has been activated.

Editing a Plan

With the exception of the activity plan's owner, the types of modifications that can be performed on an activity plan are dependent on the activity plan's status.

- Pending Activation - Indicates the plan is scheduled to start some time in the future. This plan can be modified up to and until the start date at which time it becomes Active.

The **Edit** button is enabled for activity plans that have not yet been activated. Clicking the **Edit** button shows a wizard that is essentially the same as the Create Plan wizard. All edits are enabled. Follow the steps outlined in [Creating a Change Activity Plan](#).

- Active - Once an activity plan has been activated, editing is available only from the Manage Plan page. From this page, it is possible to update the activity plan's scheduled end date and depending on a task's status, assign a task, update a task's start or scheduled end date, and cancel a task.
- Completed, Deactivated - No edits are allowed
- Definition in Progress - Edits are allowed

To edit a plan, on the Change Activity Plans page, highlight a plan in the Change Activity Plans table. Click **Edit**. This activates the Create Plan Wizard.

For information on how to change the owner of a plan, see [Changing the Owner of a Plan](#).

Deleting a Plan

Delete a plan when it is no longer useful.

To delete a plan:

1. On the Change Activity Plans page, highlight a plan in the Change Activity Plans table.
2. Click **Delete**.
3. In the resulting Confirmation dialog box, click **Yes**.

Deactivating a Plan

Deactivate a plan when the plan is no longer relevant to the original task or project, but might prove to be useful for future informational purposes or auditing. Also consider that the structure of this plan could be useful for future projects.

Deactivating a plan results in the plan's tasks being canceled and the plan's progress is no longer monitored. However, for viewing purposes, the plan continues to be available from the Change Activity Plans pages.

If you no longer need to view the plan, consider deleting the plan instead of deactivating it.

Note:

After a plan is deactivated, it cannot be re-activated.

1. On the Change Activity Plans page, highlight a plan in the Change Activity Plans table.
2. From the **Actions** menu, select **Deactivate**.

3. In the resulting Confirmation dialog box, click **Yes, deactivate plan**.

Exporting Plans

Exporting plans captures in an external document the status of the Change Activity Plans.

To export activity plans:

1. From the **Actions** menu on the Change Activity Plans page, select **Export**.
2. Either save the file or open it as a spreadsheet.

The export function exports a list of all the plans, not an individual plan.



Note:

The resulting spreadsheet is in Read Only format.

Printing Plans

Because reporting is a major part of Project tracking and Auditing, print the plans to provide information to management.

To print plans:

1. From the **Actions** menu on the Change Activity Plans page, select **Print**.
2. Print creates a printable page in another tab. Select the tab and then Print content. The content is the list of all Plans, not an individual plan.

Changing the Owner of a Plan

The owner of an activity plan can be updated regardless of the activity plan's status. To change the owner of an activity plan:

1. In the Change Activity Plans table located on the Change Activity Plans page, highlight the plan whose owner you want to change.
2. From the **Actions** menu, select **Change Owner**.
3. In the resulting table, select the name of the new owner. Click **Select**.



Note:

The owner of the plan and Super Administrator can change the owner.

You can change the Owner of a task based on the Target Owner. If the task has a target, and that Target has an Owner defined, then you can set the task Owner to the Target Owner.

Assigning Tasks to the Group Owner

For project purposes, you might need to assign tasks to a manager. This enables the manager to evaluate the work load, the availability of the team, and then assign tasks accordingly.

To assign tasks to a group owner, you must first create a target property. The following steps explain how to create a target property.

1. Create an emcli property:

```
emcli login -username=sysman
emcli add_target_property -property=CAP_assignments -target_type=""
```

where CAP_assignments represents any property (**NOTE:** The property MUST begin with CAP_)
and * represents all targets in the Enterprise Manager Repository

2. From the Targets menu, select Host. From the table, click the host in which you are interested. On the resulting Host home page, from the Host menu located in the upper-left corner, select **Target Setup**, then select **Properties**. Type the target property value for the target.

Using the CAP UI, you can reassign tasks using this target property. For example, on the Tasks tab located on the Manage Plan page, click **Change Owner**, then select **Assign Tasks by Target Property**.

Managing a Change Activity Plan

After you create a plan, you need to manage the plan. The Change Activity Plans page displays a table of plans, along with their status and progress. It is the place to manage a plan once a plan is activated.

Using the manage capabilities of Change Activity Planner, you can:

- View all activated plans; cancel activated plans.
- Drill into the plan and get more details and manage individual tasks. This helps you to identify any issues that may delay the activity plan completion deadline.
- View the audit trail on who made changes to the plan and what changes were made to the plan.
- View and update links. Links associated with the plan during creation can be viewed during plan management. Plan owners and privileged users can update existing links.
- Reassign a plan to a different plan owner; reassign a task to a different task owner.
- Change the scheduled end date of the plan.
- Manage tasks
 - Manually update task status to: acknowledged, in progress, or cancelled.
 - Update status to completed for manual tasks.
 - Maintain audit trail on who made changes to the plan and what changes were made to the plan.
 - View dependent tasks.

To manage a plan, on the Change Activity Plans page, highlight a plan that has been activated and click **Manage** or click the plan name.

The Manage Plan page provides the following features to help you manage your plan:

- **Summary Tab**
Provides an overall summary of the plan's status including the status of tasks and plan summary.
- **Tasks Tab**
Enables you to manage all the plan's tasks. Tasks can be updated individually or in bulk.
- **Comments and Audit Trail Tab**
Enables you to view the audit trail to determine who made changes to the plan and what changes were made to its tasks.

Summary Tab

The Summary tab provides an overall summary of the plan's status. The following information is provided:

- **Overview** - Shows the progress that has been made in the accomplishment of the plan. Plan details are also included: status, scheduled end date, priority, owner, and links.

The estimated end date is the maximum scheduled end date of all the tasks that make up the plan.
- **Plan Summary** - Visual depiction of the task definitions and task groups that make up the plan. Dependencies are indicated by arrows between the objects in the plan and show the order in which tasks should be performed.

You can view the Plan Summary in either graph or table format. The topology graph (and alternate table view) show all the task definitions and task groups that make up the plan and the order in which the tasks should be executed.

The task definitions and task groups all have an associated status, indicated by an icon. The status is determined by the status of the tasks that make up the task definition or task group.

Example: If a task definition has four associated Tasks: Task 1 is Unacknowledged, Task 2 is Complete, Task 3 is In Progress and Task 4 is Canceled, then the task definition status will be In Progress.

Table View

In Table view, there are counts associated with each Task Definition / Task Group. Clicking the links displays the Tasks tab which automatically adds search filters to allow you to see details on that task definition / task group based on the link you clicked. For example, when you click the Overdue link, the task tab displays with the task name set and the Overdue flag checked (in requires attention). You will then see all overdue tasks outstanding for the task definition you were looking at in the table.

From the Table view of the Plan Summary, you can Print and Export the plan.

Graph View

When using the graph format, you can choose the graph format to be either top-down or left-to-right. You can opt for the following annotations to be displayed on the graph: Scheduled Start Date, Scheduled End Date, or Total Tasks.

In graph view, arrows will be drawn between task definitions that have a dependency relationship. Arrows will point from a task definition that needs to be completed first, toward a task definition that must wait for its completion.

Task groups are drawn as an encompassing box around other task definitions.

Hover over a task definition name and the resulting arrow to display information about the tasks created as a result of this task definition.

Right-click on an empty area of the graph to bring up a menu that allows you to switch the orientation of the graph, or print the contents of the graph.

The Display option allows you to toggle between the Graph and Table views.

- If a plan is not yet completed, you can change its scheduled end date by clicking on the calendar icon next to the Scheduled End Date in the Overview section.
- You can view the task completion trend by clicking on the line chart icon next to the Progress(%) in the Overview section.

Tasks Tab

Use the Tasks tab to manage all the plan's tasks. Tasks can be updated individually or in bulk.

- **Status of Tasks**

Pie chart breakdown of the plan by task status.

- **Views**

This section provides quick links to allow you to quickly find tasks that require attention. The supported views are: Unassigned Tasks, Unacknowledged Tasks, Overdue Tasks, Due Within One Week Tasks, and All Active.

All Active shows all open tasks. There is also a Show All link that shows all plan tasks, including completed and canceled tasks.

- **Tasks**

This section displays a table of tasks and allows you to search, view and update the plan tasks. The Actions menu and table buttons provide bulk task update support for operations like setting dates, changing task owners, acknowledging and canceling tasks. By default, the tasks are filtered to show only open tasks.

Select a single task to see and update task details or select multiple tasks, and use the bulk operation buttons at the top of the table, and in the table's Action menu.

The plan task view automatically filters out closed tasks but search filters can be set by expanding the search section, or using the View Links.

Single Task

When a single task is selected, you can view and edit task information using the General tab, for example, target name. You can add task comments and use the Comments and Audit Trail tab to determine who made changes to the plan and what changes were made to the plan. The information is specific to the task you are looking at. Therefore it tracks all changes made to that task including comments that were manually added.

Multiple Tasks (Bulk Operations)

When multiple tasks are selected, you can perform the following:

- Acknowledge a task. Acknowledging a task means you have seen the task but have not started working on it. The only users who can set a task to Acknowledged are: the task owner, the plan owner, and super administrator. When a user is assigned a task, he is allowed to see the plan that contains the task.
- Cancel a task
- Set scheduled start and end dates
- Change owners
- Submit a job. You can submit a job only if the tasks are part of the same task definition.
 - Associate with existing job execution. Choose an existing job execution to associate with all the selected tasks. Note from engineer: Didn't have time to implement the bulk case.
- Create a patch plan. A patch plan can be created when patch plan-based tasks are selected. The tasks must be part of the same task definition. Use the **Patches and Updates** page to deploy the patch plan.
- Associate with existing patch plan. Note that when the patches in the patch template for the task definition have been applied to the task's target, the task will close automatically whether or not a patch plan is associated with the task.
- Associate with existing procedure run. This option is enabled for tasks in the same task definition. Once you create the association, use the **Deployment Library** page to launch the deployment procedure.

The task table also supports Export and Print of the task list. Use the Action menu for these operations.

Requiring Attention

When you click the **Requires Attention** icon in the table, the General tab and the Comments and Audit Trail tabs appear. Study the available information to determine what you should do next. For example, if the plan is overdue, determine what task is causing the delay.

You can study the audit trail for a particular task, for example, dates, whether dependencies exist, and owner or status of the task changes.

Changing Owner

One especially helpful thing you can do is change the owner. This is particularly useful when job responsibilities change and user roles change. Change Owner shows you a list of users; select the new user.

Comments and Audit Trail Tab

View the audit trail to determine who made changes to the plan and what changes were made to its tasks.

Viewing My Tasks

The My Tasks page provides you a way to manage the tasks assigned to you. My Tasks provides an overview of your tasks and allows you to search, view, and update your tasks.

To access the page:

1. From the **Enterprise** menu, select **Configuration**, then select **My Tasks**.
2. Highlight the task on which you want to perform any of the following operations:
 - Change the owner of the task.
 - Cancel the task.
 - Set scheduled start and end dates for the task.

Note: You can select one or more tasks in the table and access bulk operations, such as Acknowledge and Change Owner.

By default the tasks are filtered to show only open tasks.

The following sections are provided on the My Task page:

- Status of Tasks - A pie chart summary of the user's task, by status.
- Advanced search options.
- Views - Provides quick links to allow you to quickly find tasks that require attention. The supported views are: Overdue Tasks, Due Within One Week Tasks, Unacknowledged Tasks, and All Active.

All Active shows all open tasks. The Show All link shows all tasks, including completed, and canceled tasks.

- Tasks - Displays a table of tasks and allows you to search, view, and update the tasks. The Actions menu and table buttons provide bulk task update support for operations like setting dates, changing task owners, acknowledging and canceling tasks.

Selecting a Single Task

Selecting a single task allows you to view and edit the details of the task, and lets you add comments to the task, as well as review the comments and audit trail for the task.

When you select a single task, the data associated with the task is displayed in the following tabs:

- General: Provides task details and tracking information
- Task Details: Provides basic task information like the plan, target and task action. In this section you can update the task's scheduled start date and scheduled end date.

From the task details section, you can submit the job, create a patch plan, associate the task with a patch plan, associate a job with a job execution, and associate a job with a deployment procedure run.

Selecting Multiple Tasks

If you multi-select tasks in the task table, you can use the table buttons (and Actions menu) to perform actions across tasks. For example: Acknowledge, Change Owner, Cancel, Set Scheduled Start Date, Set Scheduled End Date for multiple tasks at one time. In addition, you can Submit Job, Associate with Existing Job Execution, Create Patch Plan, Associate with Existing Patch Plan, and Associate with Existing Deployment Procedure Execution.

- The bulk Change Owner feature supports the following ways of setting the owner:
 - * Assign Tasks to User (Select using the Enterprise Manager user selector.)

- * Assign Tasks to Target Owners: (Only applies to tasks that have targets. The Target Owner must be set for that target for the assignment to take place.) **Note:** This label can change to **Assign Tasks by Target Property**, if you add any Change Activity Planner target properties to your environment. For additional information, see the 'Overview of Change Activity Planner' chapter in the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.
- * Set Tasks to Unassigned (Removes the owner from the task.)
- For job-based tasks, you can select multiple tasks and submit one job to complete the tasks, provided the tasks have the same task definition.

If the task definitions for the selected tasks do not all refer to the same job, the *Submit Job* option will be disabled.
- For deployment procedure-based tasks, all the tasks must have the same task definition. If the task definitions for the selected tasks do not all refer to the same deployment procedure, the *Associate with Existing Deployment Procedure Execution* option will be disabled.
- For patch template based tasks, the plan creator associates the task definition with a patch template. The task owner can either associate the task with an existing patch plan, or create a new patch plan. The patch template must be the same for all selected tasks.
- Tracking: Allows you to see and edit the current owner and the task status. If this task depends on another task, the dependency will be displayed and, in the case where the task is waiting on this dependency, the waiting icon will be visible. If this task requires attention, details of the issues will be provided in this section.
- Comments and Audit Trail: Displays all comments and the audit trail annotations.



Note:

For tasks involving jobs and deployments procedures, if you do not see your assigned task in the Tasks region, ensure you have necessary permissions to see the job and deployment procedure. If you do not have the appropriate access, contact the person who created the job or deployment procedure.

Example of Using Change Activity Planner

This section provides use cases when using Change Activity Planner.

Sections include:

- [Automating Activity Planning](#)
- [Additional Steps in Automating Activity Planning](#)
- [Using Change Activity Planner for Patching](#)

Automating Activity Planning

For example, the System Manager or DBA wants to track and ensure the application of critical security patches to all production databases within three months.

The System Manager or DBA would perform the following steps:

1. Create a new plan and provide the priority of the activity plan, activity plan description, and the target type to which the activity plan applies.

The new activity plan is associated to the selected target type.

2. Create new task definitions/task groups under this plan. This is an iterative process until you finish entering all task definitions and task groups.

Task definitions can be set to complete automatically based on the evaluation of compliance standard rules, or require manual status update from task owners.

For information regarding compliance standard rules, see [Managing Compliance](#).

3. Optional step. Assign at least one target to these tasks on which tasks need to be performed.

Task definitions that have a specified target type must be assigned at least one target of the specified type. If multiple targets are assigned to a task definition, a separate task will be created for each target. **Note:** Task definitions with target type 'None' do not require a target selection.

Note: For tasks involving jobs and deployments procedures, if you do not see the associated job or deployment procedure for your assigned task in the task details region, ensure you have the required permissions to see the job and deployment procedure. If you do not have access, contact the person who created the job or deployment procedure. You may also need to request permission to see specific job or procedure runs if they were submitted by another user.

4. Decide whether Enterprise Manager should automatically assign owners to the new tasks or you can manually assign owners to tasks after you have decided to activate the plan.

5. Activate the activity plan now or at some later point in time.

System either activates the activity plan immediately or schedules it to be activated at the time provided by the manager. Whenever the activity plan is activated, tasks are created and assigned to task owners. Tasks are automatically assigned to a target owner or will be assigned to owners the user has assigned.

If a task is defined such that the task owner updates the task status, the task owner has to manually mark it as completed. If the task status is updated based on a compliance standard rule, the task will be automatically marked as complete when the target for the task is compliant with the rule.

Task groups will be marked complete when all sub-tasks are complete.

Additional Steps in Automating Activity Planning

Here are additional steps you can perform when automating activity planning:

- Supply URL links to the activity plan or task.
- Get graphical topological view of the activity plan.
- Automatically close the task using a compliance standard rule. The system will automatically close the task if the rule check passes on the chosen targets.
- Close the task manually.
- When you assign targets to a task group, all tasks in the task group will get the assignment. You can assign targets at the task group level or the individual task

level. So if all tasks should be performed on Database1, and you select the task group folder and assign target type Database1, all tasks in the group will be assigned Database1. You can then continue to select a single task within the group to be assigned to an additional database target if needed.

- Close the task group automatically. The system automatically closes the task group when all subtasks are complete.
- During the task definition phase, rearrange tasks using delete, move and copy functions. Task definitions and task groups can be relocated in the plan hierarchy to ensure proper dependencies and flow.
- During the task creation phase, create dependency between two tasks, for example task B depends on task A implying that task B cannot start until task A is completed. The system allows users to create dependencies at the same level (dependency only between siblings).

Using Change Activity Planner for Patching

To create patching tasks, select a patch template and associate targets to your patching tasks. Once tasks are assigned to target owners, target owners can apply the patches present in the patch template using the Patch Plan tool.

For information on how to patch targets, refer to [Patching Software Deployments](#).

Part XII

Deployment Procedures

This part contains the following chapters:

- [About Deployment Procedures](#)
- [Customizing Deployment Procedures](#)

About Deployment Procedures

This chapter provides an overview of Deployment Procedures and describes the key aspects you need to know about them. In particular, this chapter covers the following:

- [Overview of the Provisioning Page](#)
- [Granting Roles and Privileges to Administrators](#)
- [Components of a Procedure](#)
- [Creating a Procedure](#)
- [Managing Deployment Procedures](#)
- [Creating, Saving, and Launching User Defined Deployment Procedure \(UDDP\)](#)
- [Procedure Instance Execution Page](#)

Overview of the Provisioning Page

Enterprise Manager provides a framework for automating, orchestrating, and tracking tasks that can be run on multiple Oracle homes. You can perform complex software life cycle management activities such as provisioning, patching, upgrade, and so on from the Cloud Control console. The workflow of all the tasks that need to be performed for a particular life cycle management activity is encapsulated in a Procedure. A Procedure is a hierarchical sequence of provisioning steps, where each step may contain a sequence of other steps. It provides a framework where specific applications and procedures can be built.

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

The Provisioning page has three tabs: Procedure Library, Procedure Activity, and Recycle Bin.

- **Procedure Library Tab:** Use this tab to view a list of all available procedures. For executing procedures or creating new procedures or creating procedures from Oracle supplied ones, a deployment procedure is created. A deployment procedure is a sequence of provisioning steps and phases, where each phase can contain sequence of steps.

Oracle provides best practice deployment procedures that are marked as Oracle under the Created By field. You cannot edit or delete these procedures.

For information about the tasks that can be performed from the Procedure Library page, refer to [Managing Deployment Procedures](#).

- **Procedure Activity Tab:** Use this tab to view a list of all procedure runs that have been submitted for execution and all executing procedures. You can also view the status of a procedure run. In addition to this, you can perform a number of actions

on the submitted procedure like Stop, Suspend, Resume, Retry, Delete, and Reschedule. To understand that actions that you qualify to perform on a procedure, select the procedure. For example, you can Stop or Retry a failed procedure. Starting with Enterprise Manager 12.1.01.3, a new option called **Reschedule** has been introduced, which is enabled for a job that is in progress and has a repeating schedule.


When a procedure is executed, an instance of that procedure is created. This instance keeps track of which step is currently being executed and stores any data collected from the user or any data automatically gathered by executing the action steps.

For an overview of the Procedure Activity tab, see [Overview of the Procedure Instance Execution Page](#)

- **Recycle Bin Tab:** You can delete procedures and runs. When procedures or runs are deleted, they will be internally marked as deleted and will be displayed in the Recycle Bin tab.

The following graphic shows you how you can access the Provisioning screen from within Cloud Control.

Figure 55-1 Accessing the Provisioning Page



Granting Roles and Privileges to Administrators

Administrators are Enterprise Manager users who can login to Enterprise Manager to perform management tasks. The breadth of management tasks available in Enterprise Manager depends on the privileges and roles assigned to the administrators. Roles allow grouping of Enterprise Manager secure resource privileges and can be granted to

administrators or to other roles. Based on the roles, and privileges granted to an Administrator, they can be broadly classified into Designers, and Operators. Normally, the roles and privileges are granted to users or other roles at deployment procedure level, and Software Library level.

This section describes how administrators are granted the predefined roles and privileges that Oracle provides:

- [Granting Roles and Privileges to Administrators on the Deployment Procedure](#)
- [Granting Roles and Privileges to Administrators on Software Library](#)

Granting Roles and Privileges to Administrators on the Deployment Procedure

In a typical data center, the main users of Deployment Procedures are Designers (Lead Administrators) and Operators. Deployment Procedure privileges enable users to perform some design-time activities like setting Privilege Delegation, customizing Deployment Procedure, and run-time activities like running the Deployment Procedure to provision or patch software applications.

Following are the primary users/roles predefined by Oracle for a Deployment procedure, and their associated privileges:

- Super Administrator role allows you to perform all the Administrative operations, and provides full privileges on all the targets.
- EM_ALL_DESIGNER (Designer): This role allows you to perform design time operations on entities. For example, Creating and Monitoring Deployment Procedure templates.

The following table lists all the roles predefined for Designers, and their corresponding descriptions:

Table 55-1 Predefined Roles for Designers

Roles	Description
EM_PATCH_DESIGNER	Role has privileges for creating and viewing any patch plan
EM_PROVISIONING_DESIGNER	Role has privileges for provisioning designer
EM_TC_DESIGNER	Role has privileges for creating Template Collections

Users can be granted any of the following Target Privileges:

- Create Privilege Propagating Group. Privileges granted on a privilege propagating group will be automatically granted to the members of the group.
- Add any target in Enterprise Manager.

Users can be granted any of the following Resource Privileges:

- Create Compliance Entity.
- Create Enterprise Rule Set, basically collection of rules that apply to Enterprise Manager elements, for example, targets and job.

- Create Metric Extension. Metric Extensions allows extending monitoring for a target type by adding new metrics.
- Create new Named Credential that are required to perform Enterprise Manager Administrative Operations.
- Create Any Software Library Entity, Import Any Software Library Entity, Export Any Software Library Entity, and so on.
- Create Template Collection.

**Note:**

For information about EM_PATCH_DESIGNER see, [Table 44-4](#), and for information about EM_PROVISIONING_DESIGNER see, [Table 4-3](#).

- EM_ALL_OPERATOR (Operator): This role has restricted access, and allows you to perform only the run-time activities. For example, Launching a Deployment Procedure.

The following table lists all the roles predefined for Operators, and their corresponding descriptions:

Table 55-2 Predefined Roles for Operators

Roles	Description
EM_ALL_VIEWER	Role Desc
EM_HOST_DISCOVERY_OPERATOR	Role has privileges to execute host discovery
EM_PATCH_OPERATOR	Role for deploying patch plans
EM_PROVISIONING_OPERATOR	Role has privileges for provisioning operator
EM_TARGET_DISCOVERY_OPERATOR	Role has privileges to execute target discovery
EM_USER	Role Desc

Users can be granted any of the following Target Privileges:

- Connect and manage any of the viewable target.
- Add any target in Enterprise Manager.
- Perform administrative operations on all managed targets.
- Run any Operating System Command at any Management Agent

Users can be granted any of the following Resource Privileges:

- Application Replay Operator. Application Replay Entities include captures, replay tasks, and replays.
- Manage custom configurations owned by the user.
- Create new Named Credential that are required to perform Enterprise Manager Administrative Operations.

 Note:

For information about EM_PATCH_OPERATOR see, [Table 44-4](#), and for information about EM_PROVISIONING_OPERATOR see, [Table 4-3](#).

Granting Roles and Privileges to Administrators on Software Library

Software Library is a centralized media storage for all Enterprise Manager entities. Super Administrator is responsible for configuring the Software Library, once the Enterprise Manager installation is complete. After the Software Library is configured with Storage Locations, it becomes usable to store entities. Designers and Operators are the main users of Software Library who perform the design-time and run-time activities respectively. The design-time activities include Customizing entities, Creating entities, Importing entities, Exporting entities, and so on. The run-time activities performed by Operators include running deployment procedures which in turn use any the entities stored in the Software Library.

For more information about Software Library users, roles, and their associates privileges, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

 Note:

To run the procedure on a Windows host which involves executing some Software Library entities (for example, directive scripts), you (**the Windows user**) must be granted the following privileges:

- Act as part of the operating system
- Adjust memory quotas for a process
- Logon as batch job
- Replace a process level token

If not, the execution of the directive steps in the procedure may fail.

Components of a Procedure

This section describes the following:

- [Target List](#)
- [Procedure Variables](#)
- [Phases and Steps](#)

Target List

Target List is a pre-populated list of targets on which you can run your job. Phases operate on a set of Enterprise Manager targets, collectively known as a target list. Each phase must be associated with a target list. When the Deployment Procedure is selected for execution, the Deployment Procedure Manager will prompt the user to assign targets to the target list.

Starting with Enterprise Manager 12c (12.1.0.2), Custom Target Lists have been introduced. In addition to the default target list, you can now have your own customized lists of targets on which designated Phases can run. The advantage of this approach is that you can have multiple custom target lists, and assign it to the different phases in your procedure. This allows you to choose the Target List on which you want the Phase to iterate.

The following examples describes the various scenarios which employs one, two, or more target lists in a procedure:

- For copying a jar file to multiple hosts, you will need just one target list. You may choose to use the default target list for this purpose.
- For cloning an Oracle Home, and provisioning it on multiple targets, you will need a minimum of two target lists: one target list for the source which contains only a single target, and a second target list which contains all the destination targets.
- For provisioning or patching a WebLogic Server, you might require three separate target lists one for the Administration Server, one for the Managed Servers, and one for the Database.

Procedure Variables

Procedure Variables are user-defined variables that can be used while customizing a procedure. Normally, when you add a custom step to a User-owned procedure or customize an Oracle-owned procedure, then you might need to declare procedure variables that you can later use in your custom step or phase.

To access the Procedure Variable tab, from the **Enterprise** menu, select **Provisioning and Patching** and then select **Procedure Library**. From the menu, select **Create New**, then click **Go**. The following page is displayed:

Select	Name	Display Name	Description	Type	Value Options	Required
☐	isPingSuccessful	Is Ping Successful Variable		String	Text	☐

To declare the Procedure Variable, you must enter a unique name, a description for it. Optionally, you can select the password check box to make the variable secure.

You can create two types of Procedure Variables that can be later used while launching the deployment procedure. They are as follows:

- **String:** This variable once declared at design-time, can be used by operator to specify the values at run time. For example, DBVER (the version of the database).
 - * Text, allows you to enter one value for the variable. For example, staging location, host name, profile name, and so on.
 - * Password, allows you to provide a password variable. For example, host password, WLS password, and so on.
 - * List of Values, allow you to enter many values for a variable. To provide multiple values for a variable, click **Add**, then enter the details like **Value**, **Display Name**, and a **Description** for the variable. For example, let's declare a variable called country with multiple values as follows

Value	Display Text
IDC	India
US	America
IE	Ireland

- **Software Library Entity:** This variable allows you the flexibility of binding the variable to Software Library Directive or Component at the time of launching the procedure. Earlier these values had to be specified at design-time while creating the procedure, now with the introduction of Software Library entity variable you can specify the values dynamically at the time of launching the procedure.

**Note:**

You cannot add Procedure Variables to a Deployment Procedure that is owned by Oracle.

Phases and Steps

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

- [Types of Phases](#)
- [Types of Procedure Steps](#)
- [Performing Tasks on Procedure Steps](#)

Types of Phases

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.

Types of Procedure Steps

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 that task that requires user interaction and cannot be automated. Typically, Deployment Manager would display the instructions that need to be performed by the user. After the operation is performed, the user proceeds to the next step.

Examples of a Manual Step:

- Log on to a system and update the kernel parameter.
- Reboot a system.

- Provide special privileges to the user. For example, SSH Setup.

- *Computational Step*

Computational Step is that task whose operations are performed within the Deployment Engine and does not require any user intervention. This step gathers additional information for executing a procedure. This step cannot be inserted by a user, and only Oracle Corporation can insert this step.

Examples of Computational Step:

- Executing SQL query against an Enterprise Manager schema to gather more data for other steps to run.
- Retrieving target properties from the repository and updating the runtime information.

- *File Transfer Step*

File Transfer Step is a step used for copying files and/or directories from one host to one or more hosts. You can archive files and directories transferred from a source host to the destination hosts. When this step is inserted within a phase, you can set the Source and Destination Targets using existing variables.

For example, to copy a directory from a host X to the hosts associated with the phase, then for Source Target select "Set Value" and assign host X, and for Destination Target select "Choose Variable" and assign it to the option "TargetVariable:Current Target".

- *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.

Examples of Job Step:

- * Cloning an existing Oracle home.
- * Staging a patch.
- * Starting a database.

- *Library: 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 the Software Library and reuse them in a Deployment Procedure.

For more information about Directives, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Examples of Directive Step:

- * Executing root scripts.
- * Applying `catpatch.sql` and restarting the database.
- * Confirming if the prerequisites have been met.

— *Library: Component*

A 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 has one directive associated with it. 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 will be ignored while executing the step.

For more information about Components, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Examples of Component Step:

- * Applying a patch.
- * Performing prerequisites before performing an installation.
- * Installing Oracle software on target machines.

— *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.

Examples of Host Command Step:

- * Starting a Management Agent (`emctl start agent`).
- * Shutting down OPMN (`opmnctl stopall`).
- * Restarting OID.

Performing Tasks on Procedure Steps

In the Procedure Steps tab, you can perform the following operations on the selected step or phase:

- **Enable/Disable:** 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. To do so, follow these steps:
 1. To disable a phase or step, select the phase or step you want to disable, and click **Disable**.
 2. To enable a phase or step, select the phase or step you want to enable, and click **Enable**.
- **Delete:** Select the step or phase you want to delete, and click **Delete**.

**Note:**

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.

- **Insert:** To add a new Step or Phase, click **Insert**. In the Create wizard, [do one of the following](#):
 - Add a Phase. See [Adding Rolling or Parallel Phase](#)
 - Add a Step. See [Adding Steps](#)
- **Edit Step:** To edit a Step or Phase, click **Edit Step**. Depending upon your selection either the Edit Phase or Edit Step wizard is displayed. Accordingly, follow the steps available in:
 - Edit a Phase.
 - Edit a Step.

Creating a Procedure

To create a custom procedure, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. In the Procedure Library, from the list of actions, select **Create New** and click **Go**.
3. On the Create New Procedure page, in the **General Information** tab, provide a Name and description for the procedure, Procedure Utilities Staging Path, and Environmental Variables.
4. On the Create New Procedure page, click **Target List** tab. You can create your own custom target lists on which designated Phases can run. See [Target List](#).

The advantage of this approach is that you can have multiple custom target lists, and assign it to the different phases in your procedure.
5. On the Create New Procedure page, click **Procedure Variable** tab. You can create your own procedure variables. See [Procedure Variables](#).
6. On the Create New Procedure page, click Procedure Steps tab. This tab allows you to add phases and steps to your procedure. For more information about adding a phase or a step, see sections [Adding Rolling or Parallel Phase](#) and [Adding Steps](#).

Adding Rolling or Parallel Phase

Ensure that you have created a Target List before creating a phase.

To insert a phase, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching** and then select **Procedure Library**.
2. On the Provisioning page, in the Procedure Library tab, from the menu, select **Create New**, then click **Go**.

3. On the Create New Procedure page, select **Procedure Steps** tab.
4. Select the Default Phase, and click **Insert** to add a phase.

 Note:

When creating a phase inside another phase, for the insert location, select **After "Default Phase"** or **Before "Default Phase"**. **Inside "Default Phase"**, you will not be able to select any target in the next page.

5. In the Create wizard, do the following:
 - a. On the Create page, specify general information about the phase as described in the following table:

Table 55-3 Field Description - Adding Rolling Phase

Field Name	Description
Select	Select <i>Phase</i> .
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> . To insert it inside the phase or step selected, select Inside<phase or step> , Otherwise, select Before <phase or step> .
Type	If you are adding a rolling phase, then select <i>Rolling</i> . If you are adding a parallel phase, then select <i>Parallel</i> .
Error Handling	Select the error handling mode you want to set for the custom phase. 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 Cloud 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.

- b. On the Select Target List page, select a target list to indicate the type of targets on which the new phase should run.

All the target lists declared while creating the procedure is listed in the drop down menu, select the target list to use for this phase. The actual targets can be chosen when the procedure is being launched.
 - c. On the Review page, review the information you have provided for creating a new phase, and click **Finish**.

Adding Steps

This section explains how you can add different types of steps to a Deployment Procedure. In particular, it covers the following:

- [Adding a Job Step](#)
- [Adding a Directive Step](#)
- [Adding a Component Step](#)
- [Adding a File Transfer Step](#)
- [Adding a Host Command Step](#)
- [Adding a Manual Step](#)

Adding a Job Step

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).
2. 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**.
3. On the Map Properties page, specify values for the parameters that are required by the selected job type. Additionally, you can set the target List to be applied for this step.
4. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Adding a Directive Step

To add a directive step to a Deployment Procedure, follow these steps:

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).
2. On the Select Directive page, if you have selected a Software Library variable in the Procedure Library tab, then you can select one of the following options. If not, you can directly select a directive from the table, and click **Next**.
 - **Select New Directive:** This option lists all the directives available in Software Library, select a directive from the list that you want to run on the targets. Provide necessary values in the Select Directive section to narrow down the search results.
 - **Select New Software Library Entity Variable:** Select the Software Library variable that you declared while creating the procedure from the **Software Library Entity Variable** drop down menu. This variable behaves as a place holder and enables you the flexibility of binding it with the directives dynamically while launching the procedure. Essentially, you can use this variable to select certain entities which do not need any parameters.

For example, user-defined scripts like a Perl to print the current directory location that does not need any parameters to be passed.

- **Select New Software Library Entity Variable with Directive Properties:** This option allows you to bind a Software Library entity variable with directives that are available in Software Library. Ensure that you choose a directive whose properties (signature) matches with the entity declared.

Select **Always Use Latest Revision** so that only the latest revision of the selected directive will be used at all times.

3. On the Map Properties page, provide the following details:
 - By default the **Run Directive**, and **Perform Cleanup** options are enabled to run the script, and remove the files after the steps has run.
 - In the Directive Properties section, specify values for the properties associated with the selected directive. You have the option of providing or not providing the property values at this stage. If you do not provide the property values now, then they are prompted at the time of launching the procedure.
 - In the Credentials section, set the target List to be applied for this step.
 - In the Time limit properties section, you can set a max time allowed for an operation to complete in seconds. For example, let's assume that you have a huge procedure with numerous steps and you do not want to block the whole execution if one step fails (because agent is down). In such a scenario, setting a time limit on a step is very effective. If you set a time limit of 75 seconds on a step, then if the job exceeds this set time, the step is skipped.
4. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Adding a Component Step

To add a generic component step to a Deployment Procedure, follow these steps:

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).
2. On the Select Component page, select a component name from the table, and click **Next**. However, if you have set a Software Library variable in the Procedure Library tab, then you must select one of these following options:
 - **Select New Component:** This option lists all the components available in Software Library, select a component from the list that you want to run on the targets. Provide necessary values in the Select Component section to narrow down the search results.
 - **Select New Software Library Entity Variable:** Select the Software Library variable that you declared while creating the procedure from the **Software Library Entity Variable** drop down menu. This variable behaves as a place holder and enables you the flexibility of binding it with the components dynamically while launching the procedure. Essentially, you can use this variable to select certain entities which do not need any parameters.

For example, user-defined scripts like a Perl to print the current directory location that does not need any parameters to be passed.

- **Select New Software Library Entity Variable with Component Properties:** This option allows you to bind the Software Library variable with the components that are available in Software Library. Ensure that you choose a component whose properties (signature) matches with the entity declared.

If you check **Always Use Latest Revision**, then only the latest revision of the selected component will be used at all times.

3. On the Select Directive page, you can select a directive from the table, and click **Next**. However, if you have set a Software Library variable in the Procedure Library tab, then you must select one of these following options:

- **Select New Directive:** This option lists all the directives available in Software Library, select a directive from the list that you want to run on the targets. Provide necessary values in the Select Directive section to narrow down the search results.

- **Select New Software Library Entity Variable:** Select the Software Library variable that you declared while creating the procedure from the **Software Library Entity Variable** drop down menu. This variable behaves as a place holder and enables you the flexibility of binding it with the directives dynamically while launching the procedure. Essentially, you can use this variable to select certain entities which do not need any parameters.

For example, user-defined scripts like a Perl to print the current directory location that does not need any parameters to be passed.

- **Select New Software Library Entity Variable with Directive Properties:** This option allows you to bind a Software Library entity variable with directives that are available in Software Library. Ensure that you choose a directive whose properties (signature) matches with the entity declared.

Select **Always Use Latest Revision** so that only the latest revision of the selected directive will be used at all times.

4. On the Map Properties page, provide the following details:
 - By default the **Run Directive**, and **Perform Cleanup** options are enabled to run the script, and remove the files after the steps has run.
 - In the Directive Properties section, specify values for the properties associated with the selected directive. You have the option of providing or not providing the property values at this stage. If you do not provide the property values now, then they are prompted at the time of launching the procedure.
 - In the Credentials section, set the target List to be applied for this step.
 - In the Time limit properties section, you can set a max time allowed for an operation to complete in seconds. For example, let's assume that you have a huge procedure with numerous steps and you do not want to block the whole execution if one step fails (because agent is down). In such a scenario, setting a time limit on a step is very effective. If you set a time limit of 75 seconds on a step, then if the job exceeds this set time, the step is skipped.
5. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Adding a File Transfer Step

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).
2. On the Map Properties page, select the Source Target from which you want to transfer files, the source target path, the Target Destination for file transfer and the

destination path. Specify the Source and Destination Credential Usage, whether Host or Privileged Host credentials. Click **Next**.

If you select **Transfer all the files in this path** option, then all the files in the source path are transferred. If uncheck this option, then the **Source File Name** field becomes mandatory.

3. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Adding a Host Command Step

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).
2. 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 menu. 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 menu. 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"`.

3. In the Time limit properties section, you can set a max time allowed for a script to run in seconds. For example: Let's assume that you have set this value to 75 seconds, then when the script runs if it exceeds the set time, then this step is skipped.
4. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Adding a Manual Step

In the Create wizard, do the following:

1. On the Create page, specify general information about the step as described in [Table 55-4](#).

2. 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.

3. On the Review page, review the information you have provided for creating a new step, and click **Finish**.

Table 55-4 Field Description - Adding Steps

Field Name	Description
Select	Select <i>Step</i> .
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: Component. • For a file transfer step, select File Transfer • 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 phase. 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 Cloud 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.

Managing Deployment Procedures

This section contains the following:

- [Viewing, Editing, Deleting a Procedure](#)
- [Editing and Saving Permissions of a Procedure](#)
- [Tracking the Procedure Execution and Status of Deployment Procedures](#)
- [Rescheduling a Procedure](#)
- [Reverting a Procedure](#)

Viewing, Editing, Deleting a Procedure

To view, edit, or delete an existing procedure, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, do the following:
 - For Viewing the procedure, select the deployment procedure, and from the actions menu, click **View Procedure Definition**.
 - For Editing the procedure, select a user-defined procedure, and from the actions menu, select **Edit Procedure Definition** and click **Go**. If you want to customize an Oracle-provided procedure, from the actions menu, select **Create Like** and click **Go**. Save the procedure, and then customize it.
 - For Deleting the procedure, select the deployment procedure, and from the actions menu, click **Delete**.

Editing and Saving Permissions of a Procedure

A designer with Super Administrator privileges has the access to edit the permissions of a Deployment Procedure, and save it.

To edit the permissions on a Deployment, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, from the actions menu, select **Edit Permissions**, and then click **Go**.
3. On the Edit Permissions: <target name> page, click **Add**. From the Search and Select Administrator or Role dialog box, select the administrators or roles to which you want to grant the permissions, and click **Select**.
4. On the Edit Permissions: <target name> page, select the Role and the privileges that you want to grant to each of these roles. A *full* privilege will let the Operator edit the Deployment Procedure, and a *Launch* privilege will only allow an Operator to view and run the Deployment Procedure. Click **OK** to save these grants.

Provisioning
Edit Permissions: 'Provision Oracle Database Client'

Owner

Remove Add

Select All Select None

Select	Name	Type	Description	Manage Target Privilege Grants
☒	OPER	Administrator		Launch
☐	PUBLIC	Role	PUBLIC role is granted to all administrators. This role can be customized at site level to group privileges that need to be granted to all administrators	Launch Full Launch

Tracking the Procedure Execution and 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 Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, click the **Procedure Activity** tab.
3. On the Procedure Activity page, click the procedure to view the status of all deployment procedures in various stages of their lifecycle.

Table 55-5 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.
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 remaining steps. However, you always have the option of retrying the failed steps. Alternatively, you can ignore the error, and proceed further.
Succeeded	Implies that the Deployment Procedure has successfully completed its execution.
Skipped	Implies that the Deployment procedure has skipped the execution of this step. Primarily, a step is skipped when the condition attached to the step evaluates to false.
Stopped	Implies that the Deployment Procedure has been permanently stopped from execution by the user.
Saved	Implies that the Deployment Procedure has not been submitted 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 <i>Skip target/Continue on Error</i> 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**.

Rescheduling a Procedure

Use this page to reschedule a procedure. You can reschedule a procedure only when the status of the procedure is **Scheduled**. In all the other cases, this button will appear grayed out. To reschedule a procedure, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity** tab.
2. On the Procedure activity page, select the scheduled procedure from the table, and click **Reschedule**.
3. On the Reschedule procedure page, select the date and time when you want to run the procedure.
4. From the Repeat menu, select an option from the menu to run the job at the selected frequency.
5. After you set all the desired parameters, click **Reschedule**.

Reverting a Procedure

Use this page to revert to a previous version of a procedure. To revert to another version of the procedure, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Procedure Library, from the menu, select **Revert**. Note that the revert option is enabled only if the procedure was edited by user.
3. Click **Go**.

For example, if you have version 1.5 as the latest and you revert to version 1.3, a new version 1.6 is created, which will be the same as version 1.3.

Setting Step Level Grace Period

You can now set an OMS-wide step level grace period using the Enterprise Manager command line (emctl) utility. For a procedure submission, if an agent is not reachable for a time of period longer than grace period, the step will be marked as failed.

To set the step level grace period, run the following command:

```
emctl set property -sysman_pwd sysman -name  
oracle.sysman.core.procedure.steplevel_graceperiod -value <value of the grace period>
```

For example, to set a grace period of 12 mins for a step, run the following command:

```
emctl set property -sysman_pwd sysman -name  
oracle.sysman.core.procedure.steplevel_graceperiod -value 12
```

Submit any procedure using the Cloud Control UI, if the agent machine is not reachable for a time period longer than 12 minutes, then the step is marked as failed.

 Note:

Things to keep in mind:

- The grace period is counted in minutes.
- The minimum default value of grace period is 10 mins.
- The grace period is only applicable to the steps in a procedure. The procedure level grace period is assigned when submitting the procedures using Cloud Control or EMCLI.

Creating, Saving, and Launching User Defined Deployment Procedure (UDDP)

Creating a procedure from scratch by inserting the required phases, steps, variables, and so on is possible with User Defined Deployment Procedure. This functionality has been introduced in Enterprise Manager 12c to allow users to completely customize a procedure to suit their requirements.

 Note:

For a video tutorial on creating and using the User Defined Deployment Procedure, see:

[Oracle Enterprise Manager 12c: Implement User-Defined Deployment Procedures](#)

Broadly, the process can be divided into two subcategories as follows:

- [Step 1: Creating User Defined Deployment Procedure](#)
- [Step 2: Saving and Launching User Defined Deployment Procedure with Default Inputs](#)
- [Step 3: Launching and Running the Saved User Defined Deployment Procedure](#)
- [Step 4: Tracking the Submitted User Defined Deployment Procedure](#)

 Note:

For a workflow example on User Defined Deployment Procedure with illustrations on how to provision JRE6 on a Linux host `abc.example.com`, see [Provisioning User Defined Deployment Procedure](#).

Step 1: Creating User Defined Deployment Procedure

Log in to Enterprise Manager Cloud Control with designer privileges to create a UDDP template. To do so, follows these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, select **Procedure Library**.
2. On the Provisioning page, from the list of actions, select **Create New**, and click **Go**.
3. On the Create New Procedure page, in the General Information tab, provide a unique name and description for your procedure.
4. In the **Target Lists** tab, you can use the default `host_target_list` variable or add any number of new custom target lists. Adding new custom target lists enables you to group the targets which in turn allows phases to use separate target lists (targets) that they can iterate on.
5. In the Procedure Variables tab, click **Add Row** to define procedure variables. In addition to String type, you can add Software Library Entity variable. For more information about this, refer to [Procedure Variables](#).

Specify the **Variable Name**, **Display Name**, **Description**, and **Type** from the drop down menu. Also define whether the variable is a password and a mandatory field.

6. In the **Procedure Steps** tab, select the default phase, and do the following:
 - a. Select Default Phase, and click **Insert**. For information on inserting a phase, see [Adding Rolling or Parallel Phase](#).

 Note:

Without declaring a Target List, you can not proceed with the creation of a phase.

- b. Select the phase you created, and then click Insert to insert steps. For information on inserting steps, see [Adding Steps](#).
7. Repeat Step 6 to insert steps and phases according to the procedure you want to create.
8. Click **Save and Close** to save the procedure. You can run the configuration for future deployments.

Step 2: Saving and Launching User Defined Deployment Procedure with Default Inputs

Log in to Enterprise Manager Cloud Control with Operator privileges to launch the saved UDDP with default values. To do so, follows these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, select **Procedure Library**.
2. On the Provisioning page, select the saved UDDP, and click **Launch**.
3. On the Select Targets page, select the target list from the drop down menu, and click **Add** to populate the target list. Click **Next**.
4. If you declared variables that you did not define during the procedure creation, then you will have to provide the details in the Set variable page. All the unbound variables are displayed, enter appropriate values for the same.

If you have declared a Software Library Entity variable, then you could search and select the desired entity from Software Library. Once the value is populated, you may even choose to lock this value so that any other Operator who does not have privileges on your procedure will not be able to update this values. For more information on different types of variables, see [Procedure Variables](#). Click **Next**.

5. On the Set Credentials page, you need to set the credentials for the target host machines that you have included as a part of the host_target_list variable. Click **Next**.
6. On the Set Schedule and Notification page, you can schedule the job to run immediately or at a later preferred time.
7. Click **Save**, and provide a configuration name to save the job template with default values.

Some of the procedures allow you to not just save the procedure with default values, but also lock them. For example, the following Database Provisioning procedure describes how to save and launch a procedure with lock downs.

Saving and Launching the Deployment Procedure with Lock Down

Lock Down is a new feature introduced in Oracle Enterprise Manager Cloud Control 12c that enables administrators with Designer privileges to standardize the Deployment Procedures across the enterprise. If Designers with Super Administrator privileges create Deployment Procedure templates with lock downs, and save them, then these templates can be used by Operators who can launch the saved Deployment Procedures, make changes to the editable fields, and then run them.

To create a Deployment Procedure with lock downs, an administrator logs in with designer privileges, and launches a Deployment Procedure. In the interview wizard of the Deployment Procedure, the designer enters the values for certain fields and locks them so that they cannot be edited when accessed by other users. For example, in the following graphic, fields such as **Database Version**, **Database type** are locked by the designer, and when an operator launches the same deployment procedure, these fields will be grayed out, and cannot be edited:

Create Database : Database Version and Type

Select database version

Version

Select database type

☒ Oracle Single Instance Database
☐ Oracle Real Application Clusters (Oracle RAC) Database
☐ Oracle RAC One Node Database

Prior to database creation, ensure that Oracle Home is available on the host. If Oracle Home is not already available, use the Pr

Hosts

You can create the database on multiple hosts. Select hosts in the table below. If you are selecting multiple hosts, you have a

[Specify common settings for Oracle Home and Credentials](#)

[+ Add Hosts](#) [X Remove Host](#)

Hostname	Oracle Home	Host Credentials
scag01db01.us.oracle.com	/u01/app/oracle/product/11.2.0/dbhome_2	NC_SCAG01DB_2012

In the following use case, user logs in with designer privileges to provision a Single Instance Database on a Linux host. *Designer* updates most of the values prompted in the wizard and locks them as he/she does not want other users such as Operators to have edit privileges on them. Some of the fields such as adding targets, and some additional configuration details are not locked. The Deployment Procedure is then saved with a unique procedure name, but not submitted. A user with *Operator* privileges logs in and runs the saved procedure after updating all the editable fields, such as adding targets, additional configuration details.

Broadly, it is a two-step process as follows:

- [Step 1: Saving a Single Instance Database Deployment Procedure with Lock Downs](#)
- [Step 2: Launching the Saved Single Instance Database Deployment Procedure](#)

Step 1: Saving a Single Instance Database Deployment Procedure with Lock Downs

In the following section user logs in to Cloud Control as a designer (ATBARB0Z1), and provisions a Single Instance Database with lock downs as follows:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
2. On the Database Provisioning page, select **Provision Oracle Database**, and click **Launch**.
3. In the Select Hosts page, in the Select hosts section, click **Add** to select the destination host where you want to deploy and configure the software.

If you want to use a provisioning profile for the deployment, choose **Select a Provisioning Profile** and then select the profile with previously saved configuration parameters.

In the Select Tasks to Perform section, do the following and lock the values:

- Select **Deploy Database software** to provision single instance databases.
- Select **Create a New Database** to create a new database and configure it after installing the standalone Oracle Database.

Provision Oracle Database: prov_db_template : Select Hosts [Save] [Back] Step 1 of 5

Select provisioning profile

Provisioning profiles allows you to record inputs and use them later while performing deployments using standard values. Selecting a profile at this stage allows you pre-populate the interview with previously saved values of deployment and configuration parameters.

☒ Do not use a Provisioning Profile
☐ Select a Provisioning Profile

Name	Description
Profile for Single Inst:	This is a reference profile created for provisioning database on file system
Profile for Single Inst:	This is a reference profile created for provisioning database on ASM

Select tasks to perform

Specify the tasks to perform as part of the provisioning process.

Deploy software

☐ Deploy Grid Infrastructure for standalone server
☒ Deploy Database software

Configure software

☐ Configure Grid Infrastructure
☒ Create a new database

Select destination hosts

View Refresh Hosts

Target Name	Last Collection	Host Name
slc00ean.us.oracle.com		slc00ean.us.oracle.com

4. On the Configure page, the following configuration options appear:

- On the Configure page, click **Setup Hosts**. On the Operating System Users page, specify the operating system user for the Oracle Home for the database. For Oracle Home User for the database, select the Normal User and Privileged User to be added to the OS group, and lock the values. Click **Next** to proceed. On the Specify Operating System groups page, specify the OS Groups to use for operating system authentication and lock the values, as appears in the following graphic.

Provision Oracle Database : Specify OS users

Operating system users

Specify the operating system users required to provision the software.

☒ Use Preferred Credentials ☐ Override Preferred Credentials

Oracle Home User	Normal user	Privileged User
Oracle Database	Privileged Host: Credentials	Privileged

Provision Oracle Database : Specify OS groups

Operating system groups

Specify the operating system groups required to provision the software. Local groups will be automatically added to the 'Inventory group (OINSTALL)' if necessary. If you have an external group, you can also specify it as part of the provisioning procedure.

Group label	OS Group name
Inventory Group (OINSTALL)	oinstall
Database Administrator (OSDBA)	dba
Database Operator (OSOPER)	oper

Configure


Use the following sections to provide configuration details for the part of this provisioning operation.

Task No.	Task	Status
1	Setup hosts	
2	Deploy software	
3	Create databases	
4	Compliance Standards	

Click **Next** to come back to the Configure page.


- On the Configure page, click **Deploy Software**. On the Select Software Locations page, specify the source and destination locations for the software binaries of Oracle Database. Update the values for all the fields, and click the Lock

icon so that the fields can not be edited by a user with Operator privileges, as appears in the following graphic:



Click **Next** to come back to the Configure page.

- On the Configure page, click **Create Databases**, the following screen appears. Update only the mandatory fields, and click **Next** to proceed. Do not lock any of the values in this wizard:



- On the Configure page, click **Compliance Standards**. On the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**. Do not lock the values.

- On the Schedule page, specify a Deployment Instance name. In the Schedule section, select **Immediately**. You can set the notification preferences according to deployment procedure status, and click **Next**.

- In the Review page, review the details you have provided for the deployment procedure. Click **Save** to save the deployment procedure with a unique name **prov_db_template**, and then click **Cancel**. The Procedure library page appears with the saved procedure.


Select	Procedure	Type	Parent	Version	Last Updated	Description	Last Modified By	Owner
☒	prov_db_template	Database Provisioning	Provision Oracle Database		Jan 19, 2012 6:07:09 AM UTC			ATBARBOZ1

Step 2: Launching the Saved Single Instance Database Deployment Procedure


In the following section user logs in as a Operator (SSIRAJUD1), and runs the saved Deployment Procedure **prov_db_template** to provision a Single Instance Database.

- In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.


2. In the Procedure Library, select a procedure **prov_db_template**, and click **Launch**.




3. On the Select Hosts page, in the Select hosts section, click **Add** to select the destination host where you want to deploy and configure the software, and then click **Next**.



4. On the Configure page, the following configuration options appear:
 - On the Configure page, click **Setup Hosts**. Since the values here are locked by the designer, you will not be able to edit them. Click **Next** to come back to the Configure page.



- On the Configure page, click **Deploy Software**. Since the values here are locked by the designer, you will not be able to edit them. Click **Next** to come back to the Configure page.



- On the Configure page, click **Create Databases**. The following screen appears. Update all the fields, and click **Next** to proceed.

For information about updating the Creating Database wizard, see [Provisioning and Creating Oracle Databases](#).

Create Database : Database Template

Specify Template

Database template contains database configuration and structural information. 'Structure and Data' templates are useful to deploy pre-created databases with seeded data across the enterprise quickly. 'Structure Only' Templates should be used when you want to customize parameters such as database size, datafiles, tablespace attributes, and database options when creating the database.

☐ Select Template From Software Library

Template No template selected

Temporary Storage Location on Managed Host(s) /tmp/

☒ Select Template From Oracle Home

Template	Includes Data
General Purpose or Transaction Processing	☒
Custom Database	☒
Data Warehouse	☒

- On the Configure page, click **Compliance Standards**. On the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.

Configure

Configure Configuration Standards Target Association Configure

Provision Oracle Database : Configuration Standards Target Association

Configuration Standards Target Association

Associate a Compliance Standard with the database(s) that will be created as part of this procedure.

Compliance Standard None

- On the Schedule page, specify a Deployment Instance name. In the Schedule section, select **Immediately**. You can set the notification preferences according to deployment procedure status, and click **Next**.
- In the Review page, 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.

Step 3: Launching and Running the Saved User Defined Deployment Procedure

Log in to Enterprise Manager Cloud Control with Operator privileges to run the saved UDDP. To do so, follows these steps:

- In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, select **Procedure Library**.
- On the Provisioning page, select the saved UDDP configuration template that you saved as a part of the previous step, and click **Launch**.

Note:

While creating the UDDP template, if you have locked any of the values, then they will appear greyed out since they are read-only values that cannot be edited now.

3. On the Select Targets page, select the target list from the drop down menu, and click **Add** to populate the target list. Click **Next**.
4. If you declared variables that you did not define during the procedure creation, then you will have to provide the details in the Set variable page.

If you have declared a Software Library Entity variable, then you could search and select the desired entity from Software Library. Once the value is populated, you may even choose to lock this value so that any other Operator who does not have privileges on your procedure will not be able to update this values. For more information on different types of variables, see [Procedure Variables](#). Click **Next**.
5. On the Set Credentials page, you need to set the credentials for the target host machines that you have included as a part of the host_target_list variable. Click **Next**.
6. On the Set Schedule and Notification page, you can schedule the job to run immediately or at a later preferred time.
7. Click **Submit**, and provide a unique **Submission Name** for your job.

Step 4: Tracking the Submitted User Defined Deployment Procedure

Follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, select **Procedure Activity**.
2. On the Procedure Activity page, click the job that you submitted.
3. The new Instance Execution page for the job is displayed which will give you information about the success or failure of your job for on all the targets.

For more information about the new Instance Execution page, see [Overview of the Provisioning Page](#).

Procedure Instance Execution Page

The following tasks can be performed from the Procedure Instance Execution page:

- [Comparison Between the Existing Design and the New Design for Procedure Instance Execution Page](#)
- [Overview of the Procedure Instance Execution Page](#)
- [Investigating a Failed Step for a Single or a Set of Targets](#)
- [Retrying a Failed Step](#)
- [Creating an Incident](#)
- [Viewing the Execution Time of a Deployment Procedure](#)
- [Searching for a Step](#)
- [Downloading a Step Output](#)
- [Accessing the Job Summary Page](#)

**Note:**

For a video tutorial on the Procedure Execution Tracking page, see:

[Oracle Enterprise Manager 12c: View Deployment Procedure Execution Details](#)

Comparison Between the Existing Design and the New Design for Procedure Instance Execution Page

The current Procedure Activity page provides the status of all the steps executed in a deployment procedure. This page also gives you information of the failed step and the necessary action to be taken to rectify it.

Before you understand the new Procedure Activity page, take a moment to review the challenges you might be facing while using the current Procedure Activity Page:

Table 55-6 Comparison Between the Existing Procedure Activity Page and the New Procedure Activity Page.

Category	Existing Procedure Activity Page	New Procedure Activity Page
Screen Design	The screen design allows you to access one step at a time.	Optimal screen design which allows you to access all the steps and targets from the same page without having to drill down
Multiple Selects	Multiple selects are not supported.	Multiple selects are possible from a single page. For example, if you want to select only the failed steps you can do so using the new design.
Target-Centric Design	Step-centric approach restricts your access to all the targets from a single screen. Which means that in the earlier approach you could drill down to only one failed step at a time, and would have to repeat the whole procedure for the other failed steps	Target-centric design with the introduction of filters have made it easy to analyze all failed steps from the same page and perform the required action on the step.
Step Output	The step-centric design requires traversing through a number of pages to drill down to the actual step.	Target-centric design now allows you to view all the step details from the same page. unlike the earlier step-centric.
Detailed Output	Detailed Output for a step was not available in the earlier design. You had to download the entire log.	Detailed Output is a new option available at step-level which captures the log information pertaining to that step selected, only making it easy to view and debug the step in case of a failure.
Incident Creation	Incident Creation was not available in the earlier design.	Incident Creation is a new feature that has been introduced at Procedure-Level which enables you to create an incident for the execution which can later be used to debug the procedure in case of a failure.

Cloud Control addresses the challenges of the existing Procedure Activity page with its much-improved target-centric procedure management solution that allows access to

all the targets and steps from one single page with maximum ease and minimum time. The new Procedure Activity page offers the following benefits

Overview of the Procedure Instance Execution Page


The Procedure Activity page primarily helps you track the execution of the procedure instance submitted. Click any instance running to view the Procedure Instance Execution page which is broadly divided into the following regions as shown in [Figure 55-2](#):

Note:

Starting with Enterprise Manager 12.1.0.3 (Patch Set 2), the following enhancements have been made to the Procedure Execution Page:

- More filters have been introduced, and filtering of steps is now possible at the table level instead of a submenu in view menu. For details refer to point 6 in the [Figure 55-2](#).
- By default, the step details are now available in a tab layout. However, option is still available to switch to stack view. The benefit of using the tab layout are:
 - a> All the log details are displayed on the screen itself.
 - b> There is a provision to download the log files.
 - c> You can click **Job Summary** link to get more information about the underlying job.

Figure 55-2 Procedure Execution Page



Note:

For information about the tasks that can be performed from the Procedure Instance Execution Page, see [Procedure Instance Execution Page](#).

1. [Breadcrumb Trail](#)
2. [View Data](#)
3. [Procedure Actions Section](#)
4. [Switch To Stack View](#)
5. [Step Details Section](#)
6. [Procedure Steps Section](#)

Breadcrumb Trail

Enables you to go back to the Procedure Activity page with a single click.

View Data

Enables you to refresh the page after making some procedure-level or step-level updates. To do so, you must select the refresh options available in the **View Data** menu. To run the procedure refresh in the background, you can select any of the auto refresh options like: **30 seconds Refresh**, **1 minute Refresh**, or **5 minute Refresh** and continue to work on other areas.

Procedure Actions Section

If the procedure has successfully completed with a status **Succeeded**, then the Procedure Actions menu items are greyed out, you can not perform any actions on a successful procedure. However, if the procedure was stopped or suspended for some reason, then corresponding menu items are enabled so that you can control the procedure execution.

At Procedure-level, you can perform the following actions:

Action	Description
Debug	Debugs the errors in the procedure. This is a one time action, which means that the menu is disabled after using this option the first time.
Stop	Stops the procedure execution.
Suspend	Suspends the procedure execution.
Resume	Resumes the procedure from the stage it was stopped or suspended.
Retry	Executes all the failed steps and phases in the deployment procedure instance once again.
Incident	Creates an incident for the execution which enables you to debug / understand all the steps and phases executed as a part of the deployment procedure.

Switch To Stack View

Select **Switch To Stack View** option to change the view from the default Tab view to Stack view.

The advantages of using the tab view are as follows:

- You can view the log file details on the same page.

- You can access the Job details page to get more information about the underlying job.
- You can download the log files to your local system.

Step Details Section

Enables you to view the details of the selected step. Essentially, information like Step Type, Start date, Completed Date, and Elapsed Time for the step is displayed. However, the template for each step type is not the same. For example, the manual step requires user intervention, which means you might need to confirm some details for the job to proceed. Closing the step window, deselects the step from the Procedure Steps section.

After a step has run, you can click **Job Summary** option to access the job details page to get more information about the underlying job.

Click **Download** option to download the log file for the corresponding step.

If a step has failed, then you can perform the following actions on the step using the Actions menu:

Action	Description
Ignore	Ignore the failure of a step, and continue with the other steps in the deployment procedure.
Retry	Executes the step once again
Update and Retry	Enables you to edit the step, and then executes the step when submitted.

Procedure Steps Section

Enables you to view all the steps that are run when the procedure instance is submitted for execution. From the **View** menu, select **Expand All** to view the step details like: step name, the type of step, and the status of the step.

For example, as displayed in if you select a step called *Check for Supported Configurations*, then its corresponding execution details are displayed in a new tab.

Investigating a Failed Step for a Single or a Set of Targets

Now that the design is target-centric, which means that all the targets and its corresponding steps are listed in the Procedure Steps section, you can select one target or a set of target (multiple select) from the same page to view the status of the step. To do so, perform the followings steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. On the Procedure Activity page, click the procedure name to select the procedure.
3. In the Procedure Steps section, from the **Show** menu, select **Failed Steps**.

All the steps that have failed are displayed in the Procedure Steps section. You can now select the steps that you want to retry, ignore, or update, and the corresponding details are displayed in the Step Details section

Retrying a Failed Step

To retry a failed step, perform the following steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. On the Procedure Activity page, click the procedure name to select the procedure.
3. Select the failed step from the Procedure Steps section. For information on selecting failed steps, see [Investigating a Failed Step for a Single or a Set of Targets](#).

The details of the step are displayed in the Setup Details section.

4. In the Step Details section, from the **Actions** menu, click **Retry**. To make changes to the step, select **Update and Retry** option.
5. In the Retry confirmation dialog box, click **OK** to run the step again.

Creating an Incident

To create an incident for the procedure execution, perform the following steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. On the Procedure Activity page, click the procedure name to select the procedure.
3. In the Procedure Instance Execution page, from the **Procedure Actions** menu, select **Incident**.
4. In the incident confirmation dialog box, click **OK** to create an incident for your execution.

A confirmation dialog box appears once the incident is created. For more information about creating, packaging, and uploading an incident to an SR, see *Oracle Database Administrator's Guide*.

Viewing the Execution Time of a Deployment Procedure

The execution time of the deployment procedure is displayed on top of the page in the Procedure Actions section as **Elapsed Time**. The time elapsed continues to be updated until the procedure has successfully completed or has been stopped. You can resume a stopped procedure by selecting **Resume** from the **Procedure Actions** menu.

Searching for a Step


To search for a step that is embedded deep inside, you can use the **Expand All** option available in the **View** menu. Once the expanded list is displayed in the Procedure Step section, you can easily find the step you are looking for.

Downloading a Step Output

To download a step output, perform the following:

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

2. On the Procedure Activity page, click the procedure name to select the procedure.
3. In the Procedure Steps section, select a step. Click **Download** in the step details section, as shown in the following graphic to download the step output:



Accessing the Job Summary Page

To access the job summary page, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Activity**.
2. On the Procedure Activity page, click the procedure name to select the procedure.
3. In the Procedure Steps section, select a step. Click **Job Summary** link available in the step details section to get more information about the underlying job.

Customizing Deployment Procedures

The Deployment Procedures offered by Oracle Enterprise Manager Cloud Control (Cloud 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:

- [About Deployment Procedure Customization Types](#)
- [Customizing a Deployment Procedure](#)
- [A Workflow Example for Assigning Values to Deployment Procedure Variables at Runtime](#)
- [Changing Deployment Procedure Error Handling Modes](#)
- [Setting Up email Notifications Regarding the Status of Deployment Procedures](#)
- [Copying Customized Provisioning Entities from One Enterprise Manager Site to Another](#)
- [A Workflow Example for Customizing a Directive](#)

About Deployment Procedure Customization Types

The following describes the types of customization you can perform with Deployment Procedures:

Type 1

Editing Custom Deployment Procedures

You can edit an existing custom Deployment Procedure that is offered by Cloud Control to add new phases and steps. However, for patching the steps that can be added are restricted to either a Directive step or a Host command step.

You can perform the following tasks:

- Add your own phases and steps to the pre-defined blocks of the procedure structure.
- Enable and disable phases and steps
- Delete phases and steps
- Change privilege levels
- Change error handling modes
- Enable e-mail notifications

Note: You can not edit an Oracle-owned deployment procedure. To do so, you must clone the Oracle-owned procedure using Create-like functionality, and then edit the copy to include your changes.

Type 2

Creating a User Defined Deployment Procedures

You can create your own Deployment Procedure with new steps, phases, privilege levels, and so on.


You can perform the following tasks:

- Add your own phases and steps to the pre-defined **Default phase** of the procedure structure.
- Enable and disable phases and steps
- Change privilege levels
- Change error handling modes
- Enable e-mail notifications

Note: For steps to Create a User Defined Deployment Procedure, see [Creating, Saving, and Launching User Defined Deployment Procedure \(UDDP\)](#).

The following graphic shows how you can use the Customizing Deployment Procedure page to create a copy of the default Deployment Procedure that is offered by Cloud Control. You can then add new steps and phases or edit the existing steps and phases in the copy to customize your procedure.

For more information on adding steps and phases, see [Customizing a Deployment Procedure](#).



Customizing a Deployment Procedure

The first step towards customizing a Deployment Procedure is to create a copy of the default Deployment Procedure that is offered by Cloud 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.

You can add additional phases or steps to a Deployment Procedure to run additional custom scripts, host commands, or jobs, see [Phases and Steps](#).

 Note:

If a step is added outside a phase, then the type of step that can be added is restricted to a Job Step or a Manual Step. You can not add other steps outside a phase. However, within a phase all the steps discussed in this section can be added.

This section explains how you can edit different types of phases or steps to a Deployment Procedure. In particular, it covers the following:

- [Editing the Rolling and Parallel Phase of a Deployment Procedure](#)
- [Editing a Job Step of a Deployment Procedure](#)
- [Editing a Directive Step of a Deployment Procedure](#)
- [Editing a Component Step of a Deployment Procedure](#)
- [Editing a File Transfer Step of a Deployment Procedure](#)
- [Editing a Host Command Step of a Deployment Procedure](#)
- [Editing a Manual Step of a Deployment Procedure](#)

Editing the Rolling and Parallel Phase of a Deployment Procedure

To edit the general information of the phase, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching** and then select **Procedure Library**.
2. On the Provisioning page, in the Procedure Library tab, from the menu, select **Create New**, then click **Go**.
3. On the Create New Procedure page, select **Procedure Steps** tab.
4. Select the phase, and click **Edit Step** to edit a phase.
5. In the Create wizard, do the following:
 - a. On the Create page, the general information of the phase as displayed. You can change them if you want, and click **Next**.
 - b. On the Select Target List page, select a target list to indicate the type of targets on which the new phase should run.

All the target lists declared while creating the procedure is listed in the drop down menu, select the target list to use for this phase. The actual targets can be chosen when the procedure is being launched.

- c. On the Review page, review the information you have provided for creating a new phase, and click **Finish**. The changes to the phase are saved.

Editing a Job Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. On the Select Type page, you can view the job type details.
3. On the Map Parameters page, you can update the variable values or set a new one. Additionally, you can also update the credentials to be used for the target list.
4. On the Review page, review the updates, and click **Finish**.

Editing a Directive Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. On the Select Directive page, you can select one of the following options:
 - **Select New Directive:** This option lists all the directives available in Software Library, select a directive from the list that you want to run on the targets. Provide necessary values in the Select Directive section to narrow down the search results.
 - **Retain Directive Selection:** This option allows you to use the same directive that you selected while creating this step.
3. On the Map Properties page, review the following details:
 - By default the **Run Directive**, and **Perform Cleanup** options are enabled to run the script, and remove the files after the steps has run. You may disable this is you want.
 - In the Directive Properties section, specify values for the properties associated with the selected directive. You have the option of providing or not providing the property values at this stage. If you do not provide the property values now, then they are prompted at the time of launching the procedure.
 - In the Credentials section, set the target List to be applied for this step.
 - In the Time limit properties section, you can update the max time allowed for an operation to complete in seconds.
4. On the Review page, review the updates, and click **Finish**.

Editing a Component Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. On the Select Component page, you can select one of the following options:

- **Retain Selection:** This option allows you to use the same component that you selected while creating this step.

- **Select New Component:** This option lists all the components available in Software Library, select a component from the list that you want to run on the targets. Provide necessary values in the Select Component section to narrow down the search results.

If you have set a Software Library variable in the Procedure Variables tab, you may additionally notice these two options:

- **Select New Software Library Entity Variable:** Select the Software Library variable that you declared while creating the procedure from the **Software Library Entity Variable** drop down menu. This variable behaves as a place holder and enables you the flexibility of binding it with the components dynamically while launching the procedure. Essentially, you can use this variable to select certain entities which do not need any parameters.

For example, user-defined scripts like a Perl to print the current directory location that does not need any parameters to be passed.

- **Select New Software Library Entity Variable with Component Properties:** This option allows you to bind the Software Library variable with the components that are available in Software Library. Ensure that you choose a component whose properties (signature) matches with the entity declared.

3. On the Select Directive page, you can select one of the following options:

- **Retain Software Library Entity Variable Selection:** This option allows you to use the same directive that you selected while creating this step.

- **Select New Directive:** This option lists all the directives available in Software Library, select a directive from the list that you want to run on the targets. Provide necessary values in the Select Directive section to narrow down the search results.

If you have set a Software Library variable in the Procedure Variables tab, you may additionally notice these two options:

- **Select New Software Library Entity Variable:** Select the Software Library variable that you declared while creating the procedure from the **Software Library Entity Variable** drop down menu. This variable behaves as a place holder and enables you the flexibility of binding it with the directives dynamically while launching the procedure. Essentially, you can use this variable to select certain entities which do not need any parameters.

For example, user-defined scripts like a Perl to print the current directory location that does not need any parameters to be passed.

- **Select New Software Library Entity Variable with Directive Properties:** This option allows you to bind a Software Library entity variable with directives that are available in Software Library. Ensure that you choose a directive whose properties (signature) matches with the entity declared.

4. On the Map Properties page, review the following details:

- By default the **Run Directive**, and **Perform Cleanup** options are enabled to run the script, and remove the files after the steps has run. You may disable this is you want.
 - In the Directive Properties section, specify values for the properties associated with the selected directive. You have the option of providing or not providing the property values at this stage. If you do not provide the propertie values now, then they are prompted at the time of launching the procedure.
 - In the Credentials section, set the target List to be applied for this step.
 - In the Time limit properties section, you can update the max time allowed for an operation to complete in seconds. If there are two options, the interview time will display a radio button. If there are more than two options, then a display menu appears.
5. On the Review page, review the updates, and click **Finish**.

Editing a File Transfer Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. On the Map Properties page, select the Source Target from which you want to transfer files, the source target path, the Target Destination for file transfer and the destination path. Specify the Source and Destination Credential Usage, whether Host or Privileged Host credentials. Click Next.

If you select **Transfer all the files in this path** option, then all the files in the source path are transferred. If uncheck this option, then the **Source File Name** field becomes mandatory.

3. On the Review page, review the updates, and click **Finish**.

Editing a Host Command Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. 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"`.

3. In the Time limit properties section, you can set a max time allowed for a script to run in seconds. For example: Lets assume that you have set this value to 75 seconds, then when the script runs if it exceeds the set time, then this step is skipped.
4. On the Review page, review the updates, and click **Finish**

Editing a Manual Step of a Deployment Procedure

In the Edit wizard, do the following:

1. On the Edit page, review the general information about the step. If you want to update the values, follow the general information about the step as described in [Table 56-1](#).
2. 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.

3. On the Review page, review the updates, and click **Finish**.

Table 56-1 Field Description - Customizing Steps

Field Name	Description
Select	Select <i>Step</i> .
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: Component. • For a file transfer step, select File Transfer • For a manual step, select Manual. • For a host command step, select Host Command.

Table 56-1 (Cont.) Field Description - Customizing Steps

Field Name	Description
Error Handling	Select the error handling mode you want to set for the custom phase. 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 Cloud 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.

A Workflow Example for Assigning Values to Deployment Procedure Variables at Runtime

Starting with Enterprise Manager 12.1.0.3 (Patch Set 2), a new feature has been enabled using command blocks within the PAF functional domain. Command blocks are a set of strings that needs to be printed to the output stream by the script running on a host. They follow this format:

```
$$$--*$$'
<commandBlock>
<executeProc name="MGMT_PAF_UTL.UPDATE_RUNTIME_DATA">
<scalar>%job_execution_id%</scalar>
<scalar>${data.<DP_VARIABLE_NAME>}</scalar>
<scalar><DP_VARIABLE_VALUE></scalar>
</executeProc>
</commandBlock>
$$$*--$$
```

Where, you need to enter the values for <DP_VARIABLE_NAME> and <DP_VARIABLE_VALUE>

Example:

In the following example, **isPingSuccessful** is the DP variable, and **true** is the value.

```
print '$$$--*$$';
print '<commandBlock>';
print '<executeProc name="MGMT_PAF_UTL.UPDATE_RUNTIME_DATA"> ';
print '<scalar>%job_execution_id%</scalar> ';
print '<scalar>${data.isPingSuccessful}</scalar> ';
print '<scalar>true</scalar> ';
print '</executeProc>';
print '</commandBlock>';
print '$$$*--$$';
```


Any language (shell, Perl, and so on) can be used to write the script that will run on the host. As long as the above command block gets printed on the output stream, procedure variables will be updated.

This feature enables you to update the values of procedure declared variables at run-time. The command block feature allows some automatic post-processing of the output generated from a remote machine (agent), which is then uploaded to the OMS, and processed by Job Systems. In PAF, you can use this feature to define steps in a procedure with command block formatted output. Within the command block, SQL procedure can be invoked to assign values to procedure variables.

 Note:

Only a directive step, or a component step, and a job step can use the command block to update the values of a procedure declared variables at run-time.

The following example demonstrates how a command that is run on the target can update the variable in a procedure at runtime. In this example, a ping command is run on a host using the directive step in a procedure. If the ping command succeeds, then the directive sets the value of the procedure to true. If not, the directive sets the value of the procedure to false.



Here are the high level steps:

Table 56-2 Assigning Variable Values at Runtime

Step	Details
Step 1	Create a Perl File called <code>TestPingAndSetDPVariable.pl</code> on your local host. For details, see Step 1: Creating a Perl Script to Assign Values to Deployment Procedure Variables at Runtime .
Step 2	Upload the Perl script to the Software Library. For a detailed list of steps, see Step 2: Uploading TestPingAndDPvariable.pl to Software Library .
Step 3	Create a procedure to add the variable <code>isPingSuccessful</code> , and a few procedure steps which will print the command blocks and call the directive that was created in the previous step. For details, see Step 3: Creating a Deployment Procedure .
Step 4	Run the Procedure. For details, see Step 4: Launching the Deployment Procedure, and Providing the Variable Values at Runtime .
Step 5	Verify the variable details. For details, see Step 5: Verifying the Deployment Procedure Variable Values .

Step 1: Creating a Perl Script to Assign Values to Deployment Procedure Variables at Runtime

Open any editor, copy the following lines of code, then save the file as a *Perl* script:

```
system("ping", "-c", "1", "127.0.0.1");
if ( $? == 0 )
{
    # below 9 lines of print statements set the DP variable "isPingSuccessful" as
    "true" string
    print '$$$--*$$';
    print '<commandBlock>';
    print '<executeProc name="MGMT_PAF_UTL.UPDATE_RUNTIME_DATA"> ';
    print '<scalar>%job_execution_id%</scalar> ';
    print '<scalar>${data.isPingSuccessful}</scalar> ';
    print '<scalar>true</scalar> ';
    print '</executeProc>';
    print '</commandBlock>';
    print '$$$*--$$';
}
else
{
    # below 9 lines of print statements set the DP variable "isPingSuccessful" as
    "false" string
    print '$$$--*$$';
    print '<commandBlock> ';
    print '<executeProc name="MGMT_PAF_UTL.UPDATE_RUNTIME_DATA"> ';
    print '<scalar>%job_execution_id%</scalar> ';
    print '<scalar>${data.isPingSuccessful}</scalar> ';
    print '<scalar>>false</scalar> ';
    print '</executeProc> ';
    print '</commandBlock>';
    print '$$$*--$$ ';
}
}
```

Step 2: Uploading TestPingAndDPvariable.pl to Software Library

To create a directive entity in Software Library, follow these steps:

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library**.
2. On the Software Library home page, right-click any User-owned directive folder. From the context menu, select **Create Entity**, then click **Directive**. The Create Directive page is displayed.
3. On the Describe Page, enter a unique name for the folder. For example, TestPingAndSetDPVariable. Click **Next**.
4. On the Configure page, in the Configure Properties section, from the Shell Type list, select **Perl**. Click **Next**.
5. On the Select Files page, in the Specify Destination section, select any Software Library Upload Location for uploading the specified files.

In the Specify Source section, select File Source as **Local Machine**. Click **Add**, and select the Perl script TestPingAndDPVariable.pl from your local host. Click **Next**.

6. Review all the details, and click **Save and Upload**.

See Also:

For more information, see *Enterprise Manager Cloud Control Administrator's Guide*.

Step 3: Creating a Deployment Procedure

To create a procedure, follow these steps:

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then click **Procedure Library**.
2. On the Procedure Library page, from the menu, select **Create New**, then click **Go**.
3. On the General Information page, enter a unique name for your procedure.
4. On the Create New Procedure page, click **Procedure Variables**. Click Add Row, and enter the following details:

General Information Target Lists Procedure Variables Procedure Steps						
Delete Row Add Row						
Select	Name	Display Name	Description	Type	Value Options	Required
☐	isPingSuccessful	Is Ping Successful Variable		String	Text	☐

5. Click Procedure Steps tab. Select the **Default phase**, then click **Insert** to add the following steps to the default phase:
 - a. Add a Host Command step to print the value of the variable before running the procedure. Here are the steps: On the Create page, select **Host Command** from the menu, enter a unique name for the step, then click **Next**. On the Enter Command page, enter the value: echo Before String is: "\${data.isPingSuccessful}".

Note: The expression to access the DP variable isPingSuccessful in DP code is \${data.isPingSuccessful}.

 Click **Next**. Review the details, click **Finish**.
 - b. Add a Directive step called ping to call the directive that was uploaded to Software Library. To do so, on the Create page, enter a unique name for the step, then select **Directive** from the Type menu. On the Select Directive page, search with a string %Ping%. Select the Directive that you uploaded, and click **Next**. You can leave all the default values as is on Map Properties page, and Review page, then click **Finish**.
 - c. Add a Host Command step to print the value of the variable after running the procedure. Here are the steps: On the Create page, select **Host Command** from the menu, enter a unique name for the step, then click **Next**. On the Enter Command page, enter the value: echo After String is: "\${data.isPingSuccessful}". Click **Next**. Review the details, click **Finish**.
6. Click **Save and Close**.

Step 4: Launching the Deployment Procedure, and Providing the Variable Values at Runtime

Run the procedure as follows:

1. In Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then click **Procedure Library**.
2. On the Procedure Library page, select the procedure you created in the previous step, then click **Launch**. The Procedure Interview page is displayed.
3. On the Select Targets page, select a target on which you want to run this procedure. Click **Next**.
4. On the Set variables Page, for the `isPingSuccessful` variable, enter the following value:

This is the Value I Entered.

Click **Next**.

5. On the Credentials page, set the credentials for the targets.
6. On the Schedule page, select **Immediately**, and click **Submit**.
7. Enter the **Submission Name**, then click **Ok**.

Step 5: Verifying the Deployment Procedure Variable Values

To verify the variable values at different stages of the procedure submission, follow these steps:

1. On the Procedure Activity page, select the procedure instance that you submitted.
2. On the Procedure Instance Execution page, select the Step **Print Before**. This step prints the value of the variable entered during submission.

The screenshot displays the Oracle Cloud Control Provisioning interface. The top navigation bar shows 'Provisioning' and 'Procedure Activity > rtrt'. The main section is titled 'Procedure Activity: rtrt' and shows 'Elapsed Time: 36 seconds'. Below this, the 'Procedure Steps' table lists the following steps: 'Default Phase', 'adc5260288.us.oracle.com', 'Print Before', 'ping', and 'Print After'. The 'Print Before' step is selected, and its details are shown on the right. The 'Print Before' step is a 'Host Command' that executed successfully. The 'Before String' is 'This is the value I entered'.

Select Name	Status
Default Phase	✓
adc5260288.us.oracle.com	✓
Print Before	✓
ping	✓
Print After	✓

Print Before ×

✓ **Print Before** Job Summary Download Actions

Type Host Command Start Date Feb 4, 2014 11:00:49 PM
Elapsed Time 5 seconds UTC
Completed Date Feb 4, 2014 11:00:54 PM
Date UTC

Step: Execute Command (Succeeded)

Start Date 2014.02.04 23:00:51
Completed Date 2014.02.04 23:00:52
Targets adc5260288.us.oracle.com

Before String is: This is the value I entered

The output that is displayed in the Execute Command block is:

Before String is: **This is the Value I entered.**

3. On the Procedure Instance Execution page, select the Step **Ping**.

Note:

The **Ping** step ran, and the ping command was successful. However, the print statements printing the command block string in the Perl script are not displayed. This is because, these special command block strings were parsed by job system and used to process the procedure variable assignment logic.

Procedure Activity > rtrt

Procedure Activity: rtrt

Elapsed Time: 36 seconds

Procedure Steps

Select Name	Status
Default Phase	✓
adc6260288.us.oracle.com	✓
Print Before	✓
ping	✓
Print After	✓

View: All Steps

ping ×

Step: Execute Command (Succeeded)

Start Date: 2014.02.04 23:01:03
Completed Date: 2014.02.04 23:01:05
Targets: adc6260288.us.oracle.com

This is Provisioning Executor Script

```
Input PropertyFilePath is /tmp/JOB_F19CE88E964F4490E0436D32F20AF8AF/properties
Input directiveScriptFilePath is /tmp/JOB_F19CE88E964F4490E0436D32F20AF8AF/TestPingAndSet
Input directiveTypeProperty is NAME_Executor
Input componentPath is Component
Input directivePath is Directive
Input componentFilename is /tmp/JOB_F19CE88E964F4490E0436D32F20AF8AF/null

Directive Type is SUB_Perl
Dir to CD is /tmp/JOB_F19CE88E964F4490E0436D32F20AF8AF

Final command line to execute is $PERL_HOME/perl "/tmp/JOB_F19CE88E964F4490E0436D32F20AF8AF

The output of the directive is:
PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data:
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=0.022 ms

--- 127.0.0.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.022/0.022/0.022/0.000 ms
```

Step: Set Parameters (Succeeded)

- On the Procedure Instance Execution page, select Step **Print After**. This step prints the value of the variable as true, this was set by the directive that ran on the host.

Provisioning

Procedure Activity > rtrt

Procedure Activity: rtrt

Elapsed Time: 36 seconds

Procedure Steps

Select Name	Status
Default Phase	✓
adc6260288.us.oracle.com	✓
Print Before	✓
ping	✓
Print After	✓

View: All Steps

Print After ×

Step: Execute Command (Succeeded)

Start Date: 2014.02.04 23:01:16
Completed Date: 2014.02.04 23:01:18
Targets: adc6260288.us.oracle.com

After String is: true

The output that is displayed in the Execute Command block is:

After String is: true

This way you can use a directive step to manipulate procedure variables at run-time depending on the output returned by some other command that ran on the target.

Changing Deployment Procedure 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 Cloud 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 [Phases and Steps](#).

To change the error handling modes, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, select the Deployment Procedure you want to customize, and click **Create Like**.
3. On the Create Like Procedure page, click the **Procedure Steps** tab, and then select a phase or step, and change the **Error Handling Mode**.

After the mode is selected from the list, Cloud 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:

Deployment Procedure Manager

Procedure Library | Procedure Activity | Recycle Bin

Procedures are best practices provided by Oracle for various Provisioning and Patching tasks. Procedures created by Oracle cannot be edited, but can be extended using 'Create Like', so that you can customize the procedure to fit your environment.

Search Text Fields Go Advanced Search

Launch Edit Procedure Definition... Create Like Launch Previous 1-25 of 40 Next 15

Select	Procedure	Type	Parent	Version	Last Updated	Description	Last Modified By	Owner
☐	Patch Oracle Fusion Application	Patch Oracle Software	None	1.0	Nov 13, 2011 1:03:08 PM UTC	Procedure for patching Oracle Fusion Applications	Oracle	Oracle
☐	Clone and Patch Oracle Database	Patch Oracle Software	None	12.1	Nov 13, 2011 1:00:59 PM UTC	Procedure for automatically cloning a standalone database (single instance database) and patching it with critical patch updates, interim patches, or patchsets. Note: Major upgrade for example, DB 10.1 to 10.2 is not supported.	Oracle	Oracle
☒	Provision Oracle Database Client	Provisioning	None	6.1	Nov 13, 2011 1:00:59 PM UTC	This procedure installs or clones a database replay client on the selected hosts. It follows the best practices specified by the Oracle Database Installation Guide.	Oracle	Oracle
☐	Upgrade Oracle Database	Upgrade	None			This procedure upgrades a single instance Oracle database to a higher version.		

Provisioning

Select	Procedure	Type	Description	Actions
☐	Initialize Deployment Procedure	Computational	Initializes the deployment procedure. Do NOT disable or delete this step.	Inherit (Stop On Error)
☐	Transfer Common Perl Module	Parallel	Objects that are shared by all steps are transferred to target hosts.	Inherit (Stop On Error)
☐	Transfer Common Perl Module	Component	Transfer a perl module from Software Library to each destination host. All procedure steps that invoke perl will consume this module.	Inherit (Stop On Error) Run using sudo
☐	Destination User Permission Checks	Parallel	Runs necessary user permission-related prerequisite checks on the destination hosts using Oracle home user credentials.	Inherit (Stop On Error)
☒	Check if Staging Area Path is writable	Host Command	Checks whether the staging area is writable.	Inherit (Stop On Error) Normal
☐	Check sudo permission	Host Command	Checks sudo permission.	Inherit (Stop On Error) Stop On Error Run using sudo
☐	Check if disk is usable	Host Command	Checks whether the disk is usable.	Inherit (Stop On Error) Skip Target Normal
☐	Prepare Create directory	Directive	Prepares a list of first-level directories to be created on the destination hosts. Do NOT disable or delete this step.	Inherit (Stop On Error) Normal
☐	Agent User Permission Checks	Parallel	Runs the necessary user permission-related prerequisite checks on the destination hosts using Agent user credentials.	Inherit (Stop On Error)

Setting Up email Notifications Regarding the Status of Deployment Procedures

Cloud Control can send email notifications to report the status of a Deployment Procedure. However, by default, Deployment Procedures do not have this feature enabled. For email notifications to be sent, your Super Administrator must configure an Outgoing Mail (SMTP) Server within Enterprise Manager, and then you as an Administrator must provide your email Address and Password for receiving notifications.

Enabling email notifications is a two-step process:

- [Configuring an Outgoing Mail \(SMTP\) Server In Enterprise Manager](#)
- [Adding E-mail Addresses for Enterprise Manager Notifications](#)

Note:

As a prerequisite, you are expected to have configured the Mail Server and set up the email address in Cloud Control.

Configuring an Outgoing Mail (SMTP) Server In Enterprise Manager

Before Enterprise Manager can send e-mail notifications, you must first set up the Outgoing Mail (SMTP) servers.



Note:

Only a privileged user can configure SMTP servers.

To set up the SMTP server, follow these steps:

1. In Cloud Control, from the **Setup** menu, select **Notification**, then select **Notifications Method**.
2. On the Setup page, in the Mail Server section, enter one or more outgoing mail servers and optional port number (if left blank, port 25 is used by default).
3. Enter the mail server authentication credentials like **UserName**, **Password**. The UserName and Password fields allow you to specify a single set of authentication credentials to be used for all mail servers. If no mail server authentication is required, leave the User Name, Password (and Confirm Password) fields blank.
4. Enter the name you want to appear as in the sender of the notification messages field **Identify Sender As**.
5. Enter the e-mail address you want to use to send your e-mail notifications in the **Sender's E-mail Address**. When using incident rules, any e-mail delivery problems will be automatically sent to the **Sender's E-mail Address**.



Note:

The e-mail address you specify on this page is not the e-mail address to which the notification is sent. You will have to specify the e-mail address (where notifications will be sent) from the Preferences General page, see [Adding E-mail Addresses for Enterprise Manager Notifications](#).

6. The Use Secure Connection option allows you to choose the SMTP encryption method to be used. Three options are provided:
 - **No:** E-mail is not encrypted.
 - **SSL:** E-mail is encrypted using the Secure Sockets Layer protocol.
 - **TLS, if available:** E-mail is encrypted using the Transport Layer Security protocol if the mail server supports TLS. If the server does not support TLS, the e-mail is automatically sent as plain text.

In the following example, two mail servers are specified--smtp01.example.com on port 587 and smtp02.example.com on port 25 (default port). A single administrator account (myadmin) is used for both servers.

Outgoing Mail (SMTP) Server smtp01.example.com:587, smtp02.example.com

User Name myadmin

Password *******Confirm Password** *******Identify Sender As** EMD Notifications**Sender's E-mail Address** mgmt_rep@example.com**Use Secure Connection:** SSL

Adding E-mail Addresses for Enterprise Manager Notifications

You specify one or more e-mail addresses to which notifications can be sent when you define the Notifications Schedule. In addition to defining notification e-mail addresses, you associate the notification message format (long or short) to be used for each e-mail address.

Each e-mail address can have up to 128 characters; there is no upper bound with the number of e-mail addresses.

To add an e-mail address:

1. In Cloud Control, from the **Setup** menu, select **Notification**, then select **My Notifications Schedule**.
2. Specify an Enterprise Manager administrator, and click **Define Schedule**.
3. If no previous e-mail addresses have been defined for the administrator, a message displays prompting your to define e-mail addresses for the administrator. Click **Click here to set e-mail addresses**. The General page appears.
4. Click **Add Another Row** to create a new e-mail entry field in the E-mail Addresses table.
5. Specify the e-mail address associated with your Enterprise Manager account. All e-mail notifications you receive from Enterprise Manager will be sent to the e-mail addresses you specify. For example, user1@example.com, user2@example.com, and so on.
6. If you need to additional e-mail addresses, click **Add Another Row**, enter the e-mail address and select the format.
7. You can test if the e-mail address is properly configured to receive e-mails from Enterprise Manager by selecting it and clicking **Test**.
8. Click **Apply** to save your changes when finished.

Once you have defined your e-mail notification addresses, they will be shown when you define a notification schedule. For example, user1@example.com, user2@example.com, user3@example.com. You can choose to use one or more of these e-mail addresses to which e-mail notifications for the Incident Rule will be sent.

Copying Customized Provisioning Entities from One Enterprise Manager Site to Another

If you have customized provisioning entities on an Enterprise Manager installation that you want to apply to another installation of Enterprise Manager, follow these steps. The provisioning entities can include procedure definitions or Software Library entities or a combination of these.

This section consists of the following:

- [Prerequisites for Copying Customized Provisioning Entities from One Enterprise Manager Site to Another](#)
- [Copying Customized Provisioning Entities from One Enterprise Manager Site to Another](#)

Prerequisites for Copying Customized Provisioning Entities from One Enterprise Manager Site to Another

Ensure that:

- Customized provisioning entities exist in the system at source site.
- Source administrator has access to the customized provisioning entities at source site.
- Source administrator has necessary privileges to export the provisioning entities.
- Destination site has similar setup as source site, that is, both have the same version of Enterprise Manager installed.
- Destination administrator has privileges to import provisioning entities.

Copying Customized Provisioning Entities from One Enterprise Manager Site to Another

Follow these steps:

1. Export the PAR file using the following command:

```
emctl partool export -guid <procedure guid> -file <file> -displayName <name> -description <desc> -metadataOnly(optional)
```
2. For importing the provisioning entities, you can use emctl partool as follows:

```
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)
```
3. Alternatively, you can import the PAR file from Cloud Control as explained in the following steps.
 - a. From the **Enterprise** menu, select **Provisioning and Patching** and then select **Procedure Library**.
 - b. On the Procedure Library page, from the list, select **Import** and click **Go**.
4. In the Upload Procedure File page, select:
 - Upload From Local Machine, if the PAR files are stored on your local machine. Click **Browse** and select the PAR File to Upload. Click **Import**.
 - Upload From Management Agent Machine, if you have stored the PAR files on the Management Agent machine. Click **Target** and select the Host. Click **Select File** and select the PAR file. Click **Import**.
5. Apply the imported entities.

A Workflow Example for Customizing a Directive

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 adding a Directive Step, see [Adding a Directive Step](#).

If you want to customize a directive offered by Cloud 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 and Uploading a Copy of a Default Directive](#)
- [Customizing a Deployment Procedure to Use the New Directive](#)
- [Running the Customized Deployment Procedure](#)

Creating and Uploading a Copy of a Default Directive

To create a new customized directive using a default directive, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library page, from the table, expand the Software Library and the 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 **Software Library** and then expand **Patching**. From this level, expand **Common**, and then **All**, and finally **Generic**. Under Generic, you should see the directive **Apply Patch**.

3. Select the directive you want to copy and click **Create Like**, and store the copy of the directive in a custom folder called **Directive**.
4. Select the custom directive, and then click **Edit**.
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 Select Files page, select **Upload Files**. In the Specify Source section, select **Local Machine**, and click **Add** to select the modified perl file.
- d. Click **Save and Upload**.

Customizing a Deployment Procedure to Use the New Directive

To customize a Deployment Procedure to use the new directive, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, select the Deployment Procedure for which you want to use this new directive, and click **Create Like**.
3. On the Create Like Procedure page, select **Procedure Steps** tab, and 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:
 - i. On the Edit page, click **Next**.
 - ii. On the Select Directive page, select **Select New Directive**. Then search and select the new directive you created, and click **Next**.
 - iii. On the Map Properties page, specify the values for the directive properties, and click **Next**.
 - iv. On the Review page, click **Finish**.
 - c. Click **Save**.

Running the Customized Deployment Procedure

To run the customized Deployment Procedure, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, select the customized Deployment Procedure and click **Schedule Deployment**.

Part XIII

Additional Information

This part contains the following appendixes:

- [Using Enterprise Manager Command Line Interface](#)
- [Checking Host Readiness Before Provisioning or Patching](#)
- [Using emctl partool Utility](#)
- [Understanding PXE Booting and Kickstart Technology](#)
- [Troubleshooting Issues](#)

A

Using Enterprise Manager Command Line Interface

This appendix explains how to use Enterprise Manager Command Line Interface (EM CLI) to deploy patches using Patch Plans, provision procedures, and perform some of the Software Library tasks.

Note:

For information about Enterprise Manager 13c verb usage, syntax, and examples, see *Oracle Enterprise Manager Command Line Interface*

In particular, this appendix covers the following:

- [Overview](#)
- [Prerequisites](#)
- [Enterprise Manager Command Line Interface Verbs](#)
- [Provisioning Using EM CLI](#)
- [Patching Using EM CLI](#)
- [WorkFlow Examples Using EM CLI Commands](#)
- [Limitations of Using Enterprise Manager Command Line Interface](#)

Note:

The entire EM CLI implementation for running the various Deployment Procedures has been revamped in Oracle Enterprise Manager Cloud Control (Cloud Control).

Overview

Enterprise Manager Command Line Interface (EM CLI) is a command line utility available for power users in Oracle Enterprise Manager Cloud Control (Cloud Control) that enables you to perform most of the console-based operations. It enables you to access Cloud Control functionality from text-based consoles (shells and command windows) for a variety of operating systems.

Using EM CLI you can:

- Perform various command line operations such as monitoring and managing targets, jobs, running deployment procedures, patching Enterprise Manager targets, and so on.
- Use the functions available with EM CLI called *verbs*, to build your custom scripts on various programming environments like Operating System shell, Perl, Python, and so on. This in turn allows you to closely integrate Oracle Enterprise Manager functionality with your own enterprise business process.
- Carry out operations with the same security and confidentiality as the Cloud Control console.

Prerequisites

Before using EM CLI, ensure that you meet the following requirements:

- EM CLI client must be set up. To do so, see *Oracle Enterprise Manager Command Line Interface*.
- Targets that will be supplied to the Deployment Procedures should be managed by Enterprise Management 13c Management Agents.
- If you are patching in the offline mode, with no internet connectivity, then ensure that the patches are available in the Software Library before running the EM CLI commands.

Enterprise Manager Command Line Interface Verbs

This section primarily lists all the EM CLI verbs used for accomplishing the various patching and provisioning tasks. Primarily, it contains:

- [Provisioning EM CLI Verbs](#)
- [Patching EM CLI Verbs](#)
- [Software Library EM CLI Verbs](#)

Provisioning EM CLI Verbs

The table below contains a list of all the EM CLI verbs used for running deployment procedures:

Table A-1 EM CLI Provisioning Verbs and their Usage

Verb	Usage	Example
confirm_instance	emcli confirm_instance [-instance={instance guid}] [-exec={execution guid}] -stateguid={state guid}	emcli confirm_instance -instance=234RTGHJ096YHN5KM2IKJM567 -state-guid=56IUJMN029IJ3ERFG09IKJ

Table A-1 (Cont.) EM CLI Provisioning Verbs and their Usage

Verb	Usage	Example
describe_procedure_input	emcli describe_procedure_input [-procedure={procedure GUID}] [-name={procedure name or procedure configuration}] [-owner={owner of the procedure or procedure configuration}] [-parent_proc={procedure of the procedure configuration. this only applies to describe a procedure configuration with the same name}]	emcli get_procedure_xml -procedure=16B15CB29C3F9E6CE040578C96093F61 > proc.properties
delete_instance	emcli delete_instance [-instance={instance guid}] [-exec={execution guid}]	emcli delete_instance -instance=234RTGHJ096YHN5KM2IKJM567
get_executions	emcli get_executions -instance={instance GUID}	emcli get_exections -instance=16B15CB29C3F9E6CE040578C96093F61
get_instances	emcli get_instances [-type={procedure type}]	emcli get_instances -type=DemoNG
get_instance_data	emcli get_instance_data [-instance={instance guid}] [-exec=execution guid]	emcli get_instance_data -instance=16B15CB29C3F9E6CE040578C96093F61 > instanceData.properties
get_instance_status	emcli get_instance_status [-instance={instance guid}] [-exec=execution guid] [-xml [-details] [-showJobOutput] [-tailLength={last N characters}]]]	emcli get_instance_status -instance=1TYUIOPLKMUHKJANGS090IJ -xml -details -showJobOutput
get_retry_argument	emcli get_retry_arguments [-instance={instance guid}] [-exec=execution guid] [-stateguid={state guid}]	emcli get_retry_argument -instance=16B15CB29C3F9E6CE040578C96093F61 -stateguid=4IUOH-NAG29KLNLOKJGA
get_procedures	emcli get_procedures [-type={procedure type}] [-parent_proc={procedure associate with procedure configuration}]	emcli get_procedures -type=DemoNG -parent_proc=ComputeStepTest
get_procedure_xml	emcli get_procedure_xml [-procedure={procedure guid}] [-name={procedure name}] [-owner={procedure owner}]	emcli get_procedure_xml -procedure=16B15CB29C3F9E6CE040578C96093F61 > proc.xml
get_procedure_types	emcli get_procedure_types	emcli get_procedure_types
ignore_instance	emcli ignore_instance [-instance={instance guid}] [-exec=execution guid] [-stateguid={state guid}]	emcli get_retry_argument -instance=16B15CB29C3F9E6CE040578C96093F61 -stateguid=4IUOH-NAG29KLNLOKJGA, 29C3F9E6CE040578C96093F61KNALK

Table A-1 (Cont.) EM CLI Provisioning Verbs and their Usage

Verb	Usage	Example
reschedule_instance	emcli reschedule_instance [-instance={instance guid}] [-exec=execution guid] -schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone id};	emcli reschedule_instance -instance=1TYUIOPLKMUHKJANGS090IJ -schedule="start_time:2012/12/25 00:00;tz:American/New York;grace_period:60"
resume_instance	emcli resume_instance [-instance={instance guid}] [-exec=execution guid]	emcli resume_instance -instance=1TYUIOPLKMUHKJANGS090IJ
save_procedure	emcli save_procedure_input -name={name of procedure configuration} -procedure={Procedure name} [-owner={owner of procedure}] -input_file=data:/file path/file name [-grants={users and their corresponding accessing levels}] [-notification={procedure status}] [-schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone id};grace_period:xxx]	emcli save_procedure_input -name=procConfiguration -procedure=ComputeStepTest -input_file=data:/home/data.properties -grants="user1:VIEW_JOB; user2:FULL_JOB" -notification="sheduled, action required, running" -schedule="start_time:2012/12/25 00:00;tz:American/New York;grace_period:60"
stop_instance	emcli stop_instance [-instance={instance guid}] [-exec=execution guid]	emcli stop_instance -instance=1TYUIOPLKMUHKJANGS090IJ
submit_procedure	emcli submit_procedure [-name={name of the procedure}] [-owner={owner of the procedure}] [-procedure={guid of the procedure}] -input_file={data:{file_path}/file name" [-instance_name={name for the procedure instance}] [-notification={procedure status}] [-grants={users and their corresponding accessing levels}] [-schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone ID}]	emcli submit_procedure -input_file=data:data.xml -procedure=16B15CB29C3F9E6CE040578C96093F61 -schedule="start_time:2006/6/21 21:23; tz:America/New_York" -grants="user1:VIEW_JOB; user2:FULL_JOB" -notification="sheduled, action required, running"
suspend_instance	emcli stop_instance [-instance={instance guid}] [-exec=execution guid]	emcli suspend_instance -instance=1TYUIOPLKMUHKJANGS090IJ
update_and_retry_step	emcli update_and_retry_step [-instance={instance guid}] [-exec=execution guid] [-stateguid={stateguid1, stateguid2, ...}] [-args="command1:value1;command2:value2;...]	emcli get_retry_argument -instance=16B15CB29C3F9E6CE040578C96093F61 -stateguid="4IUOH-NAG29KLNLOKJGA,P082NLKBSAKBNIU-POQTG" -args="command1:a; command2:b"

Table A-1 (Cont.) EM CLI Provisioning Verbs and their Usage

Verb	Usage	Example
update_procedure_input	emcli update_procedure_input -name={name of procedure configuration} -input_file="data:/file path/file name" [-notification={procedure status}][grants={users and their corresponding accessing levels}] [-schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone id};grace_period:xxx]	emcli update_procedure_input -name=procConfiguration -input_file=data:/home/data.properties -grants="user1:VIEW_JOB;user2:FULL_JOB" -notification="sheduled, action required, running" -schedule="start_time:2012/12/25 00:00;tz:American/New York;grace_period:60"
switch_swlib_oms_agent_storage	emcli switch_swlib_oms_agent_storage -name="location_name"-host="hostname" [-credential_set_name="setname"] [-credential_name="name" -credential_owner="owner"]	emcli switch_swlib_oms_agent_storage -name="myOMSAGtLocation" -host="fs1.us.example.com" -credential_name="MyAcmeCreds" -credential_owner="ACME_USER" Modifies the OMS Agent File-system storage location named 'myOMSAGtLocation' to use the specified host 'fs1.us.example.com', and the named credential 'MyAcmeCreds' owned by 'ACME_USER' for reading/writing files from/to this location.
verify_swlib	emcli verify_swlib [-report_type="storage entity upload-jobs all"]	emcli verify_swlib -report_type="all" Generates the storage, upload jobs and entities verification reports.

Patching EM CLI Verbs

Support for patching any number of targets from the same single patch plan using command line interface is now supported. Following are some of the important EM CLI verbs used for patching.

Table A-2 EM CLI Patching Verbs and their Usage

Verbs	Usage	Example
create_patch_plan	emcli create_patch_plan -name="name" -input_file=data:"file_path" [-impact_other_targets="add_all add_original_only cancel"] [-problems_assoc_patches="ignore_all_warnings cancel"]	emcli create_patch_plan -name="plan name" -input_file=data:"/tmp/patch-plan.pros" -impact_other_targets="add_all"
describe_patch_plan_input	emcli describe_patch_plan_input -name="name"	emcli describe_patch_plan_input -name="plan_name"
get_patch_plan_data	emcli get_patch_plan_data -name="name"	emcli get_patch_plan_data -name="plan_name"
set_patch_plan_data	emcli set_patch_plan_data -name="name" [-impact_other_targets="add_all add_original_only cancel"] [-problems_assoc_patches="ignore_all_warnings cancel"]	emcli set_patch_plan_data -name="plan name" -input_file=data:"/tmp/patch-plan.pros" -impact_other_targets="add_all"
list_aru_languages	emcli list_aru_languages [-name="language name" -id="language id"] [-noheader] [-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli list_aru_languages -noheader
list_aru_platforms	emcli list_aru_platforms [-name="platform name" -id="platform id"] [-noheader] [-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli list_aru_platforms -noheader
list_aru_products	emcli list_aru_products [-name="product name" -id="product id"] [-noheader] [-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli list_aru_products -id="product id"
list_aru_releases	emcli list_aru_releases [-name="release name" -id="release id" -productId="product id"] [-noheader] [-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli list_aru_releases -noheader
list_patch_plans	emcli list_patch_plans -name="name" [-noheader] [-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli list_patch_plans -name="plan name" -noheader

Table A-2 (Cont.) EM CLI Patching Verbs and their Usage

Verbs	Usage	Example
search_patches	emcli search_patches [-swlib] [-patch_name="patch_name"] [-product="product id"] [-include_all_products_in_family]] [-release="release id"] [-platform="platform id" -language="language id"] [-type="patch patchset"] [-noheader] [-script -xml -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli search_patches -patch_name="patch number" -platform="platform id"
get_connection_mode	emcli get_connection_mode	emcli get_connection_mode
set_connection_mode	emcli set_connection_mode -mode="online offline"	emcli set_connection_mode -mode="offline"
show_patch_plan	emcli show_patch_plan -name="name" [-info [-showPrivs]] [-actions [-onlyShowEnabled]] [-patches] [-targets] [-deplOptions] [-analysisResults][-conflictFree][-impactedTargets][-deploymentProcedures]	emcli show_patch_plan -name="plan name" -info -showPrivs
submit_patch_plan	emcli submit_patch_plan -name="name" -action="action name"	emcli submit_patch_plan -name="plan name" -action="analyze"
get_targets	emcli get_targets [-targets="[name1:]type1; [name2:]type2;..."] [-alerts][-noheader][-script -format=[name:<pretty script csv>]; [column_separator:"column_sep_string"]; [row_separator:"row_sep_string"];]	emcli get_targets -targets="databa%:%oracle%"
get_instance_status	emcli get_instance_status [-instance={instance_guid}] [-exec={execution_guid}] [-name={execution name}] [-owner={execution owner}] [-xml [-details]	emcli get_instance_status -instance=16B15CB29C3F9E6CE040578C96093F61 -xml -showJobOutput -tailLength=1024
get_job_execution_detail	emcli get_job_execution_detail -execution={execution_id}[-xml [-showOutput [-tailLength={length}]]]	emcli get_job_execution_detail -execution=12345678901234567890123456789012 -xml

Table A-2 (Cont.) EM CLI Patching Verbs and their Usage

Verbs	Usage	Example
create_named_creden- tial	emcli create_named_credential - cred_name=<name> -auth_tar- get_type=<authenticating target type> -cred_type=<Credential type> - cred_scope=<Credential Scope> - cred_desc=<Credential Description> - target_name=<target name> -tar- get_type=<target type> -in- put_file=<tag:value> -in- put_bfile=<tag:value> -proper- ties_file=<filename> -attrib- utes=<p1:v1;p2:v2;...>	emcli create_named_credential- cred_name=NC1 -auth_tar- get_type=host-cred_type=Host- Creds -attributes="HostUser- Name:foo;HostPassword:"
get_named_credential	emcli get_named_credential - cred_owner=<owner> - cred_name=<name>-out=<filename>	emcli get_named_credential - cred_name=NC1
set_preferred_cre- dential	emcli set_preferred_credential - set_name="set_name" -tar- get_name="target_name" -tar- get_type="ttype" -creden- tial_name="cred_name" [-creden- tial_owner ="owner"]	emcli set_preferred_credential -target_type=oracle_database - target_name=myDB - set_name=DBCredsSYSDBA -cre- dential_name=MyDBCredentials - credential_owner="Joe"
show_creden- tial_set_info	emcli show_credential_type_info [- target_type="target_type"] [- type_name="credential_type_name"]	emcli show_credential_type_in- fo -target_type=oracle_database
setup	emcli setup - url="http[s]:// host:port/em/" [-username=<EM Con- sole Username>] [-password=<EM Con- sole Password>][-licans=YES NO][- dir=<local emcli configuration di- rectory>][-trustall][-novalidate][- noautologin][-custom_at- trib_file=<Custom attribute file path>][-nocertvalidate]	emcli setup -url=https://dadv- mi0128.us.example.com:4473/em -username=sysman -pass- word=sysman
upload_patches	emcli upload_patches - from_host="host name" - patch_files="metadata file path;ZIP file path" [-cred_name="name" - cred_owner="owner"]	emcli upload_patches - patch_files="/scratch/ p13741363_112310_Linux- x86-64_M.xml;/scratch/ p13741363_112310_Linux- x86-64.zip" -from_host=abc.ex- ample.com
delete_patches	emcli delete_patches - patch_name="patch name" -re- lease="release id" -platform= "plat- form id"	emcli delete_patches - patch_name=13741363 -re- lease=80112310 -platform=226

Software Library EM CLI Verbs

Configuring Software Library, creating entities, and using them is supported in Oracle Enterprise Manager Cloud Control.

Note:

You can either use Enterprise Manager UI or the command line utility (EM CLI) to retrieve the folder id and the entity revision id. To do so, and for a comprehensive example on how to effectively use the EM CLI verbs to perform a number of Software Library tasks listed in the following table, see the workflow example [Creating a New Generic Component by Associating a Zip File](#).

Following are some of the important EM CLI verbs used to perform some Software Library actions:

Table A-3 Software Library EM CLI Verbs and Their Usage

Verb	Usage	Example
add_swlib_storage_location (Adding a Software Library storage location)	emcli add_swlib_storage_location -name="location_name" -path="location_path" [-type="OmsShared OmsAgent Http Nfs ExtAgent"] [-host="hostname"] [-credential_set_name="setname"] [-credential_name="name" -credential_owner="owner"]	emcli add_swlib_storage_location -name="myOMSAgtLocation" -path="/u01/swlib" -type="OmsAgent" -host="fs1.us.example.com" -credential_name="Myexample-Creds" -credential_owner="example_USER"
create_swlib_folder (Creating a Software Library folder)	emcli create_swlib_folder -name="folder_name" -parent_id="parent folder id" [-desc="folder description"]	emcli create_swlib_folder -name="myFolder" -parent_id="oracle:defaultService:em:provisioning:1:cat:B13B3B7B086458CFE040E80A19AA560C" -desc="myFolder description"
create_swlib_entity (Creating a Software Library entity)	emcli create_swlib_entity -name="entity_name" -folder_id="folder_id" [-type]="type internal id" [-subtype]="subtype internal id" [-desc="entity_desc"] [-attr="<attr name>:<attr value>"] [-prop="<prop name>:<prop value>"] [-secret_prop="<secret prop name>:<secret prop value>"] [-note="note text"]	emcli create_swlib_entity -name="myexampleInstall" -folder_id="oracle:defaultService:em:provisioning:1:cat:B13B3B7B086458CFE040E80A19AA560C" -desc="myexampleInstall description" -attr="PRODUCT:example" -attr="PRODUCT_VERSION:3.0" -attr="VENDOR:example Corp" -prop="DEFAULT_HOME:/u01/example3/" -note="myexampleInstall for test servers"

Table A-3 (Cont.) Software Library EM CLI Verbs and Their Usage

Verb	Usage	Example
list_swlib_entities (Listing the Software Library entities)	emcli list_swlib_entities [-name="entity_name"] [-folder_id="folder internal id"] [-desc="entity_desc"] [-attr="<attr name>:<attr value>"] [-type="type internal id"] [-subtype="subtype internal id"] [-maturity="maturity"] [-owner="owner"] [-status="status"] [-show_folder_path] [-show_folder_id] [-show_entity_rev_id]	emcli list_swlib_entities -name="myEntity" -attr="PROD-UCT=Oracle Database" -show_folder_id
list_swlib_entity_subtypes (Listing Software Library entity subtypes)	emcli list_swlib_entity_subtypes -entity_type_id="type internal name" [-show_subtype_id]	emcli list_swlib_entity_subtypes -entity_type_id="COMP_Component" -show_type_id
list_swlib_entity_types (Listing Software Library entity types)	emcli list_swlib_entity_types [-show_type_id]	emcli list_swlib_entity_types -show_type_id
list_swlib_folders (Listing Software Library folders)	emcli list_swlib_folders [-parent_id="parent folder id"] [-show_folder_path] [-show_folder_id]	emcli list_swlib_folders -parent_id="oracle:defaultService:em:provisioning:1:cat:B13B3B7B086458CFE040E80A19AA560C" -show_folder_id
list_swlib_storage_locations (Listing Software Library storage locations)	emcli list_swlib_storage_locations [-type="OmsShared OmsAgent Http Nfs ExtAgent"]	emcli list_swlib_storage_locations -type="OmsAgent"
refer_swlib_entity_files (Referring files from a Software Library entity)	emcli refer_swlib_entity_files -entity_rev_id="entity_rev_id" [-file="<relative file path>[<new file name>"] [-remove-file="<existing file name>"] -refer_storage="<storage location name>;<storage type>" [-use_latest_revision]	emcli refer_swlib_entity_files -entity_rev_id="oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Gener-ic:B1B1880C6A8C62AAE040548C42832D14:0.1" -file="scripts/perl/script1.pl;new_script.pl" -removefile="ALL" -refer_storage="myScripts;Http" -use_latest_revision
reimport_swlib_metadata (Re-Importing Software Library metadata)	emcli reimport_swlib_metadata	emcli reimport_swlib_metadata

Table A-3 (Cont.) Software Library EM CLI Verbs and Their Usage

Verb	Usage	Example
remove_swlib_storage_location (Removing a Software Library storage location)	emcli remove_swlib_storage_location -name="src location name" -type="OmsShared OmsAgent Http Nfs ExtAgent" -migrate_to_loc="dest location name" [-migrate_to_type="OmsShared OmsAgent Http Nfs ExtAgent"]	emcli remove_swlib_storage_location -name="myOMSSharedLocation" -type="OmsShared" -migrate_to_loc="myNewAGTLocation" -migrate_to_type="OmsAgent"
update_swlib_entity (Modifying a Software Library entity)	emcli update_swlib_entity -entity_rev_id="entity_rev_id" [-desc="entity_desc"] [-attr="<attr name>:<attr value>"] [-prop="<prop name>:<prop value>"] [-secret_prop="<secret prop name>:<secret prop value>"] [-note="note text"] [-use_latest_revision]	emcli update_swlib_entity -entity_rev_id="oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Generic:B1B1880C6A8C62AAE040548C4D14:0.1" -entity_desc="myexampleInstall description" -attr="PRODUCT:example" -attr="PRODUCT_VERSION:3.0" -attr="VENDOR:example Corp" -prop="DEFAULT_HOME:/u01/example3/" -note="myexampleInstall for test servers"
upload_swlib_entity_files (Uploading files to a Software Library entity)	emcli upload_swlib_entity_files -entity_rev_id="entity_rev_id" -host="hostname" [-file="<abs file path>;<new file name>"] [-removefile="<existing file name>"] [-credential_set_name="setname"] [-credential_name="name" -credential_owner="owner"] [-upload_storage="<storage location name>;<storage type>"] [-use_latest_revision]	emcli upload_swlib_entity_files -entity_rev_id="oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Generic:B1B1880C6A8C62AAE040548C42832D14:0.1" -file="/u01/example_downloads/file1.zip;newfile1.zip" -file="/u01/example_downloads/file2.zip" -removefile="ALL" -host="fs1.us.example.com" -credential_name="MyexampleCreds" -credential_owner="example_USER" -use_latest_revision

Provisioning Using EM CLI

Deployment Procedures can be run from the command line using EM CLI or from Cloud Control UI. Launching a procedure either from command line or from GUI requires a set of inputs to be provided. However, the mode of entering these inputs differ in both the cases. While running a Deployment Procedure from the UI, you can use a wizard to enter all the inputs required to run the procedure. However, in EM CLI, you use Properties File for entering the inputs. Properties File is a file which contains all the inputs required to run a Deployment Procedure. The following sections describe how to create properties file from scratch and use it in procedures, how to use properties file of a procedure that has already been executed, and how to create a template using a properties file and a few other attributes to run the deployment procedures.



Note:


- You cannot provision pluggable databases by running Deployment Procedures using EM CLI. For information on how to provision pluggable databases using EM CLI, view [Provisioning Pluggable Databases](#).
- The `-swlib` argument works for cloning only Oracle Database 9i Release 2. Do NOT use this argument for later releases.

This section covers the following scenarios:

- [Creating the Properties File to Submit a Deployment Procedure](#)
- [Using Properties File from an Existing Execution of a Deployment Procedure](#)
- [Launching a Procedure using an Existing Saved Procedure](#)
- [Provisioning Pluggable Databases](#)

Creating the Properties File to Submit a Deployment Procedure

This graphic illustrates how to create a template properties file, update values into the file, and then submit the procedure with the updated properties file as the input.



Step1: Create Template Properties File From a Procedure Definition

All the details required for the selected Deployment Procedure like variable names, targets, credentials, and so on are provided in this step to successfully submit the procedure from the command line. Generating the Properties file is a two-step process as follows:

1. To retrieve the **GUID** or the **Name** of the procedure, run the following command:

```
emcli get_procedures
[-type={procedure type}]
```

Example:

```
./emcli get_procedures -type=DBPROV
```

Output:

```
B3FCE84B1ED96791E040578CD7810EC5, DBPROV, Prov_112_db_using_SH_locked_acc_without_env_shift_ssubbura11, Prov_112_db_using_SH_locked_acc_without_env_shift_ssubbura11, 1.0, SSUBBURA1, SIHA_SIDB_PROC
B35E10B1F430B4EEE040578CD78179DC, DBPROV, DBREPLAYCLIENTDP_NG, Provision Oracle Database Client, 6.1, ORACLE
B35E10B1F427B4EEE040578CD78179DC, DBPROV, SIHA_SIDB_PROC, Provision Oracle Database, 1.0, ORACLE
```

2. Use the GUID or the name in the following command to generate a template properties file. Use the following command when you are running the Deployment Procedure for the first time, or when you do not have too many variables in your procedure to update:

```
emcli describe_procedure_input
[-procedure={procedure GUID}]
[-name={procedure name or procedure configuration}]
[-owner={owner of the procedure or procedure configuration}][--parent_proc={procedure of the procedure configuration. this only applies to describe a procedure configuration with the same name}]
```

The following examples describe how to use the procedure GUID to generate the properties file template:

```
./emcli describe_procedure_input -procedure=B35E10B1F427B4EEE040578CD78179DC >
procConfiguration.properties
```

This EM CLI verb describes the input data of a deployment procedure or a procedure configuration in a name-value pair format, which is also called as the *properties file format*. The advantage of this name-value file format for a procedure is that it is flexible enough to accept multiple destination targets.

Step 2: Entering New Values in The Properties File

Use any editor to open the properties file and enter values against the names. After updating all the fields, save and close the properties file.

The main goal of this step is to create a library of property files where the most common input values have been set as defaults, this in turn reduces the chances of operator errors, and also reduces the number of inputs expected from the operators.

For example, vi procConfiguration.properties



Note:

For example properties file, see section [Provisioning Oracle Database Software](#) or [Provisioning Oracle WebLogic Server](#).

Step 3: Submitting the Procedure With The Updated Properties File as Input

After the properties file is ready with the correct name-value pair required to run the Deployment procedure, you must use the EM CLI verb *submit_procedure*, which accepts the edited properties file as the input.


```
emcli submit_procedure
[-name={name of the procedure}]
[-owner={owner of the procedure}]
[-procedure={guid of the procedure}]
-input_file={data:{file_path}/file name" [-instance_name={name for the procedure in-
stance}] [-notification={procedure status}]
[-grants={users and their corresponding accessing levels}] [-sched-
ule=start_time:yyyy/MM/dd HH:mm; tz:{java timezone ID}]
```

Starting with Cloud Control 12c, you can submit the procedure either using the procedure GUID or using the procedure name/owner pair, as described in the following example:

- Submitting the properties file using the GUID of the procedure:

```
emcli submit_procedure -input_file=data:procConfiguration.properties -proce-
dure=B35E10B1F427B4EEE040578CD78179DC -schedule="start_time:2006/6/21 21:23;
tz:America/New_York" -grants="user1:VIEW_JOB; user2:FULL_JOB" -notifica-
tion="scheduled, action required, running"
```

- Submitting the properties file using the procedure name/owner pair:

```
emcli submit_procedure -input_file=data:procConfiguration.properties -name=SI-
HA_SIDB_PROC -owner=sysman -schedule="start_time:2006/6/21 21:23; tz:America/
New_York" -grants="user1:VIEW_JOB; user2:FULL_JOB" -notification="scheduled, ac-
tion required, running"
```

Output:

```
Verifying parameters ...
B35E10B1F427B4EEE040578CD78179DC
Deployment procedure submitted successfully
Note: The instanceId is B35E10B1F427B4EEE040578CD78179F1
```

This verb functions in a non-waiting mode, which means it submits the procedure for execution and returns without waiting for it to complete. The output of this verb indicates if the submission of the procedure was successful or if any errors were encountered. A successful submission displays the Instance GUID as the output.

Step 4: Verifying The Status Of the Procedure

The final step lets you to track the progress and status of the procedure. This is especially important since the submit procedure verb does not wait for the completion of the Deployment Procedure:

```
emcli get_instance_status
[-instance={instance guid}]
[-exec=execution guid]
[-xml]
[-details]
[-showJobOutput]
[-tailLength={last N characters}]
```


Example:

```
emcli get_instance_status -instance=B35E10B1F427B4EEE040578CD78179F1 -details -show-
JobOutput
```

Output:
B35E10B1F427B4EEE040578CD78179F1, WEBLOGIC_WSM, DANS_SCALEUP_WSM12, FAILED

Using Properties File from an Existing Execution of a Deployment Procedure

This graphic illustrates how to retrieve the properties file of a deployment procedure that has already been executed, update values into the file, and then submit the procedure with the updated properties file as the input.



Retrieving Properties File From an Existing Execution

All the inputs required for the selected Deployment Procedure like variable names, targets, credentials, and so on are provided in this step to successfully submit the procedure.

cedure from the command line. Generating the Properties file is a two-step process as follows:

1. To retrieve the **GUID** or the **Name** of the procedure, run the following command:

```
emcli get_procedures
[-type={procedure type}]
[-parent_proc={procedure associate with procedure configuration}]
```

Example:

```
./emcli get_procedures -parent_proc=SIHA_SIDB_PROC
```

Output:

```
B3FCE84B1ED96791E040578CD7810EC5, DBPROV, Prov_112_db_us-
ing_SH_locked_acc_without_env_shift_ssubbura11, Prov_112_db_us-
ing_SH_locked_acc_without_env_shift_ssubbura11, 1.0, SSUBBURA1, SI-
HA_SIDB_PROC
```

2. Use the GUID to retrieve the Instance ID of the procedure:

```
emcli get_instances
[-type={procedure type}]
```

Example:

```
./emcli get_instances -type=DBPROV
```

Output:

```
B3FE0C8302EA4A4CE040578CD781133C, B3FE0C8302F64A4CE040578CD781133C,
DBPROV, Prov_112_db_using_SH_locked_acc_without_env_shift_ssubbur,
Failed
B3FE34D472C00AD9E040578CD781107B, B3FE34D472CC0AD9E040578CD781107B,
DBPROV, Prov_112_db_using_SH_locked_acc_without_env_shift_ssubbura1,
Failed
```

3. Use the Instance ID in the following command to retrieve the input properties file of the instance:

```
emcli get_instance_data
[-instance={instance guid}]
[-exec=execution guid]
```

The following examples describe how to use the procedure GUID to generate the properties file template:

```
emcli get_instance_data -instance=B3FE0C8302EA4A4CE040578CD781133C > in-
stanceData.properties
```

4. Updating the Existing Values in the Properties File

The main goal of this step is to update the values in the properties file (if required). To do so, use any editor to open the properties file and enter the updated values

against the names. After updating the required fields, save and close the properties file.

Example:
vi instanceData.properties

5. Submitting the Procedure with the Updated Properties File as Input

To run the procedures from the command line you must use the EM CLI verb *submit_procedure* as described in [Creating the Properties File to Submit a Deployment Procedure](#)

6. Verifying the Status of the Procedure

To verify the status of the procedure, see [Creating the Properties File to Submit a Deployment Procedure](#).

Launching a Procedure using an Existing Saved Procedure

Procedures that are used repeatedly can be saved along with the properties file, job grants, schedules, and notifications, and so on with a unique name. This specially packaged procedure can be run using the unique name whenever required. This is especially useful when the procedure must be executed multiple number of times, and helps in saving a lot of time and effort. Running the verb *emcli get_procedures* fetches all the procedures which also include the Procedure Configurations.

To launch a procedure using an Existing Procedure Configuration File, do the following:

1. Run the verb *emcli get_procedures* to fetch an existing Procedure Configuration file.
2. Update the properties file if required.
3. Save the Procedure Configuration with the updated Properties file, and the other attributes like job grants, schedules, and notifications. To do so, see [Saving a Procedure Configuration of a Procedure](#).
4. Submit the Procedure Configuration file.

 Note:

You can update the Procedure Configuration using the *update_procedure_input* verb as described in [Updating the Procedure Configuration of a Procedure](#). After Updating the Procedure Configuration, follow Step 4 to resubmit the procedure.

5. Verify the status of the procedure.

Saving a Procedure Configuration of a Procedure

If you have to use a properties file repeatedly to run a procedure, then Oracle recommends that you save this Procedure with the properties file and give the saved procedure a name. Every time you want to run the procedure with the same properties file, you can run the saved Procedure by giving its name. To save the procedure, run the following command:

```
emcli save_procedure_input -name={name of procedure configuration} -procedure={Procedure name}
[-owner={owner of procedure}]
-input_file=data:/file path/file name
[-grants={users and their corresponding accessing levels}] [-notification={procedure status}]
[-schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone id};grace_period:xxx]
```

Example:

```
emcli save_procedure_input -name=procConfiguration -procedure=ComputeStepTest -input_file=data:/tmp/instanceData.properties -grants="user1:VIEW_JOB;user2:FULL_JOB" -notification="scheduled, action required, running" -schedule="start_time:2012/12/25 00:00;tz:American/New York;grace_period:60"
```

Updating the Procedure Configuration of a Procedure

To update the existing values in a saved procedure, run the following command:

```
emcli update_procedure_input -name={name of procedure configuration} -input_file="data:/file path/file name"
[-notification={procedure status}]
[-grants={users and their corresponding accessing levels}]
[-schedule=start_time:yyyy/MM/dd HH:mm;tz:{java timezone id};grace_period:xxx]
```

Example:

```
emcli update_procedure_input -name=procConfiguration -input_file=data:/tmp/instanceData.properties -grants="user1:VIEW_JOB;user2:FULL_JOB" -notification="scheduled, action required, running" -schedule="start_time:2012/12/25 00:00;tz:American/New York;grace_period:60"
```

Provisioning Pluggable Databases

This section provides information on how to provision pluggable databases (PDBs) using EM CLI. It contains the following:

- [Creating a New Pluggable Database](#)
- [Provisioning a Pluggable Database Using a Snapshot Profile](#)
- [Migrating a Non-Container Database as a Pluggable Database](#)
- [Unplugging and Dropping a Pluggable Database](#)

Creating a New Pluggable Database

To create a new PDB using EM CLI, follow these steps:

1. Meet the prerequisites.

For information on the prerequisites for creating a new PDB, view [Prerequisites for Creating a New Pluggable Database](#). For information on the prerequisites for creating a PDB by plugging in an unplugged PDB, view [Prerequisites for Plugging In an Unplugged Pluggable Database](#). For information on the prerequisites for creating a PDB by cloning an existing PDB, view [Prerequisites for Cloning a Pluggable Database](#).

2. Log in to EM CLI by running the following command on the OMS host:

```
$<OMS_HOME>/bin/emcli login -username=<name_of_user> -password=<password>
```

3. View the help file of the `create_pluggable_database` verb by running the following command:

```
$<OMS_HOME>/bin/emcli help create_pluggable_database
```

4. Run the `create_pluggable_database` verb:

```
$<OMS_HOME>/bin/emcli create_pluggable_database
-cdbTargetName=<Specify the CDB target name for creating new PDB>
-cdbTargetType=<Specify the CDB target type - oracle_database, rac_database>
-cdbHostCreds=<Specify the host credentials on which the CDB target is located>
[-cdbTargetCreds=<Specify the credentials of container database on which the new
PDB will be created.>]
-pdbName=<Specify a name for the new PDB>
[-numOfPDBs=<Specify the number of PDBs to be created>]
-sourceType=<Type of pdb to be created - DEFAULT, UNPLUGGED_PDB, CLONE, PROFILE>
[-sourceFromSWLIB=<If -sourceType is 'UNPLUGGED_PDB', specify if the dump loca-
tion is SWLIB or not.>]
[-pdbTemplateInSWLIB=<If -sourceFromSWLIB, specify the URN of pdb template compo-
nent in SWLIB.>]
[-sourcePDBTempStagingLocation=<If from SWLIB, specify fully qualified location
for staging the source pdb dump temporarily>]
[-unpluggedPDBType=<If -sourceType is 'UNPLUGGED_PDB', specify pdb dump type -
ARCHIVE, RMAN, XML.>]
[-sourcePDBArchiveLocation=<If -unpluggedPDBType=ARCHIVE, this is fully quali-
fied archive location>]
[-sourcePDBMetadataFile=<If -unpluggedPDBType=RMAN or XML, this is fully quali-
fied path of the source PDB metadata file>]
[-sourcePDBDataBackup=<If -unpluggedPDBType=RMAN, this is fully qualified path
of the source PDB datafile>]
[-sourcePDBName=<If -sourceType is 'CLONE', specify the name of an existing PDB
which is a valid em target>]
[-sourceCDBCreds=<If -sourceType is 'CLONE', specify the credentials of contain-
er database on which the -sourcePDBName is present>]
[-pdbAdminCreds=<Name of pdb credentials with admin role>]
[-useOMF=<Specifies that the datafiles can be stored in OMF location>]
[-sameAsSource=<Specifies that the datafiles of new PDB can be stored in the
same location as that of source CDB>]
[-newPDBFileLocation=<Specify the storage location for datafiles of the created
PDB.>]
[-createAsClone=<If -sourceType is 'UNPLUGGED_PDB' and if 'createAsClone' is
specified, the PDB will be created as clone.>]
[-lockAllUsers=<If -sourceType is 'UNPLUGGED_PDB' and if 'lockAllUsers' is speci-
fied, all PDB users of the new PDB will be locked.>]
[-noUserTablespace=<Specifies that the new DEFAULT PDB will not be created with
USERS tablespace.>]
[-useSnapClone=<If -sourceType is 'CLONE', specifies that Snap Clone must be
used for cloning the PDB.>]
[-sourceCDBHostCreds=<If -sourceType is 'CLONE' and -useSnapClone is specified,
this specifies the host credentials for the source CDB.>]
[-mountPointPrefix=<If -sourceType is 'CLONE' and -useSnapClone is specified,
this specifies the mount point prefix for the cloned volumes.>]
[-writableSpace=<If -sourceType is 'CLONE' and -useSnapClone is specified, this
specifies the writable space, in GB, for the cloned volumes.>]
[-privHostCreds=<If -sourceType is 'CLONE' and -useSnapClone is specified, this
specifies the privileged host credentials required to mount the cloned volumes
at the specified locations.>]
```

For example, you can run the following commands to create new PDBs:

```
$<OMS_HOME>/bin/emcli create_pluggable_database -cdbTargetName=database -cdbTargetType=oracle_database -pdbName=pdb -sourceType=UNPLUGGED_PDB -unpluggedPDBType=ARCHIVE -sourcePDBArchiveLocation=/u01/app/oracle/product/12.1.0/dbhome_2/assistants/dbca/templates/a.tar.gz

emcli create_pluggable_database -cdbTargetName=database -cdbTargetType=oracle_database -cdbHostCreds=HOST_CREDS -cdbTargetCreds=CDB_SYS_CREDS -pdbName=pdb -sourceType=CLONE -sourcePDBName=source_pdb -sourceCDBCreds=CDB_SYS_CREDS -useSnapClone -srcCDBHostCreds=HOST_CREDS -mountPointPrefix=/oracle -writableSpace=1 -sourcePDBTempStagingLocation=/tmp -privHostCreds=ROOT_CREDS
```

Provisioning a Pluggable Database Using a Snapshot Profile

When a PDB is cloned using the Snap Clone method, you can choose to create a snapshot profile out of the created snapshot, save the snapshot profile in Software Library, then use the snapshot profile later to provision PDBs. Using this method, you can provision a stable, standardized, and up-to-date PDB on a large number of hosts, without having to create a snapshot of the source PDB during each provisioning operation.

This section consists of the following:

- [Prerequisites for Provisioning a Pluggable Database Using a Snapshot Profile](#)
- [Procedure for Provisioning a Pluggable Database Using a Snapshot Profile](#)

Prerequisites for Provisioning a Pluggable Database Using a Snapshot Profile

Ensure that you meet the following prerequisites before provisioning a PDB using a snapshot profile:

- The 12.1.0.6 Enterprise Manager for Oracle Database plug-in, or a higher version, must be downloaded and deployed in your system.
- The prerequisites for cloning a PDB using the Snap Clone method, (described in [Prerequisites for Cloning a Pluggable Database](#)) must be met.

Procedure for Provisioning a Pluggable Database Using a Snapshot Profile

To provision a PDB using a snapshot profile, follow these steps:

1. Log in to EM CLI by running the following command on the OMS host:

```
$<OMS_HOME>/bin/emcli login -username=<name_of_user> -password=<password>
```
2. View the help file of the create_pluggable_database verb by running the following command:

```
$<OMS_HOME>/bin/emcli help create_pluggable_database
```
3. Run the create_pluggable_database verb, specifying the -sourceType parameter as CLONE, and the -sourcePDBName, sourceCDBCreds, -useSnapClone, -srcCDBHostCreds, -mountPointPrefix, -writableSpace, -privHostCreds, -saveProfile, -profileName, and -profileLocation parameters, to create a snap clone and a snapshot profile of the source PDB:

```
$<OMS_HOME>/bin/emcli create_pluggable_database
-cdbTargetName=<Specify the CDB target name for creating new PDB>
-cdbTargetType=<Specify the CDB target type - oracle_database, rac_database>
```

```
-cdbHostCreds=<Specify the host credentials on which the CDB target is located>
-pdbName=<Specify a name for the new PDB>
-sourceType=CLONE
-sourcePDBName=<Specify the name of the existing PDB that you want to clone>
-sourceCDBCreds=<Specify the credentials of the CDB within which the source PDB
is present>
-useSnapClone
-sourceCDBHostCreds=<Specify the host credentials for the source CDB.>
-mountPointPrefix=<Specify the mount point prefix for the cloned volumes.>
-writableSpace=<Specify the writable space, in GB, for the cloned volumes.>
-privHostCreds=<Specify the privileged host credentials required to mount the
cloned volumes at the specified locations.>
-saveProfile
-profileName=<The name of the profile that you want to create>
-profileLocation=<The location in software library where you want to create the
profile>
[-cdbTargetCreds=<Specify the credentials of container database on which the new
PDB will be created.>]
[-numOfPdfs=<Specify the number of PDBs to be created>]
[-pdbAdminCreds=<Specify the PDB credentials having the admin role>]
[-useOMF=<Specifies that the datafiles can be stored in OMF location>]
[-sameAsSource=<Specifies that the datafiles of new PDB can be stored in the
same location as that of source CDB>]
[-newPDBFileLocation=<Specify the storage location for datafiles of the created
PDB.>]
[-noUserTablespace=<Specifies that the new DEFAULT PDB will not be created with
USERS tablespace.>]
```

This command clones the source PDB using the Snap Clone feature, and creates a snapshot profile of the source PDB, which is stored at the specified location in Software Library.

4. Run the `list_swlib_entities` verb, specifying the `-name` and the `-show_entity_rev_id` parameters, to obtain the Uniform Resource Name (URN) of the created snapshot profile:

```
$<OMS_HOME>/bin/emcli list_swlib_entities
                        -name="<The name of the snapshot profile>"
                        -show_entity_rev_id
```

5. Run the `create_pluggable_database` verb, specifying the `-sourceType` parameter as `PROFILE`, and the `-profileURN` parameter, to provision PDBs using the created snapshot profile:

```
$<OMS_HOME>/bin/emcli create_pluggable_database
-cdbTargetName=<Specify the CDB target name for creating new PDB>
-cdbTargetType=<Specify the CDB target type - oracle_database, rac_database>
-cdbHostCreds=<Specify the host credentials on which the CDB target is located>
-pdbName=<Specify a name for the new PDB>
-sourceType=PROFILE
-profileURN=<URN of the snapshot profile that you want to use to provision PDBs>
[-cdbTargetCreds=<Specify the credentials of container database on which the new
PDB will be created.>]
[-numOfPdfs=<Specify the number of PDBs to be created>]
[-pdbAdminCreds=<Name of pdb credentials with admin role>]
[-useOMF=<Specifies that the datafiles can be stored in OMF location>]
[-sameAsSource=<Specifies that the datafiles of new PDB can be stored in the
same location as that of source CDB>]
[-newPDBFileLocation=<Specify the storage location for datafiles of the created
PDB.>]
```


`[-noUserTablespace=<Specifies that the new DEFAULT PDB will not be created with USERS tablespace.>]`

Migrating a Non-Container Database as a Pluggable Database

To migrate a non-container database (CDB) as a PDB, follow these steps:

1. Meet the prerequisites.

For information on the prerequisites for migrating a non-container database as a PDB, view [Prerequisites for Migrating a Non-CDB as a Pluggable Database](#).

2. Log in to EM CLI by running the following command on the OMS host:

```
$<OMS_HOME>/bin/emcli login -username=<name_of_user> -password=<password>
```

3. View the help file of the `migrate_noncdb_to_pdb` verb by running the following command:

```
$<OMS_HOME>/bin/emcli help migrate_noncdb_to_pdb
```

4. Run the `migrate_noncdb_to_pdb` verb:

```
$<OMS_HOME>/bin/emcli migrate_noncdb_to_pdb
-cdbTargetName=<EM CDB target into which the database will be added as PDB>
-cdbTargetType=<EM CDB target type (oracle_database|rac_database)>
-cdbDBCreds=<Named DB credentials of CDB user having sysdba privileges>
-cdbHostCreds=<Named host credentials for Oracle Home owner of CDB>
-migrationMethod=<Migration method to be used (DATAPUMP|PLUG_AS_PDB)>
-noncdbTargetName=<EM non-CDB target to be migrated>
-noncdbTargetType=<EM non-CDB target type (oracle_database|rac_database)>
-noncdbDBCreds=<Named DB credentials for non-CDB user having sysdba privileges>
-noncdbHostCreds=<Named host credentials for Oracle Home owner of non-CDB>
-pdbName=<Name of the PDB to be created on the CDB>
-pdbAdminName=<Username of the PDB administrator to be created>
-pdbAdminPassword=<Password for the PDB administrator>
[-exportDir=<Temporary file system location on the non-CDB host where the export-
ed files will be stored>]
[-importDir=<Temporary file system location on the CDB host used to stage the mi-
gration metadata and/or datafiles>]
[-useOMF=<Use OMF for datafile location if CDB is OMF enabled (Y|N)>]
[-dataFilesLoc=<Location on the CDB host where datafiles for the newly created
DB will be stored. Disk Group name in case of ASM>]
[-encryptionPwd=<Password to decrypt/encrypt datapump dump file. Mandatory if
non-CDB contains encrypted tablespaces>]
[-cdbWalletPwd=<Wallet password of the CDB. Mandatory if non-CDB contains en-
crypted tablespaces>]
[-objectExistsAction=<Action to be taken when the exported object with same name
is found on the newly created PDB (SKIP|REPLACE). Defaulted to SKIP>]
[-precheck=<Perform pre-requisite checks (YES|NO|ONLY). Defaulted to YES>]
[-ignoreWarnings=<Ignore the warnings from precheck (Y|N)>]
```

For example, you can run the following command to migrate a non-CDB as a PDB:

```
$<OMS_HOME>/bin/emcli migrate_noncdb_to_pdb -migrationMethod=datapump -noncdbTar-
getName=NON_CDB_NAME -noncdbTargetType=oracle_database -noncdbHost-
Creds=NON_CDB_HOST_CREDS -noncdbDBCreds=NON_CDB_DB_CREDS -cdbTargetName=CDB_NAME
-cdbTargetType=oracle_database -cdbHostCreds=CDB_HOST_CREDS -
cdbDBCreds=CDB_DB_CREDS -pdbName=NEW_PDB -pdbAdminName=pdbAdmin -pdbAdminPass-
word=welcome -precheck=ONLY -ignoreWarnings
```

Unplugging and Dropping a Pluggable Database

To unplug and drop a PDB, follow these steps:

1. Meet the prerequisites.

For information on the prerequisites for unplugging and dropping a PDB, see [Pre-requisites for Unplugging and Dropping a Pluggable Database](#).

2. Log in to EM CLI by running the following command on the OMS host:

```
$<OMS_HOME>/bin/emcli login -username=<name_of_user> -password=<password>
```

3. View the help file of the `unplug_pluggable_database` verb by running the following command:

```
$<OMS_HOME>/bin/emcli help unplug_pluggable_database
```

4. Run the `unplug_pluggable_database` verb:

```
$<OMS_HOME>/bin/emcli unplug_pluggable_database
-cdbTargetName=<Specify the CDB target name from which PDB needs to be unplugged>
-cdbTargetType=<Specify the CDB target type - oracle_database, rac_database>
-cdbHostCreds=<Specify the host credentials on which the CDB target is located>
-cdbTargetCreds=<Specify the credentials of container database on which the PDB
will be unplugged.>
-pdbName=<Specify name of the PDB that needs to be unplugged>
[-unplugPDBToSWLIB=<Specifies that unplugged PDB should be uploaded to Software
Library (SWLIB)>]
[-pdbTemplateNameInSWLIB=<If -unplugPDBToSWLIB, specify the name to be used for
PDB Template component in SWLIB.>]
[-tempStagingLocation=<If -unplugPDBToSWLIB, specify a temporary working directo-
ry for copying staging SWLIB files.>]
-unplugPDBTemplateType=<Specify the PDB template type - ARCHIVE, RMAN, XML.>
[-pdbArchiveLocation=<If -unplugPDBTemplateType=ARCHIVE, this is fully qualified
archive location with file name>]
[-pdbMetadataFile=<If -unplugPDBTemplateType=RMAN or XML, this is fully quali-
fied path for the PDB metadata file>]
[-pdbDataBackup=<If -unplugPDBTemplateType=RMAN, this is fully qualified path
for the PDB datafile backup>]
```

For example, you can run the following command to unplug and drop a PDB:

```
$<OMS_HOME>/bin/emcli unplug_pluggable_database -cdbTargetName=db -cdbTarget-
Type=oracle_database -cdbHostCreds=HOST_CREDS -cdbTargetCreds=CDB_CREDS -
pdbName=db_pdb -unplugPDBTemplateType=ARCHIVE -pdbArchiveLocation=/u01/app/
unplugged/db_pdb.tar.gz
```

Patching Using EM CLI

This section contains the following:

- [Before You Begin Patching](#)
- [Patching Using EM CLI](#)

Before You Begin Patching

Keep the following points in mind before patching the targets using EM CLI:

1. Target information like: target name, target type, target version, release number, platform, product and so on ready.
2. Patch information like: Patch Name (Patch Number), Release ID, Platform ID, and Language ID ready.
3. Set at least one of the following credentials on the Oracle home of the target host machines before beginning the patching process:
 - Oracle Home Named Credentials
 - Privileged Oracle Home Named Credentials
4. Setup Privilege Delegation through Sudo or PowerBroker, and apply the templates to the host target when you do not have the access (Username/Password) for the Oracle account or root account.
5. Set one of the following modes before patching:
 - **Online Mode:** This mode is helpful when you have internet connectivity. However, to search and download the patches from *My Oracle Support*, you need to set the preferred credentials for *My Oracle Support*.
 - **Offline Mode:** This mode can be used for patching provided you have already downloaded the patches to the Software Library. You can then search for them on Software Library.

Patching Using EM CLI

Targets can be patched using Enterprise Manager Command Line Interface. You do not necessarily need Cloud Control to download and apply patches.

The following table describes EM CLI patching scenarios:

Table A-4 EM CLI Patching Scenarios

Case No	Scenario	High Level Steps
Case 1	Creating a new properties file for patching targets.	To patch targets using a fresh properties file, follow these steps: 1. Select the targets, and search for the patches that you want to add to the plan. 2. Create a properties file, and save it in a temporary location. 3. Create the plan using the properties file as the input. 4. View the patch plan before submitting to know if the given patch plan can be submitted. If the action is : - Analyze, then you can submit your patch plan for analysis. - Deploy, then you can submit your patch plan for deployment. 5. Verify the status of the submitted plan. For details about how to use the EM CLI commands to perform each of the steps see Creating a New Properties File for Patching Targets.

Table A-4 (Cont.) EM CLI Patching Scenarios

Case No	Scenario	High Level Steps
Case 2	Updating the properties file of an existing patch plan to patch the targets.	To update a properties file retrieved from an existing patch plan, follow these steps: 1. Get the user-editable data for a given patch plan, and save the output as a properties file. 2. Edit the properties file using an editor. For example, vi editor. 3. Save the edited user-data. 4. View the patch plan before submitting to know if the given patch plan can be submitted. If the action is : - Analyze, then you can submit your patch plan for analysis. - Deploy, then you can submit your patch plan for deployment. 5. Verify the status of the submitted plan. For details about how to use the EMCLI commands to perform each of the steps see Using the Properties File of an Existing Patch Plan to Patch the targets.

Creating a New Properties File for Patching Targets

If you are creating the patch plan from scratch, then you need to create the properties file afresh, and submit this properties file as input for creating the plan. To do so, follow these steps:

1. Select the targets that need to be patched. To do so, run the following EM CLI command:

```
emcli get_targets
[-targets="[name1:]type1;[name2:]type2;..."]
```

For example:

```
emcli get_targets -targets=oracle_emd
```

Output:

Displays all the Management Agent targets.

Status ID	Status	Target Type	Target Name
2	Metric collection on Error	oracle_emd	h1.us.example.com:5125
2	Metric collection on Error	oracle_emd	h2.us.example.com:5125
1	Up	oracle_emd	slc01nha.us.example.com:11852
1	Up	oracle_emd	slc00bng.us.example.com:1833
1	Up	oracle_emd	adc2101349.us.example.com:1832

2. Search for the patches that you want to apply. To find the relevant patches for your plan, you either need to use the Patch ID (Basic Search), or use a combination of Release ID, Platform ID, and Product ID (Advanced Search) and drill down to the patches required. To do so, run the following EM CLI command:

```

emcli search_patches
    [-swlib]
    [-patch_name="patch_name"]
    [-product="product id" [-include_all_products_in_family]]
    [-release="release id"]
    [-platform="platform id" | -language="language id"]
    [-type="patch | patchset"]
    [-noheader]
    [-script | -xml | -format=
                                [name:<pretty|script|csv>;
                                [column_separator:"column_sep_string"];
                                [row_separator:"row_sep_string"];
    ]

```

Note:

You can search for patches in one of the following locations:

- ARU Site
- Software Library

If you have internet connectivity, then you are in online mode, and by default can look for patches on the ARU site. However, if you are in offline mode, then you must ensure that the patches are already uploaded to Software Library so you can use them.

You can perform searches in one of the following modes using EM CLI:

- Simple Search: This mode allows you to search the ARU site or Software Library using the patch ID information.
- Advanced Search: This mode allows you to provide a combination of key values like platform ID, Language ID, Release ID, and/or product ID to drill down to the patch that you are looking for.

You can use the following syntax, and the corresponding examples to perform simple and advanced search using EM CLI commands:

- (Basic Search)** To search for the patches using the **Patch ID**, do the following:

```

emcli search_patches
    [-swlib]
    [-patch_name="patch_name"]
    [-product="product id" [-include_all_products_in_family]]
    [-release="release id"]
    [-platform="platform id" | -language="language id"]
    [-type="patch | patchset"]
    [-noheader]
    [-script | -xml | -format=
                                [name:<pretty|script|csv>;
                                [column_separator:"column_sep_string"];
                                [row_separator:"row_sep_string"];
    ]

```

Example 1: Basic Search (Online Mode)

To search for patches on My Oracle Support using the Patch ID:

```
emcli search_patches -patch_name=11993573
```

Output:

```
11993573      Agent Plugin PATCH      Cloud Control (Agent)
12.1.0.1.0    Linux x86-64      American English      General Enterprise
Manager Base Platform - Plugin
```

Example 2: Basic Search (Offline Mode)

To search for patches on Software Library using the patch ID:

```
emcli search_patches -patch_name=11993573 -swlib -script
```

Output:

```
11993573      Agent Plugin PATCH      Cloud Control (Agent)
12.1.0.1.0    Linux x86-64      American English      General Enterprise
Manager Base Platform - Plugin
```

- b. (Advanced Search)** Use the **Product ID, Release ID, and Platform ID (or Language ID)** to get the patch details that you want to add to the patch plan.

Example:

To search for patches using a combination of Product ID, Release ID, and Platform ID (obtained from the earlier steps):

```
emcli search_patches -product=12383 -release=9800371121010 -platform=226
```

Output:

```
13491785      ENTERPRISE MANAGER BASE PLATFORM - AGENT 12.1.0.1.0 BP1
(PORT) Cloud Control (Agent) 12.1.0.1.0      Linux x86-64      American Eng-
lish      Recommended      Enterprise Manager Base Platform13481721
WRONG ERROR MESSAGE RETURNED FROM NMO      Cloud Control (Agent)
12.1.0.1.0      Linux x86-64      American English      General Enterprise
Manager Base Platform
```

- 3.** Create a patch-target map (stored in the properties file) using any editor, and supply information like Patch ID, Release ID, Platform ID, and Language ID. Here is a sample properties file:

```
vi demo.props
```

```
patch.0.patch_id=13426630
patch.0.release_id=9800371121010
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=abc1.example.com:1836
patch.0.target_type=oracle_emd
patch.1.patch_id=13426630
patch.1.release_id=9800371121010
patch.1.platform_id=2000
patch.1.language_id=0
patch.1.target_name=abc2.us.example.com:1839
patch.1.target_type=oracle_emd
```

- 4.** Run the `create_patch_plan` command to create the plan, and supply the newly created properties file (`demo.props`) as input:

```
emcli create_patch_plan
  -name="name"
  -input_file=data:"file_path"
  [-impact_other_targets="add_all | add_original_only | cancel"]
```

Example:

```
emcli create_patch_plan -name=demo_agent -input_file=data:demo.props -impact_oth-
er_targets=add_all
```

 Note:

If the selected target impacts other targets, then you need to add `impact_other_targets` with the value `"add_all"`. For example, if one of the agents running on the NFS home is selected for patching, other agent based on the same NFS home will also be impacted while patching, so they are all required to present in the patch plan.

5. After you have created the patch plan with all the relevant data, you can submit your patch plan in the Analyze mode to verify if the plan is deployable or not. To do so, run the following command:

```
emcli submit_patch_plan -name=demo_agent -action=analyze
```

Output:

The action "analyze" is successfully submitted on the Patch Plan "demo_agent", now "analyze" is in progress.

The **Analyze** mode facilitates the plan to perform all the validations to ensure that the plan is deployable. Only once the analysis is successful you should deploy the plan.

6. To verify the status of the patch plan, run the following EM CLI command:

```
emcli show_patch_plan -name=demo_agent -info | grep plan_status
```

Output:

```
<plan_status>CONFLICTS</plan_status>
```

If you see any conflicts, then you must resolve them before deploying the plan. You can use the User Interface to resolve the issues, and then rerun the plan until the status is **CLEAN**.

7. After a successful analysis, you can deploy the patch plan. To do so, run the following command with action **deploy**:

```
emcli submit_patch_plan -name=agent_demo -action=deploy
```

Output:

The action "deploy" is successfully submitted on the Patch Plan "demo_agent", now "deploy" is in progress

8. To verify the status of the plan, run the EM CLI command `show_patch_plan` as mentioned in step 6. Only when the output of the command is `DEPLOY_SUCCESS`, it means that the plan has been successfully deployed, and the targets mentioned in the patch plan have been patched.

```
emcli show_patch_plan -name=demo_agent -info
```

Output:

```
<plan>
  <planDetails>
    <plan_id>79CAF6A6DAFCFEE6654C425632F19411</plan_id>
    <name>demo</name>
    <type>PATCH</type>
    <description/>
    <conflict_check_date>Tue Feb 21 18:04:04 PST 2012</conflict_check_date>
    <conflict_check_date_ms>1329876244000</conflict_check_date_ms>
```

```
<is_deployable>1</is_deployable>
<plan_status>CONFLICTS</plan_status>
<review_status>CONFLICT_FREE</review_status>
<created_date>Tue Feb 21 17:40:47 PST 2012</created_date>
<created_date_ms>1329874847000</created_date_ms>
<created_by>SYSMAN</created_by>
<last_updated_date>Tue Feb 21 17:58:29 PST 2012</last_updated_date>
<last_updated_date_ms>1329875909000</last_updated_date_ms>
<last_updated_by>SYSMAN</last_updated_by>
<grant_priv>yes</grant_priv>
<user_plan_privilege>FULL</user_plan_privilege>
<see_all_targets>N</see_all_targets>
<homogeneousGroupLabel>Database Instance 10.2.0.1.0 (Linux x86-64)</homogeneous-
GroupLabel>
<executeGuid/>
<executeUrl/>
<planDetails/>
```

9. To get the details of the patching procedure/job that you submitted, use the GUID of the execution in the command `get_job_execution_details` as follows:

```
emcli get_job_execution_detail
    -execution={execution_id}
    [-xml [-showOutput [-tailLength={length}]]]
```

For Example:

```
emcli get_job_execution_detail -execution=79CAF6A6DAFCFEE6654C425632F19411 -xml
```

Using the Properties File of an Existing Patch Plan to Patch the targets

To edit an existing patch plan after it has been created for updating the patch-target pairs or generic information or deployment options, you can follow the steps listed here:

1. To view the user-editable fields of an existing plan, run the `get_patch_plan_data` command, and save the output to a properties file as follows:

```
$ emcli get_patch_plan_data -name=demo_agent >demo_agent.props
```

Output:

```
name=demo_agent
description=
deployment_date=
planPrivilegeList=VIEWER:ADMIN:VIEW;OPER:ADMIN:VIEW;DESIGNER:ADMIN:VIEW
patch.0.patch_id=13426630
patch.0.release_id=9800371121010
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=abc1.example.com:1836
patch.0.target_type=oracle_emd
patch.1.patch_id=13426630
patch.1.release_id=9800371121010
patch.1.platform_id=2000
patch.1.language_id=0
patch.1.target_name=abc2.example.com:4473
patch.1.target_type=oracle_emd
deploymentOptions.StageLocation=%emd_emstagedir%
deploymentOptions.AdvancedOPatchOptions=null
deploymentOptions.StagePatches=true
```


2. Edit the properties file (`demo_agent.props`) using any editor. You can change the storage location as follows:

```
name=demo_agent
description=
deployment_date=
planPrivilegeList=VIEWER:ADMIN:VIEW;OPER:ADMIN:VIEW;DESIGNER:ADMIN:VIEW
patch.0.patch_id=13426630
patch.0.release_id=9800371121010
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=abc1.example.com:1836
patch.0.target_type=oracle_emd
patch.1.patch_id=13426630
patch.1.release_id=9800371121010
patch.1.platform_id=2000
patch.1.language_id=0
patch.1.target_name=abc2.example.com:4473
patch.1.target_type=oracle_emd
deploymentOptions.StageLocation=%emd_emstagedir%/demo
deploymentOptions.AdvancedOPatchOptions=null
deploymentOptions.StagePatches=true
```

3. To save the patch plan with the new edited data, run `set_patch_plan_data` command as follows:

```
emcli set_patch_plan_data -name=demo_agent -input_file=data:demo_agent.props
```

Output:

It is successfully on updating deployment options from the patch plan.



Note:

If the selected target impacts other targets, then you need to add `impact_other_targets` with the value "add_all". For example, if one of the agents running on the NFS home is selected for patching, other agent based on the same NFS home will also be impacted while patching, so they are all required to present in the patch plan.

4. Follow Steps 5, 6, 7, 8, and 9 mentioned in the [Creating a New Properties File for Patching Targets](#) to complete the patching process.

WorkFlow Examples Using EM CLI Commands

The following sections describe some of the provisioning, patching, and Software Library tasks that can be performed using EM CLI commands:

- [Provisioning Oracle Database Software](#)
- [Provisioning Oracle WebLogic Server](#)
- [Provisioning User Defined Deployment Procedure](#)
- [Patching WebLogic Server Target](#)
- [Creating a New Generic Component by Associating a Zip File](#)
- [Migrate and Remove a Software Library Storage Location](#)

- [Adding ATS Service Test from Using EM CLI](#)

Provisioning Oracle Database Software

This use case describes how to provision an Oracle Database Software using the EM CLI commands available in Cloud Control. The first step is to filter out the database procedures running in your enterprise, from the list, select the Single Instance Database procedure and its corresponding GUID. For the SI DB procedure, a new properties file is created from scratch. Initially, the name-value pair in the template will be empty, you must edit the attributes in the properties file to update the values. Following which the procedure is submitted with the updated properties file as the input, and tracked to completion.



Note:

The following verb clones only Oracle Database 9i Release 2: `emcli clone_database_home -swlib true`.

Here is the step-by-step procedure with the outputs:

1. To retrieve the GUID of the Deployment Procedure, run the following command:

```
./emcli get_procedures | grep DB_
```

For Example:

```
./emcli get_procedures | grep DB
B3F.CE84B1ED96791E040578CD7810EC5, DBPROV, Prov_112_db_using_SH_locked_acc_with-
out_env_shift_username1, Prov_112_db_using_SH_locked_acc_without_env_shift_ssub-
bura11, 1.0, USERNAME1, SIHA_SIDB_PROC
B35E10B1F42AB4EEE040578CD78179DC, DB_PROV_UPGRADE, DbProvUpgradedDP, Upgrade Ora-
cle Database, 1.0, ORACLE
B35E10B1F427B4EEE040578CD78179DC, DBPROV, SIHA_SIDB_PROC, Provision Oracle Data-
base, 1.0, ORACLE
```

Select the GUID corresponding to the SIHA_SIDB_PROC , which is
B35E10B1F427B4EEE040578CD78179DC

2. Create the Properties File template using the following command:

```
./emcli describe_procedure_input -procedure=B35E10B1F427B4EEE040578CD78179DC >
sihasidb.properties
```

3. Use an editor to view the contents of the generated properties file `sihasidb.properties` file, and enter the required values.

For example, here is a sample properties file used with the values updated in them:

```
# The Procedure Configuration with name emcli_11202 has input and arguments as
follows:
# Input properties are:
DB_COMPONENT=11<ADMIN_NAME>/Oracle Database Installation Media
DB_HOST_NORMAL_CREDNAMES=AIME_USER1:<USERNAME>
DB_HOST_ROOT_CREDNAMES=AIME_ROOT:<USERNAME>
DB_ORACLE_BASE_LOC=/scratch/db11202
DB_ORACLE_HOME_LOC=/scratch/db11202/app/product/11.2.0/db
```

```
DB_PRODUCT_VERSION=11.2.0.2.0
DEPLOY_MODE=DEPLOY_DB
OINSTALL_GROUP=svrtech
OSDBA_GROUP=dba
OSOPER_GROUP=oper
PAUSE_AFTER_PREREQ=false
RAC_HOME_SHARED=false
SOURCE_TYPE=SOFTWARE_LIBRARY
TARGET_HOST_LIST=host.us.example.com
WORK_DIR_LOC=/tmp
```

4. Submit the procedure using the following command:

```
./emcli submit_procedure -input_file=data:sihasidb.properties
-instance="emcli_db1" -procedure=B35E10B1F427B4EEE040578CD78179DC
Verifying parameters ...
Schedule not specified, defaults to immediate
A8F7700333BAE9FAE040E40A45D866F1
Deployment procedure submitted successfully
```

Provisioning Oracle WebLogic Server

This use case describes how to provision an Oracle WebLogic Server, and how to Scale up and Scale out Middleware procedures using the EM CLI commands available in Cloud Control.

Cloud Control supports the following usecases for provisioning Oracle WebLogic Server using EM CLI commands:

- [Provisioning Oracle WebLogic Server Using the Provisioning Profile](#)
- [Scaling Up or Scaling Out Middleware Deployment Procedure](#)

Prerequisites for Provisioning Oracle WebLogic Server

- Ensure that you have setup a WebLogic Domain with Administrator Server and Managed Server, and registered your targets with the OMS so that your host target is discovered on the Middleware Provisioning Page.
- Create the WebLogic Domain Provisioning Profile, this ensures that the domain selected and its Middleware Home are archived and stored in the software library for future cloning operations. You can use this profile while cloning a WebLogic domain.

Provisioning Oracle WebLogic Server Using the Provisioning Profile

The first step is to filter out the FMW procedures running in your enterprise, from the list, select the FMWPROV procedure and its corresponding GUID. For the FMWPROV procedure, a new properties file template is created from scratch. Initially, the name-value pair in the template will be empty, you must edit the attributes in the properties file to update the values. Following which the procedure is submitted with the updated properties file as the input, and tracked to completion.

Follow these steps:

1. To retrieve the GUID of the Deployment Procedure, run the following command:

```
./emcli get_procedures | grep FMWPROV_
```

The output appears in the following format:

<proc_guid>, <procedure_type>, <Procedure_name>, <Display name>, <version>, <Parent procedure name>

For example:

```
./emcli get_procedures | grep FMWPROV_
B35E10B1F154B4EEE040578CD78179DC, FMW Provisioning, FMWPROV_DP, Provision Middle-
ware, 2.0, ORACLE
```

2. Use the GUID retrieved in the previous step to prepare the Properties File template using the following command:

```
./emcli describe_procedure_input - procedure=<proc_guid> -name = <proc_name>
```

For example:

```
./emcli describe_procedure_input -procedure=B35E10B1F154B4EEE040578CD78179DC >
instanceFMWData.properties
A properties file with the name instanceFMWData.properties is created
```

3. Use an editor to open the generated properties file *instanceFMWData.properties*, and enter the required values.

For example, here is a sample properties file used with the values updated. For information about these parameters, see [Table A-5](#).

```
FMW_PROFILE_LOCATION=Fusion Middleware Provisioning/Profiles/WLS 11g IM Profile
DEST_MIDDLEWARE_BASE=/scratch/oracle_wls/mwBase
DEST_ADMIN_SERVER_PASSWORD=<password>
DEST_ADMIN_SERVER_USERNAME=weblogic
DEST_JDK_HOME=/usr/local/packages/jdk7
ANALYZE_MODE=false
DEST_HOST_LIST.0=example.com
DEST_HOST_CREDENTIAL_LIST.0=user1:<password>
PROVISIONING_MODE=BASIC
SUBMITTED_FROM_UI=true
```

[Table A-5](#) describes the parameters used in the aforementioned example.

Table A-5 Description of the Parameters Used in a Properties File That Is Used for Provisioning Oracle WebLogic Server with a Provisioning Profile

Parameter	Description
FMW_PROFILE_LOCATION	Absolute path to the fusion middleware profile. For example, Fusion Middleware Provisioning/Profiles/WLS 11g IM Profile
DEST_MIDDLEWARE_BASE	Absolute path to the middleware base directory on the destination host. For example, /scratch/oracle_wls/mwBase
DEST_ADMIN_SERVER_PASSWORD	Password of the administration server on the destination host.
DEST_ADMIN_SERVER_USERNAME	User name of the administration server on the destination host. For example, weblogic
DEST_JDK_HOME	Absolute path to the JDK home on the destination host. For example, /usr/local/packages/jdk7

Table A-5 (Cont.) Description of the Parameters Used in a Properties File That Is Used for Provisioning Oracle WebLogic Server with a Provisioning Profile

Parameter	Description
ANALYZE_MODE	Prerequisite Mode. If set to true, then the deployment procedure only runs the prerequisite checks and pauses for you to examine the checks.
DEST_HOST_LIST.0	List of destination hosts on which you want to provision the Oracle WebLogic Server.
DEST_HOST_CREDENTIAL_LIST.0	List of Host Credentials.
PROVISIONING_MODE	For BASIC mode, you provide only some basic details and the other details are determined by the deployment procedure. For ADVANCED mode, you provide all the details.
SUBMITTED_FROM_UI	If set to True, then the deployment procedure is submitted from the UI. If set to False, then the deployment procedure is submitted from the command line (emcli).

4. Submit the procedure with the generated instanceFMWData.properties properties file as the input:

```
emcli submit_procedure -input_file=data:<input_properties_file> -procedure=<proc_guid> -instance_name=<optional_DP_Instance_Name>
```

```
./emcli submit_procedure -input_file=data:instanceFMWData.properties -procedure=B35E10B1F154B4EEE040578CD78179DC
```

Scaling Up or Scaling Out Middleware Deployment Procedure

The process of increasing a cluster's capacity by adding additional server instances to the cluster on an existing machine, or adding machines to the cluster to host the new server instance, is called Scaling up. Scaling Up and Scaling Out Managed Server can be achieved through the command line using EM CLI commands available in Enterprise Manager 13c.

In this use case, the Instance GUID of the SCALEUP procedure is retrieved, which in turn is used to retrieve the input properties file of this instance of the procedure. After making necessary updates to the properties file, like adding another user-friendly so on, the procedure is submitted with the updated properties file as the input:

Here is the step-by-step process:

1. To retrieve the GUID of the Deployment Procedure, run the following command:

```
./emcli get_procedures | grep SCALEUP_
```

The output appears in the following format:

```
<proc_guid>, <procedure_type>, <Procedure_name>, <Display name>, <version>, <Parent procedure name>
```

For example:

```
./emcli get_procedure | grep SCALEUP
B35E10B1F154B4EEE040578CD78179DC, FMW Provisioning, SCALEUP DP, Scale up/Scale out Middleware, 2.0, ORACLE
```

2. Use the Instance GUID retrieved in the previous step to get input properties of an instance of the procedure:

```
./emcli get_instance_data -instance=<instance_guid> -exec=<execution_guid>
```

For example:

```
emcli get_instance_data -instance=B35E10B1F140B4EEE040578CD78179DC > instanceData.properties
A properties file with the name instanceData.properties is created.
```

 Note:

This step is valid only if the instances of the procedure is available, which means that the procedure should have been submitted at least once in the past. If you have never submitted the procedure, then you may see an error message as follows:

Instance with GUID=<guid> is not found in repository. Please make sure the value is correct and try again.

3. Use an editor to open the generated properties file instanceData.properties, and update the existing values in the properties file.

For example, here is a sample properties file used with the values updated. For information about these parameters, see [Table A-6](#).

```
COHERENCE_ENABLED=false
DOMAIN_TARGET_NAME=/Farm03_base_domain/base_domain
ONLINE_MODE=true
DOMAIN_TYPE=General
TEMPLATE=template.jar
TEMPLATE_NAME=mytemplate
APPS_ARCHIVE=apps.zip
ADMIN_HOST_NAME=example.com
ADMIN_LISTEN_ADDRESS=10.240.34.37
ADMIN_LISTEN_PORT=7002
ADMIN_PROTOCOL=t3
ADMIN_WLS_USERNAME=weblogic
ADMIN_WLS_PASSWORD=<password>
ADMIN_WORK_DIR_LOC=/tmp/scaleUpSrc
FARM_PREFIX=Farm03
OMS_WORK_DIR_LOC=/tmp
ARCHIVE_FILE_NAME=archive.jar
CLONING_JAR_NAME=cloningclient.jar
SESSION_TS_LOC=/20151013122409043
REF_PREREQ_ZIP=prereq.zip
REF_SIZE_FILE=sizeprereq
IS_CLONE=false
IS_PLAIN=true
NOT_WINDOWS=true
USE_OWNER_CREDENTIALS=true
APPS_LIST_FILE_NAME=files.list
PREREQ_ONLY_DP=false
COMPUTE_SKIP_CLONE=true
USE_EXISTING_HOME=false
IS_OHS_EXIST=false
IS_CUSTOM_LOCAL_DOMAIN=false
IS_SERVER_MIGRATION=false
```

```

WLS_VERSION=10.3.6.0
IS_EXEC_PREREQ=false
CONFIGURE_JOC=false
ADMIN_HOST.0.name=example.com
ADMIN_HOST.0.type=host
ADMIN_HOST.0.normalHostCreds=NAME:<password>:user1
ADMIN_HOST.0.SERVER_ADDRESS=example.com
ADMIN_HOST.0.IS_SERVER_MIGRATABLE=false
ADMIN_HOST.0.ORACLE_HOME=n/a
ADMIN_HOST.0.SERVER_PORT=19030
ADMIN_HOST.0.DOMAIN_HOME_ADMIN_HOST=/scratch/oracle_wls/mwBase/domains/
base_domain
ADMIN_HOST.0.MIDDLEWARE_HOME=/scratch/oracle_wls/mwBase/middleware
ADMIN_HOST.0.TLOG_DIR=/scratch/oracle_wls/mwBase/domains/base_domain/servers/
Managed_Server1/tlogs
ADMIN_HOST.0.CLUSTER_NAME_ADMIN_HOST=Cluster_1
ADMIN_HOST.0.JDK_LOC=/ade_autofs/gd29_3rdparty/nfsdo_generic/JDK7/MAIN/
LINUX.X64/150608.1.7.0.85.0B015/jdk7/jre
ADMIN_HOST.0.SSL_PORT=0
ADMIN_HOST.0.EX_PATTERN=*.mem
ADMIN_HOST.0.MS_OVERRIDE_PORT=0
ADMIN_HOST.0.CLONE_MODE=true
ADMIN_HOST.0.WORK_DIR_LOC_ADMIN_HOST=/tmp/scaleUpSrc
ADMIN_HOST.0.MACHINE_NAME_ADMIN_HOST=Machine_1
ADMIN_HOST.0.WLS_HOME_ADMIN_HOST=/scratch/oracle_wls/mwBase/middleware/wlserv-
er_10.3
ADMIN_HOST.0.SOURCE_SERVER_NAME=Managed_Server_1
ADMIN_HOST.0.SERVER_NAME_ADMIN_HOST=Managed_Server1
DEST_MANAGED_SERVERS.0.name=example.com
DEST_MANAGED_SERVERS.0.type=host
DEST_MANAGED_SERVERS.0.normalHostCreds=NAME:<password>:user1
DEST_MANAGED_SERVERS.0.START_SERVER_REQUIRED=true
DEST_MANAGED_SERVERS.0.JOC_PORT=9988
DEST_MANAGED_SERVERS.0.WLS_HOME_DEST_MANAGED_SERVERS=/scratch/oracle_wls/mwBase/
middleware/wlserver_10.3
DEST_MANAGED_SERVERS.0.MACHINE_NAME=Machine_1
DEST_MANAGED_SERVERS.0.SERVER_NAME_DEST_MANAGED_SERVERS=Managed_Server1
DEST_MANAGED_SERVERS.0.JRE_LOC=/usr/local/packages/jdk7
DEST_MANAGED_SERVERS.0.START_NM=false
DEST_MANAGED_SERVERS.0.START_MS_USE_NM=true
DEST_MANAGED_SERVERS.0.NM_LISTEN_ADDRESS=example.com
DEST_MANAGED_SERVERS.0.CLUSTER_NAME_DEST_MANAGED_SERVERS=Cluster_1
DEST_MANAGED_SERVERS.0.DOMAIN_HOME_DEST_MANAGED_SERVERS=/scratch/oracle_wls/
mwBase/domains/base_domain
DEST_MANAGED_SERVERS.0.NM_LISTEN_PORT=5558
DEST_MANAGED_SERVERS.0.FMW_HOME_DEST_MANAGED_SERVERS=/scratch/oracle_wls/mwBase/
middleware
DEST_MANAGED_SERVERS.0.WORK_DIR_LOC_DEST_MANAGED_SERVERS=/tmp/scaleUpDest
DEST_MANAGED_SERVERS.0.MS_PORT_DETAILS=19030:Listen Port
DEST_MANAGED_SERVERS.0.DEST_MANAGED_SERVER_LISTEN_ADDRESS=example.com

```

Table A-6 describes the parameters used in the aforementioned example.

Table A-6 Description of the Parameters Used in a Properties File for Scaling Up or Scaling Out a WebLogic Server

Parameter	Description
COHERENCE_ENABLED	If set to True, then Coherence is configured through deployment procedure. If set to False, then Coherence is not configured through deployment procedure.
DOMAIN_TARGET_NAME	Name of the domain target.
ONLINE_MODE	If set to True, then the scaleup operation is done when other servers and applications are up and running. If set to False, then the scaleup operation is done when other servers and applications are not up and running.
DOMAIN_TYPE	Domain type. For wls, the value is General. For SOA, the value is soa.
TEMPLATE	Server template jar file name. For example, template.jar.
TEMPLATE_NAME	Server template name. For example, mytemplate.
APPS_ARCHIVE	Applications archive file name. For example, apps.zip.
ADMIN_HOST_NAME	Name of the host where the administration server is running.
ADMIN_LISTEN_ADDRESS	Administration Server Listen Address. (HostName/IP Address where Admin Server is running).
ADMIN_LISTEN_PORT	Administration Server Listen Port.
ADMIN_PROTOCOL	Type of protocol used by the administration server.
ADMIN_WLS_USERNAME	User name to connect to the Weblogic domain.
ADMIN_WLS_PASSWORD	Password to connect to the Weblogic domain.
ADMIN_WORK_DIR_LOC	Absolute path to the working directory of the administration server.
FARM_PREFIX	Farm prefix of the domain where you want to add a server.
OMS_WORK_DIR_LOC	Absolute path to the working directory of the OMS.
ARCHIVE_FILE_NAME	Archive file name. For example, archive.jar.
CLONING_JAR_NAME	Cloning client jar name. For example, cloningclient.jar.
SESSION_TS_LOC	Session time stamp. For example, /20151013122409043.
REF_PREREQ_ZIP	Prerequisite zip file name. For example, prereq.zip.
REF_SIZE_FILE	Reference prereq size file. For example, sizeprereq.
IS_CLONE	Set to False by default.
IS_PLAIN	Set to True, if domain is a plain wls.
NOT_WINDOWS	Set to False, if domain is present on the Windows box. Set to True, if domain is not present on the Windows box.
USE_OWNER_CREDENTIALS	Set to True to set the same owner's credentials.
APPS_LIST_FILE_NAME	Set the value to files.list. files.list that contains the list of all the applications running on the server.
PREREQ_ONLY_DP	Prerequisite mode. If set to true, then the deployment procedure only runs the prerequisite checks and pauses for you to examine the checks.
COMPUTE_SKIP_CLONE	By default, set to True, and therefore skips cloning. Set to false to avoid skipping cloning.

Table A-6 (Cont.) Description of the Parameters Used in a Properties File for Scaling Up or Scaling Out a WebLogic Server

Parameter	Description
USE_EXISTING_HOME	Set to True, if the server is created on the same host where the middleware home is already present. Else, set to False and create mwh.
IS_OHS_EXIST	Set to True, if the domain is OHS frontended. Else, set it to False.
IS_CUSTOM_LOCAL_DOMAIN	Set to True if custom local domain exists. Default value is False.
IS_SERVER_MIGRATION	Set to True if the server is migratable. Set it to False if the server is not migratable.
WLS_VERSION	Version of the Oracle WebLogic Server.
IS_EXEC_PREREQ	Set to True if OUI (product specific) prerequisite needs to be executed on destination server host. Else, set it to False.
CONFIGURE_JOC	Set to True if Java object cache needs to be configured on the scaling up server. Else, set it to False.
ADMIN_HOST.0.name	Administration Server host name.
ADMIN_HOST.0.type	Administration Server host type.
ADMIN_HOST.0.normalHostCreds	Administration Server host login credential.
ADMIN_HOST.0.SERVER_ADDRESS	.Administration Server host address.
ADMIN_HOST.0.IS_SERVER_MIGRATABLE	Set to True if the Administration Server host is migratable. Set to False if the Administration Server host is not migratable.
ADMIN_HOST.0.ORACLE_HOME	Administration Server Oracle Home location.
ADMIN_HOST.0.SERVER_PORT	Administration Server listen port.
ADMIN_HOST.0.DOMAIN_HOME_ADMIN_HOST	Domain home location present in Administration Server host.
ADMIN_HOST.0.MIDDLEWARE_HOME	Middleware home location present in Administration Server host.
ADMIN_HOST.0.TLOG_DIR	TLOG directory location present in Administration Server host.
ADMIN_HOST.0.CLUSTER_NAME_ADMIN_HOST	Cluster name present in the Administration Server host.
ADMIN_HOST.0.JDK_LOC	Absolute path to the JDK home on the Administration Server host.
ADMIN_HOST.0.SSL_PORT	SSL listen port, if SSL is configured on Administration Server.
ADMIN_HOST.0.EX_PATTERN	Value is *.mem.
ADMIN_HOST.0.MS_OVERRIDE_PORT	Managed server override port, if SSL is configured on the Administration Server.
ADMIN_HOST.0.CLONE_MODE	Set to True, if the server scaling up is cloned from another server. Else, set to False.
ADMIN_HOST.0.WORK_DIR_LOC_ADMIN_HOST	Working directory location of the Administration Server host.
ADMIN_HOST.0.MACHINE_NAME_ADMIN_HOST	Machine name present in the Administration Server host.
ADMIN_HOST.0.WLS_HOME_ADMIN_HOST	Weblogic home of the Administration Server host.
ADMIN_HOST.0.SOURCE_SERVER_NAME	Set to True, and then specify the source server name.
ADMIN_HOST.0.SERVER_NAME_ADMIN_HOST	Server name of the Scaling up Server.

Table A-6 (Cont.) Description of the Parameters Used in a Properties File for Scaling Up or Scaling Out a WebLogic Server

Parameter	Description
DEST_MANAGED_SERVERS.0.name	Destination managed server host name.
DEST_MANAGED_SERVERS.0.type	Destination managed server type. Value is host.
DEST_MANAGED_SERVERS.0.normalHostCreds	Destination managed Server host login credential.
DEST_MANAGED_SERVERS.0.START_SERVER_REQUIRED	Set to True if the server needs to be started after adding it to the domain. Else, set to False.
DEST_MANAGED_SERVERS.0.JOC_PORT	Set the Java Object Cache port number to 9988, if it needs to be configured for the destination managed server.
DEST_MANAGED_SERVERS.0.WLS_HOME_DEST_MANAGED_SERVERS	WebLogic home for the destination managed server.
DEST_MANAGED_SERVERS.0.MACHINE_NAME	Name of the destination managed server machine.
DEST_MANAGED_SERVERS.0.SERVER_NAME_DEST_MANAGED_SERVERS	Name of the destination managed server.
DEST_MANAGED_SERVERS.0.JRE_LOC	Absolute path to the JDK home on the destination host.
DEST_MANAGED_SERVERS.0.START_NM	If destination managed server is associated with a existing Machine in running state, then START_NM is set to False. If destination managed server is associated with a new machine, then set to True and start the machine.
DEST_MANAGED_SERVERS.0.START_MS_USE_NM	If destination managed server is associated with a machine, then set to True. Else, set to False.
DEST_MANAGED_SERVERS.0.NM_LISTEN_ADDRESS	Listen address of node manager of the machine that is associated with the destination managed server.
DEST_MANAGED_SERVERS.0.CLUSTER_NAME_DEST_MANAGED_SERVERS	If destination managed server is a part of the cluster, you specify the cluster name.
DEST_MANAGED_SERVERS.0.DOMAIN_HOME_DEST_MANAGED_SERVERS	Domain Home Location.
DEST_MANAGED_SERVERS.0.NM_LISTEN_PORT	If destination managed server is associated with a machine, you specify the node manager listen port.
DEST_MANAGED_SERVERS.0.FMW_HOME_DEST_MANAGED_SERVERS	Middleware home of the destination managed server.
DEST_MANAGED_SERVERS.0.WORK_DIR_LOC_DEST_MANAGED_SERVERS	Working directory location of destination managed server host.
DEST_MANAGED_SERVERS.0.MS_PORT_DETAILS	Destination managed server listen port.
DEST_MANAGED_SERVERS.0.DEST_MANAGED_SERVER_LISTEN_ADDRESS	Destination managed server listen address.

4. Submit the procedure with the generated properties file as the input:

```
./emcli submit_procedure -input_file=data:<input_properties_file> -procedure=<proc_guid> -instance_name=<optional_DP_Instance_Name>
```

```
./emcli submit_procedure -input_file=data:instanceData.properties -procedure=B35E10B1F140B4EEE040578CD78179DC
```

Provisioning User Defined Deployment Procedure

This use case describes how to provision a User Defined Deployment Procedure (UDDP) using the EM CLI commands available in Cloud Control. This use case essentially covers, creating the UDDP using the Cloud Control UI, and then submitting the UDDP using the EM CLI commands.

In this use case, a User Defined Deployment Procedure to provision JRE6 on a Linux host `abc.example.com` is created using the Cloud Control UI. Steps like **Transfer JRE** and **Check JRE Version** are added to the procedure, and the procedure is submitted with a unique submission name. EM CLI command is then used to retrieve the instance GUID of the procedure submitted. Minor modifications are made to the properties file, and then submitted through EM CLI.

Prerequisites for Provisioning User Defined Deployment Procedure

Ensure that you meet the following prerequisites:

- Log in to Cloud Control as a designer.
- Create Software Library directive to install JRE6 on Linux in the following directory: `/software_library/provisioning/install_jre6_linux32`. Note, you can choose any directory that you want.
- Create Software Library component containing hotspot JRE6 for Linux in the following directory: `/software_library/provisioning/hotspot_jre6_linux32`.

Adding Steps and Phases to User Defined Deployment Procedure Using GUI

To add phases and steps to User Defined Deployment Procedure (UDDP), log in to Cloud Control as a Designer, and follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. On the Provisioning page, from the **Actions** menu select **Create New**, and click **Go**.
3. Provide a unique name for your procedure **UDDPTest**, and click **Procedure Steps** tab.
4. On the Procedure Variables tab, add a procedure variable called `destination_path`.
5. Select the Default Phase, and click **Insert** to add a new step to the phase. On the Create wizard select Type as **Library:Component**. The page refreshes, and a five-step wizard appears.
 - a. On the Create page, enter a unique name **Transfer JRE**, and then click **Next**.
 - b. On the Select Component page, select the Component `hotspot_jre6_linux32`.
 - c. On the Select Directive page, select the directive `install_jre6_linux32`.
 - d. On the Map Properties page, map the directive properties with the variables defined. For example, set the `destination_path` directive property to Choose Variable, and then choose the procedure variable that you set `destination_path`.

- e. On the review page, review the details, and click **Finish**.
6. Select the step **Transfer JRE**, and click **Insert**. On the Create Wizard, select Type **Host Command**. The page refreshes, and a three-step wizard appears.
 - a. On the Create page, enter a unique name **Check JRE Version**, and then click **Next**.
 - b. On the Enter Command page, enter the following command:


```
${data.destination_path}/jre1.6.0_23/bin/java -version
```
 - c. On the review page, review the details, and click **Finish**.
7. Go back to the Procedure Library page, and select the **UDDPTTest** procedure that you just created, and click **Launch**. To complete the wizard enter the following details: target where you want to provision your procedure, variable (destination path: /tmp), credential info, and notification information.
8. Once you have provided all the details, click Submit. Enter the a unique Submission name **FirstUDDP**.
9. After the procedure has run, verify the output of the **Check JRE Version** step. Ideally the version should be JRE6.

Using EM CLI commands to Run an Instance of the Procedure

Use EM CLI commands to submit the procedure instance:

1. Run the following command to retrieve a list of all the procedures that you have submitted, and note down the instance ID:


```
emcli get_instances
```

 For example: `emcli get_instances -type=DemoNG`
2. Run the following command to get a list of inputs submitted for your procedure:


```
emcli get_instance_data - instance=<procedure_instance_ID>
```

 For example: `emcli get_instance_data -instance=16B15CB29C3F9E6CE040578C96093F61`
`> mydp.properties`
3. Edit the file (mydp.properties), and change the values of the property destination path to /scratch.
4. Submit the procedure with the modified properties file as the input:


```
emcli submit_procedure -input_file=data:<input_file> -name=<procedure_name> -procedure=<procedure_guid>
```

 For example: `emcli submit_procedure -input_file=data:mydp.properties -name=UDDPTest -procedure=16B15CB29C3F9E6CE040578C96093F61`

Patching WebLogic Server Target

This procedure describes how to create a patch plan, update the values in them, and submit them to deploy patches on the selected targets. This workflow captures end-to-end steps on patching WLS targets. The process of patching is the same irrespective of the targets selected.

To patch WebLogic Server targets, follow these steps:

1. Run the following command to search for the release ID of the Oracle WebLogic Release 10.3.5:

```
emcli list_aru_releases -name="10.3.5"
```

Output:

Release ID	Release Name	Long Release Name
8191035020	10.3.5.0.2	WLS 10.3.5.0.2
8191035010	10.3.5.0.1	WLS 10.3.5.0.1
8191035000	10.3.5	WLS 10.3.5
95103500	10.3.5	WLS 10.3.5

2. Run the following command to search for the product ID of Oracle WebLogic:

```
emcli list_aru_products -name="Oracle WebLogic Server"
```

Output:

Product ID	Product Name
15991	Oracle WebLogic Server
16725	Oracle WebLogic Server Virtual Edition

3. Run the following command to search for the platform ID of a Generic Platform:

```
emcli list_aru_platform -name="Generic Platform"
```

Output:

Platform ID	Platform Name
2000	Generic Platform
1204	NLS Generic Platform

4. Search for the Patch ID using the product, release, and platform details that you have from the previous steps as followings:

```
emcli search_patches -product=15991 -release=8191035000 -platform=2000
```

Output:

9561331	Generic PLATFORM	-	10.3.5	Oracle WebLogic Server	10.3.5	Generic
ic	American English		Recommended			Generic Platform

5. Create a patch-target map (properties) file using the vi editor, and supply information like Patch ID, Release ID, and Platform ID, Language ID, and so on. Here is a sample properties file:

```
vi create.props
```

```
patch.0.patch_id=9561331
patch.0.release_id=8191035000
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=/Farm01_soa_domain/soa_domain
patch.0.target_type=weblogic_domain
```

6. Run the following command to create the plan, and supply the newly created properties file (create.props) as input:

```
emcli create_patch_plan -name=demo1 -input_file=data:create.props
```

Output:

The Patch Plan "demo1" is successfully created.

7. To view the user-editable fields of an existing plan, and save the output to a properties file run the following command:

```
emcli get_patch_plan_data -name=demo1 >set.props
```

```
vi set.props
```

Output:

```
name=demo1
description=
deployment_date=
planPrivilegeList=VIEWER:ADMIN:VIEW;OPER:ADMIN:VIEW;DESIGNER:ADMIN:VIEW
```

```
patch.0.patch_id=9561331
patch.0.release_id=8191035000
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=/Farm01_soa_domain/soa_domain
patch.0.target_type=weblogic_domain
```

```
deploymentOptions.StageLocation=%emd_emstagedir%
deploymentOptions.AdvancedOPatchOptions=AllNodes
deploymentOptions.StagePatches=true
deploymentOptions.rollbackMode=false
```

8. Edit the properties file (set.props) using any editor to change the rollback mode to true:

```
name=demo1
description=
deployment_date=
planPrivilegeList=VIEWER:ADMIN:VIEW;OPER:ADMIN:VIEW;DESIGNER:ADMIN:VIEW
```

```
patch.0.patch_id=9561331
patch.0.release_id=8191035000
patch.0.platform_id=2000
patch.0.language_id=0
patch.0.target_name=/Farm01_soa_domain/soa_domain
patch.0.target_type=weblogic_domain
```

```
deploymentOptions.StageLocation=%emd_emstagedir%
deploymentOptions.AdvancedOPatchOptions=AllNodes
deploymentOptions.StagePatches=true
deploymentOptions.rollbackMode=true
```

9. To save the patch plan with the new edited data, run the following command:

```
emcli set_patch_plan_data -name=demo1 -input_file=data:set.props
```

Output:

It is successfully updating deployment options from the patch plan.

10. To verify the status of the patch plan, run the following EM CLI command:

```
emcli show_patch_plan -name=demo1 -info
```

Output:

```
<plan>
  <planDetails>
    <plan_id>EDD74FFF006DD0EE6D28394B8AAE</plan_id>
    <name>demo1</name>
    <type>PATCH</type>
    <description/>
    <conflict_check_date>Tue Feb 21 18:04:04 PST 2012</conflict_check_date>
    <conflict_check_date_ms>1329876244000</conflict_check_date_ms>
    <is_deployable>1</is_deployable>
    <plan_status>CONFLICTS</plan_status>
    <review_status>CONFLICT_FREE</review_status>
    <created_date>Tue Feb 21 17:40:47 PST 2012</created_date>
    <created_date_ms>1329874847000</created_date_ms>
    <created_by>SYSMAN</created_by>
    <last_updated_date>Tue Feb 21 17:58:29 PST 2012</last_updated_date>
    <last_updated_date_ms>1329875909000</last_updated_date_ms>
    <last_updated_by><USERNAME></last_updated_by>
    <grant_priv>yes</grant_priv>
```

```
<user_plan_privilege>FULL</user_plan_privilege>
<see_all_targets>N</see_all_targets>
<homogeneousGroupLabel>Oracle WebLogic Domain 10.3.5.0 (Linux x86-64)</homogeneousGroupLabel>
<executeGuid/>
<executeUrl/>
<planDetails/>
```

11. After you have created and updated the patch plan with all the relevant data, you can submit your patch plan in the following sequence of modes. The EM CLI command used to submit the patch plan is:

```
emcli submit_patch_plan -name=demo1 -action=analyze
```

Output:

The action "analyze" is successfully submitted on the Patch Plan "demo1", now "analyze" is in progress.

12. To verify the status of the patch plan submitted, run the following EM CLI command:

```
emcli show_patch_plan -name=demo1 -info
```

Output:

```
<plan>
  <planDetails>
    <plan_id>EDD74FFF006DD0EE6D28394B8AAE</plan_id>
    <name>demo1</name>
    <type>PATCH</type>
    <description/>
    <conflict_check_date>Tue Feb 21 18:04:04 PST 2012</conflict_check_date>
    <conflict_check_date_ms>1329876244000</conflict_check_date_ms>
    <is_deployable>1</is_deployable>
    <plan_status>CONFLICTS</plan_status>
    <review_status>CONFLICT_FREE</review_status>
    <created_date>Tue Feb 21 17:40:47 PST 2012</created_date>
    <created_date_ms>1329874847000</created_date_ms>
    <created_by>USERNAME</created_by>
    <last_updated_date>Tue Feb 21 17:58:29 PST 2012</last_updated_date>
    <last_updated_date_ms>1329875909000</last_updated_date_ms>
    <last_updated_by>USERNAME</last_updated_by>
    <grant_priv>yes</grant_priv>
    <user_plan_privilege>FULL</user_plan_privilege>
    <see_all_targets>N</see_all_targets>
    <homogeneousGroupLabel>Oracle WebLogic Domain 10.3.5.0 (Linux x86-64)</homogeneousGroupLabel>
    <executeGuid/>
    <executeUrl/>
  </planDetails>
```

13. To check if there are any conflicts, run the following command:

```
emcli show_patch_plan -name=planName -analysisResults
```

Output:

```
<plan_status>CONFLICTS</plan_status>
```

You can verify the plan you have created by logging in to Enterprise Manager Cloud Control, from **Enterprise** menu, select **Provisioning and Patching**, then select **Patches and Updates**. On the home page, you will see the patch plan **demo1** that you have created using the command line as follows:

agent	Review Analysis	Patch	Not Specified	SYSMAN	Yes	Full
demo	Review Analysis	Patch	Not Specified	SYSMAN	Yes	Full
demo1	Review Analysis	Patch	Not Specified	SYSMAN	Yes	Full
hh	Review Analysis	Patch	Not Specified	SYSMAN	Yes	Full
ss	Review Analysis	Patch	Not Specified	SYSMAN	Yes	Full
vis4	Successfully Analyzed	Patch	Not Specified	SYSMAN	Yes	Full
vis5	Deployed Successfully	Patch	Not Specified	SYSMAN	Yes	Full

You can resolve the conflicts using the UI, and then submit the patch plan.

- Run the command `show_patch_plan` after resolving the conflicts to verify the status of the plan as follows:

Output:

```
<plan>
  <planDetails>
    <plan_id>EDD74FFF006DD0EE6D28394B8AAE</plan_id>
    <name>demo</name>
    <type>PATCH</type>
    <description/>
    <conflict_check_date>Tue Feb 21 18:04:04 PST 2012</conflict_check_date>
    <conflict_check_date_ms>1329876244000</conflict_check_date_ms>
    <is_deployable>1</is_deployable>
    <plan_status>INPROGRESS</plan_status>
    <review_status>CONFLICT_FREE</review_status>
    <created_date>Tue Feb 21 17:40:47 PST 2012</created_date>
    <created_date_ms>1329874847000</created_date_ms>
    <created_by>USERNAME</created_by>
    <last_updated_date>Tue Feb 21 17:58:29 PST 2012</last_updated_date>
    <last_updated_date_ms>1329875909000</last_updated_date_ms>
    <last_updated_by>USERNAME</last_updated_by>
    <grant_priv>yes</grant_priv>
    <user_plan_privilege>FULL</user_plan_privilege>
    <see_all_targets>N</see_all_targets>
    <homogeneousGroupLabel>Oracle WebLogic Domain 10.3.5.0 (Linux x86-64)</homogeneousGroupLabel>
    <executeGuid>BA8E3904DDB36CFFE040F00A5E644D13</executeGuid>
    <executeUrl>/em/console/paf/procedureStatus?execution-GUID=BA8E3904DDB36CFFE040F00A5E644D13</executeUrl>
  </planDetails/>
```

- Run the following command to determine the status of the patch plan execution:

```
emcli get_instance_status -instance=BA8E3904DDB36CFFE040F00A5E644D13
```

Output:

```
BA8E3904DDB36CFFE040F00A5E644D13, PatchOracleSoftware, demo1_Analysis_Tue Mar 06
02:08:02 PST 012, EXECUTING
```

- After a successful analysis, you can deploy/prepare the patch plan. To do so, run the following command with action **deploy**:

```
emcli submit_patch_plan -name=demo1 -action=deploy
```

Output:

```
The action "deploy" is successfully submitted on the Patch Plan "demo1", now "de-
ploy" is in progress
```

- Use the Cloud Control UI to see if the submitted plan has successfully been deployed. Alternately, you can verify the same using the EM CLI command:

```
emcli get_job_execution_detail -execution=79CAF6A6DAFCFEE6654C425632F19411 -xml
```


Creating a New Generic Component by Associating a Zip File

To upload a zip file as a new component, follow these steps:

- [Step 1: Identifying the Parent Folder in Software Library](#)
- [Step 2: Creating a Generic Component Entity](#)
- [Step 3: Associating a Zip File to the Generic Component](#)
- [Step 4: Verifying the Newly Created Entity](#)

Step 1: Identifying the Parent Folder in Software Library

Any new entity created in Software Library must be placed in a folder. You can either choose an existing folder, or create a new one. To do so, follow these sections:

- [Creating a New Folder](#)
- [Choosing an Existing Folder](#)

Creating a New Folder

To create a new folder, the parent folder should be identified. If the parent folder is the **root** folder (displayed as the top level "Software Library" folder), then use the following EM CLI verb:

```
emcli create_swlib_folder  
-name="myFolder"  
-desc="myFolder description"  
-parent_id=ROOT
```

Output:

Folder myFolder is created in Software Library folder, identifier is oracle:default-Service:em:provisioning:1:cat:C771B5A38A484CE3E40E50AD38A69D2.

You can use the identifier of the newly created folder that is part of the output message when creating or modifying entities, or for creating other sub-folders.

Choosing an Existing Folder

To choose an existing folder, you can use either of the following approaches:

- [Approach 1: Using Enterprise Manager UI](#)
- [Approach 2: Using Enterprise Manager Command Line Interface](#)

Approach 1: Using Enterprise Manager UI

Follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library home page, from **View** menu select **Columns**, and then select **Internal ID**. By default, the Internal ID column is hidden.
3. Copy the Internal ID column value of the folder you want.

Approach 2: Using Enterprise Manager Command Line Interface

Use the following EM CLI verb:

```
emcli list_swlib_folders
-parent_id=ROOT
-show_folder_id
```

Output:

```
Java EE Provisioning, Java EE Application Provisioning Entities, oracle:defaultService:em:provisioning:1:cat:C771B5AAF4A4EED9E040E50AD38A6E98
MultiOMS, List of Oracle shipped Directives, oracle:defaultService:em:provisioning:1:cat:C771B5AAF1ACEED9E040E50AD38A6E98
myFolder, myFolder description, oracle:defaultService:em:provisioning:1:cat:C771B5A38A484CE3E040E50AD38A69D2
OSBProvisioning, OSBProvisioning Entities, oracle:defaultService:em:provisioning:1:cat:C771B5AAF3F1EED9E040E50AD38A6E98
.....
```

If the folder you want to access is a sub-folder of myFolder, then use the following verb to list the sub-folders by specifying the identifier of myFolder, as follows:

```
emcli list_swlib_folders
-parent_id='oracle:defaultService:em:provisioning:1:cat:C771B5A38A484CE3E040E50AD38A69D2'
-show_folder_id
```

Output:

```
mySubFolder, mySubFolder description, oracle:defaultService:em:provisioning:1:cat:C771B5A38A494CE3E040E50AD38A69D2
```

Step 2: Creating a Generic Component Entity

To create an entity of type Component and subtype Generic Component under mySubFolder folder, follow these sections:

- [Step 1a. Identifying the Entity Type](#)
- [Step 1b. Identifying the Entity Subtype](#)
- [Step 2. Creating a Generic Component Entity](#)

Step 1a. Identifying the Entity Type

To list all the available types in Software Library, use the following verb:

```
emcli list_swlib_entity_types
-show_type_id
```

Output:

```
Component, COMP_Component
Directives, COMP_Directives
Bare Metal Provisioning, BMPType
Virtualization, Virtualization
```

Step 1b. Identifying the Entity Subtype

To list all the subtypes for the component type, use the following verb:

```
emcli list_swlib_entity_subtypes
-entity_type_id=COMP_Component
-show_subtype_id
```

Output:

```
Generic Component, SUB_Generic
Oracle Database Software Clone, SUB_OracleDB
Configuration Template, SUB_ConfigTpl
SUB_OracleAS
Self Update, SUB_SelfUpdate
Oracle Clusterware Clone, SUB_OracleCRS
Service Bus Resource, SUB_OSBResource
Oracle Software Update, SUB_OraSoftUpdate
Java EE Application, SUB_JavaEEApplication
Installation Media, SUB_InstallationMedia
Database Template, SUB_DbCreateTemplate
Database Provisioning Profile, SUB_DbProfile
WebLogic Domain Provisioning Profile, SUB_FMWBundle
WebLogic Domain Clone, SUB_WLSTemplate
Oracle Middleware Home Gold Image, SUB_FMWImage
```

Step 2. Creating a Generic Component Entity

To create a generic component, run the following verb:

```
emcli create_swlib_entity
-name=myEntity
-type=COMP_Component
-subtype=SUB_Generic
-folder_id='oracle:defaultService:em:provisioning:
1:cat:C771B5A38A494CE3E040E50AD38A69D2'
-desc='myEntity description'
-attr="PRODUCT:Example"
-attr="PRODUCT_VERSION:3.1"
-attr="VENDOR:Example Corp"
-note='first comment for myEntity
```

Note:

The type and subtype options are optional when creating a Generic Component, but has been used explicitly for this illustration.

Output:

```
Entity 'myEntity' is created in 'mySubFolder' folder, identifier is 'oracle:default-
Service:em:provisioning:1:cmp:COMP_Component:SUB_Gener-
ic:C77200CA9DC1E7AAE040E50AD38A1599:0.1'
```



Note:

You can use the identifier of the newly created entity that is part of the output message when uploading files or modifying the entity.

To verify the newly created entity, use the following verb:

```
emcli list_swlib_entities
-name=myEntity
-folder_id='oracle:defaultService:em:provisioning:
1:cat:C771B5A38A494CE3E040E50AD38A69D2'
```

Output:

```
myEntity,0.1,myEntity description,Ready,Component,Generic Component,Untest-
ed,SYSMAN
```

Step 3: Associating a Zip File to the Generic Component

To upload a zip file to an existing entity myEntity, use the following verb:

```
emcli upload_swlib_entity_files
-entity_rev_id='oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Gener-
ic:C77200CA9DC1E7AAE040E50AD38A1599:0.1'
-file="/scratch/user1/patch13653908.zip;newfile1.zip"
-host="host.us.example.com"
-credential_name=mycred11
-credential_owner=sysman
```



Note:

A new revision of the entity myEntity will be created after the upload is complete.

Output:

```
Upload of file(s) initiated, this may take some time to complete...
Upload of file(s) completed successfully.
Entity 'myEntity (0.2)' in 'mySubFolder' folder has been created, identifier is 'ora-
cle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Gener-
ic:C77200CA9DC1E7AAE040E50AD38A1599:0.2'.
```

Alternately, to refer to a zip file present in an **HTTP** reference location, say myScripts, use the following verb:

```
emcli refer_swlib_entity_files
-entity_rev_id='oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Gener-
ic:C77200CA9DC1E7AAE040E50AD38A1599:0.1'
-file='scripts/perl/script1.pl;new_script.pl'
-refer_storage='myScripts;Http'
```

Output:

Entity 'myEntity (0.2)' in 'mySubFolder' folder has been created, identifier is 'oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Gener-ic:C77200CA9DC1E7AAE040E50AD38A1599:0.2'.

Step 4: Verifying the Newly Created Entity

To verify the newly created entity, use the following verb:

```
emcli list_swlib_entities
-name=myEntity
-folder_id='oracle:defaultService:em:provisioning:
1:cat:C771B5A38A494CE3E040E50AD38A69D2'
```

Output:

myEntity,0.1,myEntity description,Ready,Component,Generic Component,Untested,USERNAME

Migrate and Remove a Software Library Storage Location

This use case describes how to Migrate and Remove an existing storage location. In the following example, a Software Library Upload File storage location has already been configured. Files can be migrated from one storage location to another, of the same or different storage type.

In the following example, the first storage location is named 'firstLoc' and is an OMS Agent File System storage type. The second storage location is an OMS Shared File System storage location, named 'secondLoc'.

- [Step 1: Adding a Destination Storage Location for Migrating Files](#)
- [Step 2: Migrate and Remove an existing storage location](#)

Step 1: Adding a Destination Storage Location for Migrating Files

Any file associated with an entity in Software Library can exist in only one storage location. Removal of an Upload File storage location is always preceded by migration of files to another existing Upload File storage location. As part of this step, a destination storage location of OMS Shared File System storage type is created with sufficient disk space for keeping a copy of all the files in the source storage location, 'firstLoc'.

Add Destination OMS Shared File System Storage Location

```
emcli add_swlib_storage_location
-name=secondLoc
path=/u01/swlib
type=omsShared
```

Sample Output:

Location configuration initiated, this may take some time to complete...
Location 'secondLoc' created.

List OMS Shared File System Storage Location

To verify the creation of the new storage location use the following command.

```
emcli list_swlib_storage_locations
type=0msShared
Sample Output:
secondLoc, /u01/swlib/, Active
```

Step 2: Migrate and Remove an existing storage location

To remove the storage location 'firstLoc' and migrate all the files to storage location 'secondLoc', use the following command.

```
emcli remove_swlib_storage_location
name=firstLoc
type=0msAgent
migrate_to_loc=secondLoc -migrate_to_type=0msShared
```

Sample Output:

Job 'MigrateEntityFiles_1352113174929' has been successfully submitted for migrating the files.

The location 'firstLoc' will be removed on successful execution of the job. You can see the detail of the job execution by navigating to **Enterprise, Jobs**, and then, **Activity Page**.

Verifying the Status of Source OMS Agent File System Storage Location

Immediately after initiating the migrate and remove operation for storage location 'firstLoc', the location will be marked 'Inactive' to stop any new file uploads to this location. To verify the status of the location, use the following command.

```
emcli list_swlib_storage_locations
-type=0msAgent
```

Sample Output:

```
firstLoc, /u01/swlib/, Inactive
```

Once the migrate job is complete, the 'firstLoc' location is not listed, as it is removed.

Adding ATS Service Test from Using EM CLI

This use case describes how you can customize an existing ATS Service Test instance, available in the Test Repository, using custom databank.



Note:

Ensure that you have uploaded an ATS Service Test type to the Test Repository before proceeding with this procedure. For more information on this, see Configuring and Using Services chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

To create an ATS Test instance using the service test available in the repository, and to customize the test by applying a custom databank, follow these steps:

1. Create an ATS Service Test instance called my_service for the existing service:

```
emcli apply_template_tests -targetName='my_service' -targetType=generic_service -
input_file=template:'my_template.xml' -swlibPath='/service/test/entity'
```

Where,

-input_file=template:'my_template.xml' contains all the ATS test related information.

-swlibPath: holds the path to retrieve the ATS Zip file from software library.

Here is an example for a sample input file:

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<transaction-template template_type="generic_service" xmlns="template">
  <variables/>
  <transactions>
    <mgmt_bcn_transaction>
      <mgmt_bcn_txn_with_props>
        <mgmt_bcn_txn is_representative="true" name="ats91_with_databanks"
monitoring="true" txn_type="OATS"/>
        <properties>
          <property name="Collection Interval" num_value="5.0"
prop_type="2" encrypt="false"/>
          <property name="scriptDescription" string_value="[1] Oracle Se-
cure Enterprise Search&#xA;[2] Oracle Secure Enterprise Search&#xA;[3] No Ti-
tle&#xA;[4] No Title" prop_type="1" encrypt="false"/>
          <property name="fileUploadTime" string_value="2010-09-10
13:49:46.0" prop_type="1" encrypt="false"/>
          <property name="OpenScriptJwgName" string_value="ats91_with_data-
banks.zip" prop_type="1" encrypt="false"/>
          <property name="thinkTimeMaxSec" num_value="5.0" prop_type="2"
encrypt="false"/>
          <property name="usageOptions" string_value="userDefined"
prop_type="1" encrypt="false"/>
          <property name="thinkTimeOption" num_value="1.0" prop_type="2"
encrypt="false"/>
          <property name="fileSize" string_value="93699" prop_type="1" en-
crypt="false"/>
          <property name="beaconDistributionOverride" string_value="peo-
ple=1;middle=1" prop_type="1" encrypt="false"/>
          <property name="FilePropertyValue" prop_type="7" encrypt="false"/>
          <property name="databankFilesJar" prop_type="7" encrypt="false"/>
          <property name="databankFiles" string_value="people,people.csv,
3;middle,middle.csv,3;" prop_type="1" encrypt="false"/>
          <property name="granularity" string_value="transaction"
prop_type="1" encrypt="false"/>
          <property name="thinkTimeMinSec" num_value="0.0" prop_type="2"
encrypt="false"/>
          <property name="oatsVersion" string_value="9.1.0" prop_type="1"
encrypt="false"/>
          <property name="databankValues" string_value="people.first-
name=yang.,people.lastName=wang,middle.middlename_col=x." prop_type="1" en-
crypt="false"/>
          <property name="modules" string_value="oracle.oats.scripting.mod-
ules.utilities;version=2.4.0&#xA;oracle.oats.scripting.modules.http;ver-
sion=2.4.0&#xA;oracle.oats.scripting.modules.basic;version=2.4.0" prop_type="1"
encrypt="false"/>
          <property name="databankAliasMapping" string_value="middle=mid-
dle.csv,people=people.csv" prop_type="1" encrypt="false"/>
```

```

        </properties>
        <per_bcn_properties/>
    </mgmt_bcn_txn_with_props>
    <steps_defn_with_props/>
    <stepgroups_defn/>
    <txn_thresholds/>
    <step_thresholds/>
    <stepgroup_thresholds/>
</mgmt_bcn_transaction>
</transactions>
</transaction-template>

```

2. Upload new databank file called ATSTest1 to ATS test type for the service my_service:

```
emcli upload_ats_test_databank_file -name='my_service' -type='generic_service' -
testname='ATSTest1' -testtype='OATS' -databankAlias='alias1' -input_file='data-
bank:databankFile.csv'
```

Where,

-input_file=databank: contains the databank file path

3. Enable the test ATSTest1 for the service my_service:

```
emcli enable_test -name="my_service" -type="generic_service" -testname="ATST-
est1" -testtype="OATS"
```

Limitations of Using Enterprise Manager Command Line Interface

Following are the limitations of using EM CLI for running the deployment procedures:

- You cannot add or edit steps and phases using EM CLI commands. To do so, you must log in to Cloud Control, and follow the steps described in the section Section 33.2.1.
- You cannot define new variables to be used in the deployment procedures through EM CLI, this can be done only through the Cloud Control UI. For more information about procedure variables, see Section 32.3.2.
- You cannot track the detailed execution info (such as failures) of an instance through EM CLI, which is possible through the Cloud Control UI.
- To set the *My Oracle Support* preferred credentials, you must log in to the Enterprise Manager Cloud Control. There is no command line option to do so.

B

Checking Host Readiness Before Provisioning or Patching

This appendix describes the settings you must make on the hosts before you can use them for provisioning and patching tasks. In particular, this appendix covers the following:

- [Setting Up User Accounts Before Provisioning](#)
- [Shell Limits](#)
- [Root Setup \(Privilege Delegation\)](#)
- [Environment Settings](#)
- [Storage Requirements](#)
- [Installation Directories and Oracle Inventory](#)

Setting Up User Accounts Before Provisioning

To use a host for provisioning a database, you must ensure that groups such as oinstall, dba, oper, and asmadmin are set up. Also, the user running these provisioning tasks must be added to these groups. To create the following groups, and ensure that the host user is part of these groups, you can run the following commands:

- To create the database groups:
 - `groupadd oinstall`
 - `groupadd dba`
 - `groupadd oper`
 - `groupadd asmadmin`
- To add a host user to these groups, run the following command, and enter the password when prompted.

```
useradd -u 500 -g oinstall -G dba,oper,asmdba oracle
```

Where,

 - u option specifies the user ID.
 - g option specifies the primary group, which must be the Oracle Inventory group, for example oinstall.
 - G option specifies the secondary groups, which must include the OSDBA group, and, if required, the OSOPER and ASMDBA groups, for example, dba, asmdba, or oper.

Configuring SSH

In case of a clustered environment, to configure SSH on each node in a cluster, you must log in as an oracle user, and run the following commands on every node:

```
su - oracle
mkdir ~/.ssh
chmod 700 ~/.ssh
/usr/bin/ssh-keygen -t rsa # Accept the default settings
```

Shell Limits

To improve the performance of the software on Linux systems, increase the following shell limits for the Oracle software owner users such as crs, oracle, asm, and so on. To do so, run the following commands:

- Add the following values into the `limits.conf` file located under the `/etc/security/` directory:
 - `oracle soft nproc 2047`
 - `oracle hard nproc 16384`
 - `oracle soft nofile 1024`
 - `oracle hard nofile 65536`
- Add the following line into the `/etc/pam.d/login` file, or edit the `/etc/pam.d/login` file to include the following if it does not exist already:
`session required pam_limits.so`

Root Setup (Privilege Delegation)

Provisioning Applications require some of the scripts to be run as a super user. To do so, you must ensure that host user has root privileges. To authorize other users' root privileges, you can use the authentication utilities such as SUDO, PowerBroker, and so on. This support is offered in Cloud Control using the Privilege Delegation mechanism. Technically, 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 more information about configuring Privilege Delegation Settings, see [Setting Up Credentials](#).

Environment Settings

Meet the following recommended host settings before proceeding with the provisioning tasks:

- [Kernel Requirements](#)
- [Node Time Requirements](#)
- [Package Requirements](#)
- [Memory and Disk Space Requirements](#)

- Network & IP Address Requirements

 Note:

For details about all the recommended parameters, refer the following link:
<http://www.oracle.com/technetwork/topics/linux/validated-configurations-085828.html>

Kernel Requirements

Enter the commands displayed in the following table to view the current values of the kernel parameters. Make a note of the current values and identify any values that you must change. To change any of the existing values, you will have to add or edit the variable values in the `/etc/sysctl.conf` file.

 Note:

To change the current kernel parameters, run the following command with root user privileges:

```
/sbin/sysctl -p
```

Parameter	Command
semmsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm This command displays the details of the shared memory segment sizes.
file-max	# /sbin/sysctl -a grep file-max This command displays the maximum number of file handles.
ip_local_port_range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.
rmem_default	# /sbin/sysctl -a grep rmem_default
rmem_max	# /sbin/sysctl -a grep rmem_max
wmem_default	# /sbin/sysctl -a grep wmem_default
wmem_max	# /sbin/sysctl -a grep wmem_max

**Note:**

For more information about the Kernel requirements, see the *Oracle Database Installation Guide* available in the following location: http://www.oracle.com/pls/db112/portal.portal_db?selected=11&frame=#linux_installation_guides

Node Time Requirements

In case of a clustered environment, ensure that each member node of the cluster is set as closely as possible to the same date and time. To do so, Oracle recommends using the Network Time Protocol (NTP) feature available in your operating systems, with all nodes using the same reference Network Time Protocol server.

For Oracle Cluster Time Synchronization Service (ctssd) to synchronize the times of the Oracle RAC nodes, NTP must be configured. If you are using NTP, then do the following:

1. Add the `-x` option to the `/etc/sysconfig/ntpd` file, and restart `ntpd` as follows:

```
OPTIONS="-x -u ntp:ntp -p /var/run/ntpd.pid"
```
2. Restart Network Time Protocol server:

```
# service ntpd restart
```
3. Check the configuration level as follows:

```
chkconfig --level 35 nscd on
```
4. Start the Name Service Cache Daemon (nscd):

```
service nscd start
```

Package Requirements

Run the following command as a root user to ensure that you have the required packages installed:

```
rpm -q binutils elfutils-libelf elfutils-libelf-devel glibc glibc-common  
glibc-devel gcc gcc-c++ libaio libaio-devel libstdc++ libstdc++-devel make  
compat-libstdc++ sysstat unixODBC unixODBC-devel iscsi-initiator-utils  
libgcc
```

If the packages are not installed, then refer the following link to download and install the required packages:

<http://www.oracle.com/technetwork/topics/linux/validated-configurations-085828.html>

Memory and Disk Space Requirements

Ensure that the host meets the following memory requirements:

1. A minimum of least 1 GB of physical RAM should be available. To determine the current physical RAM size on your host, run the following command:

```
grep MemTotal /proc/meminfo
```

- The following table describes the relationship between the installed RAM and the configured swap space recommendation:

Available RAM	Swap Space Requirements
Between 1 GB and 2 GB	1.5 times the size of RAM
Between 2 GB and 8 GB	Equal to the size of RAM
More than 8 GB	0.75 times the size of RAM

- To determine the amount of disk space available in the /tmp directory, run the following command:

```
df -kh /tmp
```

Network & IP Address Requirements

In case of a clustered environment, ensure that each node has at least two network adapters or network interface cards (NICs). One for the public network interface, and the other for the private network interface (the interconnect).

Following are the network configuration requirements:

Public Network Interface	Private Network Interface
The public interface names associated with the network adapters for each network must be the same on all nodes.	The private interface names associated with the network adapters should be the same on all nodes.
Each network adapter must support TCP/IP	The interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better required). Note: For the private network, the endpoints of all designated interconnect interfaces must be completely reachable on the network. There should be no node that is not connected to every private network interface. You can test whether an interconnect interface is reachable using a ping command.

Before starting the installation, you must have the following IP addresses available for each node:

- An IP address with an associated host name (or network name) registered in the DNS for the public interface. If you do not have an available DNS, then record the host name and IP address in the system hosts file, /etc/hosts.
- One virtual IP (VIP) address with an associated host name registered in a DNS. If you do not have an available DNS, then record the host name and VIP address in the system hosts file, /etc/hosts.
- A private IP address with a host name for each private interface.

For example, for a two node cluster where each node has one public and one private interface, you might have the configuration shown in the following table for your network interfaces, where the hosts file is /etc/hosts:

Node	Host Name	Type	IP Address	Registered In
node1	node1	Public	143.46.43.100	DNS (if available, else the hosts file)
node1	node1-vip	Virtual	143.46.43.104	DNS (if available, else the hosts file)
node1	node1-priv	Private	10.0.0.1	Hosts file
node2	node2	Public	143.46.43.101	DNS (if available, else the hosts file)
node2	node2-vip	Virtual	143.46.43.105	DNS (if available, else the hosts file)
node2	node2-priv	Private	10.0.0.2	Hosts file

To enable VIP failover, the configuration shown in the preceding table defines the public and VIP addresses of both nodes on the same subnet, 143.46.43.

Storage Requirements

There are two ways of storing Oracle Clusterware files:

- Oracle Automatic Storage Management (Oracle ASM): You can install Oracle Clusterware files (Oracle Cluster Registry and voting disk files) in Oracle ASM disk groups.
- A supported shared file system: Supported file systems include the NFS & OCFS.

The following table describes the various storage options for Oracle Clusterware and Oracle RAC:

Storage Option	OCR and Voting Disk Files	Oracle Clusterware binaries	Oracle RAC binaries	Oracle Database Files	Oracle Recovery Files
Oracle Automatic Storage Management (Oracle ASM)	Yes	Yes	No	Yes	Yes
Note: Loopback devices are not supported for use with Oracle ASM					
Oracle Automatic Storage Management Cluster File System (Oracle ACFS)	No	Yes	Yes	No	No
Local file system	No	Yes	Yes	No	No
NFS file system on a certified NAS filer	Yes	Yes	Yes	Yes	Yes
Note: Direct NFS does not support Oracle Clusterware files.					
Shared disk partitions (block devices or raw devices)	Not supported by OUI or ASMCA, but supported by the software. They can be added or removed after installation.	No	No	Not supported by OUI or ASMCA, but supported by the software. They can be added or removed after installation.	No

The following table displays the File System Volume Size requirements:

Oracle Clusterware Shared File System Volume Size Requirements

File Types Stored	Number of Volumes	Volume Size
Voting disks with external redundancy	3	At least 300 MB for each voting disk volume.
Oracle Cluster Registry (OCR) with external redundancy	1	At least 300 MB for each OCR volume
Oracle Clusterware files (OCR and voting disks) with redundancy provided by Oracle software.	1	At least 300 MB for each OCR volume At least 300 MB for each voting disk volume

Oracle RAC Shared File System Volume Size Requirements

File Types Stored	Number of Volumes	Volume Size
Oracle Database files	1	At least 1.5 GB for each volume
Recovery files	1	At least 2 GB for each volume
Note: Recovery files must be on a different volume than database files		

Installation Directories and Oracle Inventory

Ensure that the installation directories where you plan to provision the Oracle Products are clean. As per Optimal Flexible Architecture (OFA) standards, Oracle base directory should be available in the following path:

```
/mount_point/app/oracle_sw_owner
```

Where, mount_point is the mount point directory for the file system that will contain the Oracle software.

Note:

Ensure that the user performing the installation has write access on the mount points. To verify that the user has the required permissions, run the following command:

```
chown -R oracle:oinstall <mount point>
```

For example:

If the permission is denied while mounting:

```
[root@node2-pub ~]# mkdir -p /u01/app/test
```

```
[root@node2-pub ~]# permission denied
```

To resolve the permission issue, run the following command:

```
[root@node2-pub root]# chown -R oracle:oinstall /u01
```

C

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 as PAR files, importing PAR files, and so on. In particular, this appendix covers the following:

- [Overview of Provisioning Archive Files](#)
- [Overview of emctl partool Utility](#)
- [Checking Oracle Software Library](#)
- [Exporting Deployment Procedures](#)
- [Importing PAR Files](#)

Note:

Provisioning Archive Files that were exported from any Enterprise Manager version prior to 12c can not be imported into Enterprise Manager 12c.

Overview of Provisioning Archive Files

Provisioning Archive (PAR) files are archive files that contain a collection, or bundle of Deployment Procedures and Software Library entities that are used in numerous Life-cycle Management tasks like Provisioning and Patching applications.

In case of a Deployment Procedures, partool exports only the User-defined procedures, and not the Oracle-owned procedures. While exporting the User-defined procedure, the complete deployment procedure is not exported, only the customization (delta changes) are exported.

Also, note that in case of upgrade all the procedures that were created pre-12c can not be exported using the partool export utility.

Note:

For importing PAR files that contain Software Library entities, ensure that your Software Library is configured. For information on Configuring Software Library, see *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

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 Cloud Control. Under such cir-

cumstances, you might want to export the customized Deployment Procedures from one instance of Cloud Control, and deploy them to another instance of Cloud Control.

emctl partool utility is a tool offered by Cloud Control that helps you perform these functions using the command line interface. Essentially, 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 Cloud Control

The emctl partool utility is located in the \$ORACLE_HOME/bin directory.

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
```

[Table C-1](#) describes the additional options that can be used with the emctl partool utility.

Table C-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.
-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.

Table C-1 (Cont.) emctl partool Options

Option	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.

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
```

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](#)

Obtaining Deployment Procedure's GUID

To obtain the GUID of a Deployment Procedure using Cloud Control, follow these steps:

1. In Cloud Control, from the **Enterprise** menu select **Provisioning and Patching**, and then click **Procedure Library**.
2. On the Provisioning page, right click the deployment procedure name and from the menu select **Copy Link Location**.
3. Paste the copy the link to a notepad, and then search for **guid**.

For example: `https://adc2171248.us.example.com:14500/em/console/paf/procedureView?guid=B3B4B6C76AE46A67E040E50A65751782`

The GUID is B3B4B6C76AE46A67E040E50A65751782

Alternately you can use the following EMCLI command to retrieve the GUID of the procedure:

```
emcli get_procedures [-type={procedure type}] [-parent_proc={procedure associate with procedure configuration}]
```

Example:

```
emcli get_procedures -type=DemoNG -parent_proc=ComputeStepTest
```

Output Column:

GUID, Procedure type, name, display name, version, Parent procedure name



Note:

For information about setting up EMCLI, see *Oracle Enterprise Manager Command Line Interface Guide*.

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 Cloud 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  
"06B62B6ED5DA20BCE040578C850862A7,0C96E96D9818BC5FE040578C8508620F,  
09AEFF331025AAD0EE40578C85FB5772" -file $ENV{T_WORK}/  
tvmgf_partool_multi_dp.par -displayName "partool multi dp test" -descrip-  
tion "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.

Importing PAR Files

You can import PAR files using the command line interface or the graphical user interface offered by Cloud Control, that is, the console. This section explains the following:

- [Importing Using Command Line Interface](#)
- [Importing Using Cloud Control Console](#)

 Note:

Importing an existing PAR file (from the previous releases) into Enterprise Manager Cloud Control 12c is not supported. For example, you cannot import a Enterprise Manager 11g Grid Control Release 1 (11.1.0.0) to Enterprise Manager Cloud Control 12c.

Importing Using Command Line Interface

This section covers the following:

- [Importing Specific PAR File](#)
- [Importing All PAR Files](#)

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>
```

For example:

```
$ORACLE_HOME/bin/emctl partool deploy -parFile $ORACLE_HOME/sysman/  
prov/paf/asprov.par
```

**Note:**

- If Software Library or the procedure already exists in Enterprise Manager and you want to create a new revision of the PAR file, then you can use the `-force` attribute as follows:

```
$ORACLE_HOME/bin/emctl partool deploy -parFile $ORACLE_HOME/
sysman/prov/paf/asprov.par -force
```

- If you have multiple OMSes in your environment, then you need to run the `emctl partool` utility only once to deploy any PAR files or to perform other related operations.

While importing PAR files if the user procedure already exists in the set-up, then it will always import this procedure with revised (bumped up) version.

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
```

Importing Using Cloud Control Console

To import PAR files or deploy them to an OMS, you can use the `emctl partool` utility. Alternatively, you can import them from Cloud Control .

For importing the PAR files, follow these steps:

1. In Cloud Control, from the **Enterprise** menu select **Provisioning and Patching**, and then click **Procedure Library**.
2. On the Deployment Procedure Manager page, select the procedure, and from the drop down menu select **Import**, and then click **Go**.
3. On the Upload Procedure File page, select:
 - **Upload from Local Machine** to upload the PAR file from the local machine. Click **Browse** to select the PAR file. Click **Import** to import the file.
 - **Upload from Management Agent Machine** to select the Management Agent target. Enter the Normal or Privileged Host Credential details to access the Management Agent machine where the file is present, and then click **Import** to import the file.

**See Also:**

For a usecase on using the Upload Procedure File feature, see [Copying Customized Provisioning Entities from One Enterprise Manager Site to Another](#).

 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 C-1](#).

D

Understanding PXE Booting and Kickstart Technology

This appendix explains PXE booting and kickstart technology in the following section:

- [About PXE Booting and Kickstart Technology](#)
- [Subnet Provisioning Usecases](#)

About PXE Booting and Kickstart Technology

One of the key requirements of provisioning is the hardware server's ability to boot over the network instead of a diskette or CD-ROM. There are several ways computers can boot over a network, and Preboot Execution Environment (PXE) is one of them. PXE is an open industry standard supported by a number of hardware and software vendors. PXE is part of the "Wired for Management" (WfM) specification, which is part of a bigger PC98 specification defined by Intel and Microsoft in 1998. A detailed document on PXE specification can be found at <http://www.pix.net/software/pxeboot/archive/pxespec.pdf>.

PXE works with Network Interface Card (NIC) of the system by making it function like a boot device. The PXE-enabled NIC of the client sends out a broadcast request to DHCP server, which returns with the IP address of the client along with the address of the TFTP server, and the location of boot files on the TFTP server. The following steps describe how it works:

1. Target Machine (either bare metal or with boot sector removed) is booted.
2. The Network Interface Card (NIC) of the machine triggers a DHCP request.
3. DHCP server intercepts the request and responds with standard information (IP, subnet mask, gateway, DNS etc.). In addition, it provides information about the location of a TFTP server and boot image (pxelinux.0).
4. When the client receives this information, it contacts the TFTP server for obtaining the boot image.
5. TFTP server sends the boot image (pxelinux.0), and the client executes it.
6. By default, the boot image searches the pxelinux.cfg directory on TFTP server for boot configuration files on the TFTP server using the following approach:

First, it searches for the boot configuration file that is named according to the MAC address represented in lower case hexadecimal digits with dash separators. For example, for the MAC Address "88:99:AA:BB:CC:DD", it searches for the file 01-88-99-aa-bb-cc-dd.

Then, it searches for the configuration file using the IP address (of the machine that is being booted) in upper case hexadecimal digits. For example, for the IP Address "192.0.2.91", it searches for the file "C000025B".

If that file is not found, it removes one hexadecimal digit from the end and tries again. However, if the search is still not successful, it finally looks for a file named "default" (in lower case).

For example, if the boot file name is /tftpboot/pxelinux.0, the Ethernet MAC address is 88:99:AA:BB:CC:DD, and the IP address 192.0.2.91, the boot image looks for file names in the following order:

```
/tftpboot/pxelinux.cfg/01-88-99-aa-bb-cc-dd
/tftpboot/pxelinux.cfg/C000025B
/tftpboot/pxelinux.cfg/C000025
/tftpboot/pxelinux.cfg/C00002
/tftpboot/pxelinux.cfg/C0000
/tftpboot/pxelinux.cfg/C000
/tftpboot/pxelinux.cfg/C00
/tftpboot/pxelinux.cfg/C0
/tftpboot/pxelinux.cfg/C
```

7. The client downloads all the files it needs (kernel and root file system), and then loads them.
8. Target Machine reboots.

The Provisioning application uses Redhat's Kickstart method to automate the installation of Redhat Linux on target machines. Using kickstart, the system administrator can create a single file containing answers to all the questions that will usually be asked during a typical Red Hat Linux installation.

The host specific boot configuration file contains the location of the kickstart file. This kickstart file would have been created earlier by the stage directive of the OS image based on the input from user.

Subnet Provisioning Usecases

Following are examples of subnet provisioning usecases:

Subnet of size 256

IP Prefix: 192.168.1.0

Subnet Mask: 255.255.255.0

Covers IPs from 192.168.1.0 - 192.168.1.255

Subnet of size 16

IP Prefix: 192.168.1.0

Subnet Mask: 255.255.255.240

E

End-to-End Use Case: Patching Your Data Center

This appendix demonstrates how Enterprise Manager can be used to enable administrators to roll out patches across their data center.

The appendix contains the following sections:

- [The Challenge of Patching Your Data Center](#)
- [The Enterprise Manager Solution](#)
- [Executing the Example Scenario](#)

The Challenge of Patching Your Data Center

In any enterprise, a data center plays a critical role in keeping the IT functions alive and the business going. The data center may vary in size from one enterprise to another, but the fact that the data center is critical to the success of the business is clearly unquestionable.

The administrators in a data center carry out several data maintenance, data backup, and lifecycle management operations every day, and the challenges they face in carrying out these system management activities are sometimes immeasurable. These pain points become even more profound when the data centers span multiple geographical locations across multiple time zones.

One of the lifecycle management challenges that administrators regularly face is patching their entire ecosystem and keeping their data center secure and up to date. The requirement becomes even more complex when there are several types of patches, when it is difficult to identify the ones relevant to your data center, and when the entire patching operation is manual, error prone, and time consuming.

The Enterprise Manager Solution

The following sections provide a solution to the previously noted challenges that leverages the features of Enterprise Manager. The goal is to use a single, integrated patching workflow that not only helps you identify the patches relevant to your data center but also helps you download and roll them out in an unattended manner, and thereby ensure 100% compliance to your policies and standards.

The following is the basic flow of this use case:



Identify the Patches Relevant to Your Data Center

Use the patch recommendations offered by Enterprise Manager to identify the patches that are relevant to your data center. Patch recommendations are proactive notifications of potential system issues and recommendations that help you improve system performance and avert outages. The patches recommended for you are security patches and other patches based on your enterprise configuration.

Prepare, Test, and Certify the Patch Rollout Plan

Analyze your environment and verify if the targets in your data center can be patched. Once you are sure they can be patched, create a patch plan with the recommended patches, test the patches using the patch plan, diagnose and resolve all patch conflicts beforehand. Once the patch plan is deployable, certify the patch plan by converting it to a template.

Create a Change Activity Plan to Roll Out the Patches

Create a change activity plan to associate target types; create a series of tasks to carry out, including prepatching and postpatching tasks; select the patch plan template to use; prioritize the patching steps; and schedule the change activity plan for a formal rollout in your data center.

Monitor the Progress and Report the Status of the Change Activities

Monitor the progress of the various change activities, track the status of the patch rollout operation, and identify any drifts, and report the overall status to your higher management.

Executing the Example Scenario

The following table lists the tasks that will be performed in this example scenario, and the user roles that can perform the task.

Task	User Role
Create Administrators with the Required Roles	EM Super Administrator
Set Up the Infrastructure	EM_PATCH_DESIGNER
Analyze the Environment and Identify Whether Your Targets Can Be Patched	EM_PATCH_DESIGNER
Identify the Relevant Patches	EM_PATCH_DESIGNER
Create a Patch Plan, Test the Patches, and Certify the Patches	EM_PATCH_DESIGNER
Create a Change Activity Plan to Roll Out the Patches	EM_CAP_ADMINISTRATOR
Roll Out the Patches	EM_CAP_USER
Check and Report the Status of the Change Activities	EM_CAP_ADMINISTRATOR
Verify If the Targets Have Been Patched	EM_CAP_ADMINISTRATOR

Create Administrators with the Required Roles

Role: EM Super Administrator

[Table E-1](#) lists the roles based on which you can create administrators for the scenario described in this chapter.

Table E-1 Creating Administrators with the Required Roles

Enterprise Manager Role	Privileges
EM_PATCH_DESIGNER	CREATE_PATCH_PLAN, VIEW_ANY_PLAN_TEMPLATE
EM_CAP_ADMINISTRATOR	CREATE_JOB, CREATE_CAP_PLAN, BASIC_CAP_ACCESS
EM_CAP_USER	BASIC_CAP_ACCESS

For instructions to create administrators with these roles, see [Creating Enterprise Manager User Accounts](#).

Set Up the Infrastructure

Role: EM_PATCH_DESIGNER

Oracle recommends that you use the online patching mode for deployment of patches. Online patching mode is the default mode for patching in Enterprise Manager, and therefore, you do not have to manually set this up the first time. However, if you have set it to offline mode for a particular reason, and if you want to reset it to online mode, or if you want to verify that the online mode is indeed set, see [Setting Up the Infrastructure for Patching in Online Mode \(Connected to MOS\)](#).

In online mode, Enterprise Manager connects to My Oracle Support to download patches, patch sets, ARU seed data such as products, platforms, releases, components, certification details, and patch recommendations. For this purpose, Enterprise Manager uses the Internet connectivity you have on the OMS host to connect to My Oracle Support. However, if you have a proxy server set up in your environment, then you must register the proxy details. To register the proxy server details with Enterprise Manager, see [Setting Up the Infrastructure for Patching in Online Mode \(Connected to MOS\)](#).

Analyze the Environment and Identify Whether Your Targets Can Be Patched

Role: EM_PATCH_DESIGNER

Before creating a patch plan to patch your targets, Oracle recommends that you view the patchability reports to analyze the environment and identify whether the targets you want to patch are suitable for a patching operation. These reports provide a summary of your patchable and non patchable targets, and help you create deployable patch plans. They identify the problems with the targets that cannot be patched in your setup and provide recommendations for them.

Patchability reports are available for Oracle Database, Oracle WebLogic Server, and Oracle SOA Infrastructure targets.

To view the patchability reports, see [Analyze the Environment and Identify Whether Your Targets Can Be Patched](#).

Identify the Relevant Patches

Role: EM_PATCH_DESIGNER

View the Patch Recommendations region to identify the recommended and the relevant patches to be rolled out in your data center. Patches mentioned in the Patch Recommendation section are a collection of patches offered within MOS which can be applied as a group to one or more targets.

Using the Patch Recommendations region, you can drill down to a list of recommended patches, view their details, download them, or add them to a patch plan.

To view the recommended patches, see [About Patch Recommendations](#)

Create a Patch Plan, Test the Patches, and Certify the Patches

Role: EM_PATCH_DESIGNER

Create a patch plan with the recommended patches, test the patches using the patch plan, diagnose and resolve all patch conflicts beforehand. Once the patch plan is deployable, certify the patch plan by converting it to a template.

To create a patch plan, see [Creating a Patch Plan](#).

To access the newly created patch plan, see [Accessing the Patch Plan](#).

To add patches to the patch plan, to analyze and test the patches, and to save the patch plan as a patch template, follow Step (1) to Step (5) as outlined in the following URL, and then for Step (6), on the Review & Deploy page, click **Save as Template**. In the Create New Plan Template dialog, enter a unique name for the patch template, and click **Create Template**. For more information, see [Patching Software Deployments](#).

Create a Change Activity Plan to Roll Out the Patches

Role: EM_CAP_ADMINISTRATOR

Create a change activity plan identify the change activities, assign owners to activities, associate target types, create a series of tasks to carry out, including prepatching and postpatching tasks, select the patch plan template to use, prioritize the patching steps, and schedule the change activity plan for a formal rollout in your data center.

To do so, see [Creating a Change Activity Plan](#)

Roll Out the Patches

Role: EM_CAP_USER

Review the tasks assigned to you, monitor the task due date, complete any prepatching tasks, roll out the patch plan, complete all postpatching tasks, and update the task status.

To do so, see [Viewing My Tasks](#)

Check and Report the Status of the Change Activities

Role: EM_CAP_ADMINISTRATOR

Track the status of the tasks that are part of the change activity plan you created, and report the overall status to your higher management.

To do so, see [Managing a Change Activity Plan](#)

Verify If the Targets Have Been Patched

Role: EM_CAP_ADMINISTRATOR

Verify if the targets identified for patching have indeed been patched successfully with the selected patches.

To do so, run the Oracle-supplied configuration search titled *Search Patches Applied on Oracle Products* from the Configuration Search Library, as described in the following URL. Search for the patch ID that you applied to the targets. The search result lists all the targets with that patch ID. Verify if the targets on your list appear in the search result. To do so, see [Managing Configuration Searches](#).

F

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 Database Provisioning Issues](#)
- [Troubleshooting Linux Provisioning Issues](#)
- [Troubleshooting Patching Issues](#)
- [Troubleshooting Linux Patching Issues](#)
- [Frequently Asked Questions on Linux Provisioning](#)
- [Refreshing Configurations](#)
- [Reviewing Log Files](#)

Troubleshooting Database Provisioning Issues

This section provides troubleshooting tips for common database provisioning issues.

Grid Infrastructure Root Script Failure

See the details below.

Issue

Grid Infrastructure root script fails.

Description

After Grid Infrastructure bits are laid down, the next essential step is Grid Infrastructure root script execution. This is the most process intensive phase of your deployment procedure. During this process, the GI stack configures itself and ensures all subsystems are alive and active. The root script may fail to run.

Solution

1. Visit each node that reported error and run the following command on n-1 nodes:
`$GI_ORACLE_HOME/crs/install/rootcrs.pl -deconfig -force`
2. If the root script did not run successfully on any of the nodes, pass the `-lastNode` switch on nth node (conditionally) to the final invocation as shown below.
`$GI_ORACLE_HOME/crs/install/rootcrs.pl -deconfig -force -lastNode`

Now, retry the failed step from the Procedure Activity page.

SUDO Error During Deployment Procedure Execution

See the details below.

Issue

A SUDO error occurs while performing a deployment.

Description

While performing a deployment, all root-related operations are performed over sudo. To improve security, production environments tend to fortify sudo. Therefore, you may encounter errors related to sudo.

Solution

Make the following changes in your sudoer's file:

1. Remove entry `Default requiretty`, if it exists in your sudoer's file.
2. If sudoers file contains entry `Default env_reset`, add the following entries after this parameter:

```
Defaults env_keep="JRE_HOME PERL5LIB EMDROOT"
```

Prerequisites Checks Failure

See the details below.

Issue

Prerequisites checks fail when submitting a deployment procedure

Cause

Perform a meticulous analysis of output from prerequisite checks. While most prerequisite failures are automatically fixed, it is likely that the deployment procedure failed due to auto-fix environment requirements. Some likely cases are:

- Group membership for users that are not local to the system. Since users are registered with a directory service, even root access does not enable the deployment procedure to alter their attributes.
- Zone separation in Solaris. If the execution zone of deployment procedure does not have privilege to modify system attributes, auto-fix operations of the deployment procedure will fail.

Solution

Ensure that the deployment procedure has appropriate privileges.

Oracle Automatic Storage Management (Oracle ASM) Disk Creation Failure

See the details below.

Issue

Oracle ASM disk creation fails.

Cause

ASM disks tend to be used and purged over time. If an ASM instance is purged and physical ASM disks are left in their existing spurious state, they contain diskgroup information that can interfere with future ASM creation. This happens if the newly created ASM uses the same diskgroup name as exists in the header of such a raw disk. If such a spurious disk exists in the disk discovery path of the newly created ASM it will get consumed and raise unexpected error.

Solution

Ensure that disk discovery path is as restrictive as possible. Also, ASM disks should be zeroed out as soon as ASM is being purged. Deployment procedures that support post 11.2 RDBMS have elaborate checks to detect the use case and warn the user beforehand.

Oracle ASM Disk Permissions Error

See the details below.

Issue

Encountered an Oracle ASM Disk permissions error.

Description

Unlike NFS mounted storage wherein permissions set on any one node are visible throughout, ASM diskgroups require permissions to be set to each raw disk for all participating nodes.

Solution

For all participating nodes of the cluster, set 660 permissions to each raw disk being consumed.

Specifying a Custom Temporary Directory for Database Provisioning

To specify a temporary directory other than /tmp for placing binaries when provisioning databases, follow these steps:

1. Log in as a designer, and from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.

2. In the Database Procedures page, select the **Provision Oracle Database** Deployment Procedure and click **Create Like**.
3. In the Create Like Procedure page, in the General Information tab, provide a name for the deployment procedure.

In the Procedure Utilities Staging Path, specify the directory you want to use instead of /tmp, for example, /u01/working.

4. Click **Save**.

Use this deployment procedure for provisioning your Oracle databases.

Incident Creation When Deployment Procedure Fails

See the details below.

Issue

During deployment procedure execution, the steps to create database and Oracle ASM storage fails.

Solution

When a step in a deployment procedure executes successfully, it returns a positive exit code. If the step fails, and the exit code is not positive, it raises an incident which is stored in the OMS. All the associated log files and diagnosability information such as memory usage, disk space, running process info and so on are packaged and stored. You can access the Incident Console and package this information as a Service Request and upload it to My Oracle Support. Click Help on the Incident Manager page for more information on creating a new Service Request.

Reading Remote Log Files

In deployment procedure execution page, all remote log files relevant to the provisioning operation are displayed as hyperlinks in the job step. You can click on these hyperlinks and view the remote logs. The remote logs are stored in the OMS repository and can be accessed by My Oracle Support when troubleshooting.

Retrying Failed Jobs

See the details below.

Issue

Deployment procedure execution fails.

Solution

If you deployment procedure execution has failed, check the job run details. You can retry a failed job execution or choose to ignore failed steps. For example, if your deployment procedure execution failed due to lack of disk space, you can delete files and retry the procedure execution. For certain issues which may not affect the overall deployment procedure execution, such as cluvfy check failure, you may want to ignore the failed step and run the deployment procedure.

Retrying a job execution creates a new job execution with the status **Running**. The status of the original execution is unchanged.

To ignore a failed step and retry the procedure, follow these steps:

1. In the Procedure Activity page, click on the relevant procedure.
2. In the Job Status page, click on the status of the failed step.
3. In the Step Status page, click **Ignore**. In the Confirmation page, click **Yes**.
4. Click **Retry**. The failed step will be retried.

Troubleshooting Patching Issues

This section provides troubleshooting tips for common patching issues.

- [Oracle Software Library Configuration Issues](#)
- [My Oracle Support Connectivity Issues](#)
- [Host and Oracle Home Credential Issues](#)
- [Collection Issues](#)
- [Patch Recommendation Issues](#)
- [Patch Plan Issues](#)
- [Patch Plan Analysis Issues](#)
- [User Account and Role Issues](#)

Oracle Software Library Configuration Issues

This section describes the following patching issues:

- [Error Occurs While Staging a File](#)
- [Error Occurs While Uploading a Patch Set](#)

Error Occurs While Staging a File

See the details below.

Issue

While analyzing the patch plan, the patch plan fails with an unexpected error, although the credentials are correct and although you have write permission on the EM stage location.

For example, the following error is seen on job output page:

"Unexpected error occurred while checking the Normal Oracle Home Credentials"

Cause

You might have set up the Software Library using the OMS Agent file system, and you might not have access to the named credentials that were used for setting up the Software Library.

Solution

To resolve this issue, explicitly grant yourself access to the named credentials that were used to set up the Software Library.

Error Occurs While Uploading a Patch Set

See the details below.

Issue

When you upload a patch set using the Upload Patches to Software Library page, you might see an error stating that the operation failed to read the patch attributes.

For example,

```
ERROR:Failed to read patch attributes from the uploaded patch file <filename>.zip
```

Cause

Although you upload a metadata file along with a patch set ZIP file, sometimes the patch attributes are not read from the metadata file. As a result, the operation fails stating that it could not read the attributes.

Solution

To resolve this issue, manually enter all the patch attributes in the Patch Information section of the Upload Patches to Software Library page. For example, patch number, created on, description, platform, and language.

OPatch Update Job Fails When Duplicate Directories Are Found in the Software Library

See the details below.

Issue

When you run an OPatch Update job, sometimes it might fail with the following error:

```
2011-11-28 10:31:19,127 RemoteJobWorker 20236 ERROR em.jobs startDownload.772-
OpatchUpdateLatest: java.lang.NullPointerException: Category, 'Oracle Software Up-
dates', has no child named, 'OPatch' at oracle.sysman.emInternalSDK.core.patch.util.
ComponentUtil.getComponentCategory (ComponentUtil.java:854)
```

Even after applying the Cloud Control patch released in January 2012, you might see the following error:

```
Category, 'Oracle Software Updates' already exists.
```

Cause

The error occurs when two *Patch Components* directories are found in the Software Library. Particularly when you run two patch upload or download jobs, for example, an OPatch patch download job and a regular patch download job, a race condition is created, which in turn creates two directories with the name *Patch Components*. The Soft-

ware Library does not display any error while creating these duplicate directories, but when you run the OPatch Update job, the job fails with a `NullPointerException`.

Solution

To resolve this issue, do one of the following:

If you see two *Patch Components* directories in the Software Library, then delete the one that has fewer entries, and retry the failed patch upload or download job. To access the Software Library, from the **Enterprise** menu, select **Provisioning and Patching**, and click **Software Library**.

If you see only one *Patch Components* directory, but yet see the error that states that the Oracle Software Updates already exists, then retry the failed patch upload or download.

My Oracle Support Connectivity Issues

This section describes the following issues:

- [Error Occurs While Testing the Proxy Server That Supports Only Digest Authentication](#)


Error Occurs While Testing the Proxy Server That Supports Only Digest Authentication

See the details below.

Issue

On the Proxy Settings page, in the My Oracle Support and Proxy Connection tab, when you provide the manual proxy configuration details, you might see an exception error as shown in the following graphic.

Figure F-1 Proxy Settings Error



Cause

You might have provided the configuration details of a proxy server that supports only the *Digest* authentication schema. By default, the proxy server is mapped only to the

Basic authentication schema, and currently there is no support for Digest authentication schema.

Solution

To resolve this issue, reconfigure your proxy server to make it to use the *Basic* authentication schema.



Note:

For better understanding of connectivity issues related to HTTP Client Logging, you can perform the following steps:

1. Locate the `startup.properties` file under the GC instance directory:
`user_projects/domains/EMGC_DOMAIN/servers/EMGC_OMS1/data/node-manager/startup.properties`
2. Append the following string to the value of the property **Arguments**:
`-DHTTPClient.log.level\=ALL -DHTTPClient.log.verbose\=true`
3. Restart the OMS, and the WebLogic Server Administration Manager by running the following commands:

```
emctl stop oms -all
```

```
emctl start oms
```

4. Navigate to the following location to check the log file:

```
user_projects/domains/EMGC_DOMAIN/servers/EMGC_OMS1/logs/  
EMGC_OMS1-diagnostic.log
```

Alternately, you can also use the **grep** command to find all the HTTP connection logs as follows:

```
grep HTTPClient EMGC_OMS1-diagnostic*.log
```

Host and Oracle Home Credential Issues

This section describes the following security issues:

- [Cannot Create Log Files When You Set Privileged Credentials as Normal Oracle Home Credentials](#)

Cannot Create Log Files When You Set Privileged Credentials as Normal Oracle Home Credentials

See the details below.

Issue

While creating a patch plan, if you choose to override the Oracle home preferred credentials, and set privileged credentials as normal Oracle home credentials inadvertently,

ly as shown in the graphic below, then you will see an error stating that log files cannot be created in the EMStagedPatches directory.

For example,


```
"Unable to create the file <RAC_HOME>/EMStagedPatches/  
PA_APPLY_PATCH_09_02_2011_14_27_13.log"
```

You might also see the following error:

ERROR: SharedDeviceException.

ACTION: Please check whether the configuration is supported or not.

Figure F-2 Inadvertently Selecting Privileged Credentials as Normal Credentials



Cause

When a patch plan is deployed, the patch plan internally uses a deployment procedure to orchestrate the deployment of the patches. While some of the steps in the deployment procedure are run with normal Oracle home credentials, some of the steps are run with privileged Oracle home credentials. However, when you set normal Oracle home credentials as privileged Oracle home credentials, then the deployment procedure runs those steps as a root user instead of the Oracle home owner, and as a result, it encounters an error.

Solution

To resolve this issue, return to the Create Plan Wizard, in the Deployment Options page, in the Credentials tab, set normal credentials as normal Oracle home credentials.

Collection Issues

This section describes the following issues:

- [Missing Details in Plan Wizard](#)
- [Cannot Add Targets to a Patch Plan](#)

Missing Details in Plan Wizard

See the details below.

Issue

When you view a patch plan, sometimes the Create Plan Wizard does not display the expected information. Some pages or sections in the wizard might be blank.

Cause

The issue might be with the details in the Management Repository or with the Create Plan Wizard.

Solution

Identify whether the issue is with the Management Repository or with the Create Plan Wizard. To do so, try retrieving some details from the Management Repository using the commands and URLs mentioned in this section.

- If the URLs return the correct information but the console does not display it, then there might be some technical issue with the Create Plan Wizard. To resolve this issue, contact Oracle Support.
- If the URLs return incorrect information, then there might be some issue with the Management Repository. To resolve this issue, re-create the patch plan.

To retrieve some details from the Management Repository, do the following:

- Retrieve the GUID of the patch plan. To do so, run the following command:

```
select plan_guid from em_pc_plans where name='<name of the plan>';
```

For example,

```
select plan_guid from em_pc_plans where name='t8';
```

The result of the command looks like the following. Note down the GUID of the plan.

```
PLAN_GUID
-----
96901DF943F9E3A4FF60B75FB0FAD62A
```
- Retrieve general information about a patch plan such as its name, type, status, and plan privileges. To do so, use the following URL. This type of information is useful for debugging the *Plan Information* step and the *Review and Deploy* step.

```
https://<hostname>:<port>/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=planInfo
```

Note:

Before retrieving any information about a patch plan using the preceding URL, log in to the Cloud Control console, and from the **Enterprise** menu, select **Provisioning and Patching**, and then click **Patches & Updates**.

- Retrieve information about the patches and the associated targets that are part of the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Patches* step and the *Review* step.

```
https://<hostname>:<port>/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=patches
```

- Retrieve information about the deployment options selected in the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Deployment Options* step and the *Credentials* step.

```
https://<hostname>:<port>/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=deploymentOptions
```

- Retrieve information about the preferred credentials set in the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Credentials* step.

```
https://<hostname>:port/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=preferredCredentials
```

- Retrieve information about the target credentials set in the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Credentials* step.

```
https://<hostname>:port/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=targetCredentials
```

- Retrieve information about the conflict-free patches in the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Validation* step and the *Review & Deploy* step.

```
https://<hostname>:port/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=conflictFree
```

- Retrieve information about the suppressed patches in the patch plan. To do so, use the following URL. This type of information is useful for debugging the *Patches* step.

```
https://<hostname>:port/em/console/CSP/main/patch/plan?planId=<plan_guid>&client=emmos&cmd=get&subset=removedPatchList
```

Cannot Add Targets to a Patch Plan

See the details below.

Issue

While creating a new patch plan or editing an existing patch plan, when you add a new target, you might see the following error:

```
Wrong Platform. Expected: Oracle Solaris on SPARC (64-bit), found: null
```

Cause

The Management Repository might not have the platform information for that target. By default, for every target, the inventory details are regularly collected from the `oraclehomeproperties.xml` file that resides in the Oracle home of the target. Sometimes, the inventory collection might not have completed or might have failed, resulting in

missing data in the Management Repository. Due to these reasons, you might not be able to add those targets.

Solution

To resolve this issue, forcefully recollect the inventory details from the Oracle home of the target.

To retrieve the Oracle home details, follow these steps:

1. From the **Targets** menu, select **All Targets**.
2. On the All Targets page, from the left hand **Refine Search** pane, click **Target Type** menu to expand it.
3. From the Target Type, click **Others**, select **Oracle Home**.
4. All the targets of type Oracle Home are listed. You may search for the host name to drill down to the Oracle home details you are looking for.

To retrieve the inventory details from the Oracle Home on the target host, run the following command from the \$EMDR00T/bin directory:

```
$ emctl control agent runCollection <Oracle_Home_Target>:oracle_home oracle_home_config
```

Here, <Oracle_Home_Target> refers to the name of the Oracle home of the target whose platform information is missing.

For example,

```
$ emctl control agent runCollection db2_2_adc2170603:oracle_home oracle_home_config
```

Patch Recommendation Issues

This section describes the following issues:

- [Patch Recommendations Do Not Appear After Installing Oracle Management Agent on Oracle Exadata Targets](#)

Patch Recommendations Do Not Appear After Installing Oracle Management Agent on Oracle Exadata Targets

See the details below.

Issue

After installing the Management Agent on Oracle Exadata targets, the patch recommendations do not appear.

Cause

The patch recommendations do not appear because the Exadata plug-ins are not deployed.

Solution

To resolve this issue, explicitly deploy the Exadata plug-ins on Exadata targets. To do so, follow these steps:

1. From the **Enterprise** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, in the table, select the Oracle Exadata plug-in version you want to deploy.
3. Click **Deploy On** and select **Management Agent**.
4. In the Deploy Plug-in on Management Agent dialog, in the Selected Management Agent section, click **Add** and select one or more Management Agents where you want to deploy the plug-in, and click **Continue**. Then click **Next**, then **Deploy**.

Patch Plan Issues

This section describes the following issues:

- [Patch Plan Becomes Nondeployable and Fails](#)
- [Instances Not to Be Migrated Are Also Shown as Impacted Targets for Migration](#)
- [Cluster ASM and Its Instances Do Not Appear as Impacted Targets While Patching a Clusterware Target](#)
- [Recovering from a Partially Prepared Plan](#)
- [Error #1009 Appears in the Create Plan Wizard While Creating or Editing a Patch Plan](#)
- [Analysis Succeeds But the Deploy Button is Disabled](#)
- [Patch Plan Fails When Patch Plan Name Exceeds 64 Bytes](#)
- [Out-of-Place Patching Fails for 11.2.0.3 Exadata Clusterware](#)

Patch Plan Becomes Nondeployable and Fails

See the details below.

Issue

The patch plan fails stating it is a nondeployable plan.

Cause

You can add a patch to a target in a patch plan only if the patch has the same release and platform as the target to which it is being added. You will receive a warning if the product for the patch being added is different from the product associated with the target to which the patch is being added. The warning does not prevent you from adding the patch to the patch plan, though. However, when you try to deploy, the plan might fail.

Solution

To make a nondeployable patch plan deployable, divide the patch plan into smaller deployable plans that contain only homogenous patches and targets.


Instances Not to Be Migrated Are Also Shown as Impacted Targets for Migration

See the details below.

Issue

When you deploy a patch plan in out-of-place patching mode, sometimes even the instances that are not selected for migration are identified as impacted targets as shown in the following graphic.

Figure F-3 Instances Not to Be Migrated Are Shown as Impacted Targets



Cause

By default, the patch plan calculates the impacted targets based on only one mode, which in-place patching mode. Therefore, although you have selected out-of-patching mode, the patch plan ignores it and considers only the in-place patching mode as the option selected, and displays all the targets are impacted targets for migration.

Solution

To resolve this issue, ignore the targets you have not selected for migration. They will not be shut down or migration in any case.

Cluster ASM and Its Instances Do Not Appear as Impacted Targets While Patching a Clusterware Target

See the details below.

Issue

While creating a patch plan for patching a clusterware target, on the Deployment Options page, the What to Patch section does not display the cluster ASM and its instances are affected targets. They do not appear in the Impacted Targets section, either. And after deploying the patch plan in out-of-place mode, the cluster ASM and its instances show metric collection error.

Cause

This issue might occur if the clusterware target name in Cloud Control and the cluster-name target name in the mgmt\$target_properties table are not matching.

Solution

To resolve this issue, run the following query to verify the target property ClusterName of the clusterware target:

```
select property_value from mgmt$target_properties where target_name=<CRS Target Name> and property_name="ClusterName"
```

If the returned value is different from the clusterware target name in Cloud Control, then delete the clusterware target and other associated targets, and rediscover them. This time while rediscovering them, ensure that the clusterware target name matches with the name returned by the preceeding query.

Recovering from a Partially Prepared Plan

See the details below.

Issue

When you create a patch plan to patch multiple Oracle homes in out-of-place patching mode, and when you click **Prepare** in the Create Plan Wizard to prepare the patch plan before actually deploying it, sometimes the preparation operation fails with the message *Preparation Failed*.

Cause

The patch plan might have successfully cloned and patched some of the selected Oracle homes, but might have failed on a few Oracle homes. The overall status of the patch plan is based on the patching operation being successful on all the Oracle homes. Even if the patching operation succeeds on most of the Oracle homes and fails only on a few Oracle homes, the overall status is shown as if the patch plan has failed in one of the steps.

Solution

To resolve this issue, fix the errors on failed Oracle homes. Then, go to the procedure instance page and retry the failed steps.

Error #1009 Appears in the Create Plan Wizard While Creating or Editing a Patch Plan

See the details below.

Issue

While creating new patch plan or editing an existing patch plan, you might see the following error in the Create Plan Wizard:

Error #1009

Cause

This error occurs while accessing the Management Repository to extract any details about the patch plan, the targets, or the operation being committed. Usually, SQLException, NullPointerException, or Unhandled exceptions cause these errors.

Solution

To resolve this issue, review the following file, make a note of the exact error or exception logged, and communicate it to Oracle Support.

`$MIDDLEWARE_HOME/gc_inst/em/EMGC_OMS1/sysman/log/emoms.log`

Analysis Succeeds But the Deploy Button is Disabled

See the details below.

Issue

After you successfully analyze the patch plan, when you navigate to the Review of the Create Plan Wizard, you might see the **Deploy** button disabled. Also, the table on the Review page appears empty (does not list any patches.) As a result, you might not be able to deploy the patch plan.

Cause

This error occurs if the patches in the patch plan have already been applied on the target Oracle home. In such a case, the Validation page confirms that the patches have already been applied and therefore they have been skipped, and on the Review page, Deploy button is disabled.

Solution

The patches have already been applied, so you do not have to apply them again. If required, you can manually roll back the patch from the target Oracle home and try applying the patch again.

Patch Plan Fails When Patch Plan Name Exceeds 64 Bytes

See the details below.

Issue

On non-English locales, patch plans with long plan names fail while analyzing, preparing, or deploying, or while switching back. No error is displayed; instead the patch plan immediately reflects the *Failed* state, and logs an exception in the <INSTANCE_HOME>/sysman/log/emoms.log file.

Cause

The error occurs if the patch plan name is too long, that is, if it exceeds 64 bytes. The provisioning archive framework has a limit of 64 bytes for instance names, and therefore, it can accept only plan names that are lesser than 64 bytes. Typically, the instance name is formed using the patch plan name, the plan operation, and the time stamp (PlanName_PlanOperation_TimeStamp). If the entire instance name exceeds 64 bytes, then you are likely to see this error.

Solution

To resolve this issue, do one of the following:

If the patch plan failed to analyze, prepare, or deploy, then edit the plan name and reduce its length, and retry the patching operation.

If the patch plan was deployed successfully, then the patch plan gets locked, and if switchback fails with this error, then you cannot edit the plan name in the wizard. Instead, run the following SQL update command to update the plan name in the Management Repository directly:

```
update em_pc_plans set name = 'New shorter name' where name = 'Older longer name';  
  
commit;
```

Out-of-Place Patching Fails for 11.2.0.3 Exadata Clusterware

See the details below.

Issue

The out-of-place patching fails to unlock the cloned Oracle home in the *Prepare* phase of the patch plan, thus causing the patch plan to fail on the cloned Oracle home. The step *Run clone.pl on Clone Oracle Home* fails.

Cause

This issue occurs if the new Oracle home is different from the Oracle home mentioned in the files <gi_home>/crs/utl/crsconfig_dirs and crsconfig_fileperms that are present in the Grid Infrastructure home. For 11.2.0.3 Exadata Clusterware, the unlock framework works by operating on these files.

Solution

To resolve this issue, you can do one of the following:

- Create a new patch plan for the Exadata Cluster, select the required patch, select **In-Place** in the How to Patch section, and deploy the patch plan.
- Manually apply the patch on the Clusterware Oracle homes of all the nodes of the cluster. Then, clean up the partially cloned Oracle homes on all the nodes, and re-try the *Prepare* operation from the patch plan.

Patch Plan Analysis Issues

This section covers the following issues:

- [Patch Plan Remains in Analysis State Even After the Deployment Procedure Ends](#)
- [Patch Plan Analysis Fails When the Host's Node Name Property Is Missing](#)
- [Link to Show Detailed Progress on the Analysis Is Not Actionable](#)
- [Raising Service Requests When You Are Unable to Resolve Analysis Failure Issues](#)

Patch Plan Remains in Analysis State Even After the Deployment Procedure Ends

See the details below.

Issue

When you analyze a patch plan, sometimes the patch plan shows that analysis is in progress even after the underlying deployment procedure or the job ended successfully.

Cause

This issue can be caused due to one of the following reasons:

- Delayed or no notification from the job system about the completion of the deployment procedure. Typically, after the deployment procedure ends, the job system notifies the deployment procedure. Sometimes, there might be a delay in such notification or there might be no notification at all from the job system, and that can cause the status of the patch plan to show that is always in the analysis state.
- Delay in changing the status of the patch plan. Typically, after the job system notifies the deployment procedure about its completion, the job system submits a new job to change the status of the patch plan. Sometimes, depending on the load, the new job might remain the execution queue for a long time, and that can cause the status of the patch plan to show that is always in the analysis state.
- Failure of the job that changes the status of the patch plan. Sometimes, after the new job is submitted for changing the status of the patch plan, the new job might fail if there are any Management Repository update issues or system-related issues.
- Time zone issues with the Management Repository. If the Management Repository is configured on an Oracle RAC database, and if each instance of the Oracle RAC is running on a different time zone, then when a query is run to know the current system time, it can return incorrect time details depending on which instance serviced the request. Due to incorrect time details, the job that changes the status

of the patch plan might not run at all. This can cause the status of the patch plan to show that is always in the analysis state.

Solution

For time zone-related issue, then first correct the time zone settings on the Oracle RAC nodes, and then restart them. For all other issues, collect the logs and contact Oracle Support.


Patch Plan Analysis Fails When the Host's Node Name Property Is Missing

See the details below.

Issue

When you validate a patch plan created for patching Oracle Clusterware, the validation fails as shown in the graphic below, stating that the node name is missing for the target `<target_name>` of the type `host`. Also, the solution mentioned on the Validation page is incorrect.

Figure F-4 Analysis Fails Stating Node Name Property Is Missing for a Target



Cause

The error occurs because the Create Plan Wizard does not sync up with the actual query or the job you are running. Also, the property `nodeName` is a dynamic property for `HAS` target, which is not marked as a critical property, and therefore, this property could be missing from the Management Repository sometimes. Ideally, it should state that the node name property is missing for the `HAS` target.

Solution

To resolve this issue, run the following command to reload the dynamic properties for the `HAS` target from each node of the Oracle Clusterware.

```
emctl reload dynamicproperties -upload_timeout 600 <target_name>:has
```

For example,

```
emctl reload dynamicproperties -upload_timeout 600 <myhashtarget1>:has
```



Link to Show Detailed Progress on the Analysis Is Not Actionable

See the details below.

Issue

After you analyze a patch plan, the text *Analysis In Progress* on the Validation page appears smaller than normal, and the *here* link for progress details does not work as shown in the following graphic.

Figure F-5 Link to Show Detailed Progress Appears Broken



Cause

You see this error because of a technical issue in rendering this link.

Solution

To resolve this issue, exit the Create Plan Wizard. On the Patches & Updates page, in the **Plans** region, click on the status **Analysis in Progress** against the patch plan where you faced this issue.

Raising Service Requests When You Are Unable to Resolve Analysis Failure Issues

As described in the preceding subsections, there can be several causes for analysis failures, including My Oracle Support connectivity issues, ARU issues, or issues while accessing the Management Repository to extract any details related to the patch plan or targets or the operation being committed. If you encounter any of these issues, follow the solution proposed in the preceding sections, and if you are still unable to resolve the issue, follow these steps, and raise a service request or a bug with the information you collect from the steps.

- *(Online Mode Only)* Verify if the My Oracle Support Web site being used is currently available.
- *(Online Mode Only)* If the plan analysis is failing prior to target analysis being submitted, then verify if the patch analysis is working as expected by running the following URL. Replace `<em_url>` with the correct EM URL, and `<plan_name>` with the actual patch plan name.

```
<em_url>/em/console/CSP/main/patch/plan?cmd=getAnalysisXML&type=att&planName=<plan_name>
```

Verify if the returned XML includes conflict check request and response XMLs for each Oracle home included in the patch plan.

- Open the following file and check the exact error or exception being logged and communicate it to Oracle Support.
`$<MIDDLEWARE_HOME>/gc_inst/em/EMGC_OMS1/sysman/log/emoms.log`

User Account and Role Issues

This section describes the following:

- [Out-of-Place Patching Errors Out If Patch Designers and Patch Operators Do Not Have the Required Privileges](#)

Out-of-Place Patching Errors Out If Patch Designers and Patch Operators Do Not Have the Required Privileges

See the details below.

Issue

When you try out-of-place patching, the patch plan fails with the following error while refreshing the Oracle home configuration:

```
12:58:38 [ERROR] Command failed with error: Can't deploy oracle.sysman.oh on https://<hostname>:<port>/emd/main/
```

Cause

The error occurs because you might not have the following roles as a *Patch Designer* or a *Patch Operator*:

- **ORACLE_PLUGIN_USER**, to view the plug-in user interface
- **ORACLE_PLUGIN_OMS_ADMIN**, to deploy a plug-in on the OMS
- **ORACLE_PLUGIN_AGENT_ADMIN**, to deploy a plug-in on the Management Agent

These roles are required to submit the *Discover Promote Oracle Home Targets* job. The job deploys the Oracle home plug-in on the Management Agent if it is not already deployed.

Solution

Grant these roles explicitly while creating the user accounts. Alternatively, grant the provisioning roles because **EM_PROVISIONING_OPERATOR** and **EM_PROVISIONING_DESIGNER** already include these roles. After granting the privileges, retry the failed deployment procedure step to complete the out-of-patching preparation.

Troubleshooting Linux Patching Issues

My Staging Server Setup DP fails at "Channels Information Collection" step with the error message "Could not fetch the subscribed channels properly". How do I fix this?

This error is seen if there is any network communication error between up2date and ULN. Check if up2date is configured with correct proxy setting by following `https://`

`linux.oracle.com/uln_faq.html` - 9. You can verify if the issue is resolved or not by using the command, `up2date -nox -show-channels`. If the command lists all the subscribed channels, the issue is resolved.

My "up2date -nox -show-channels" command does not list the subscribed channels properly. How do I fix this?

Go to `/etc/sysconfig/rhn/sources` files, uncomment `up2date default` and comment out all the local RPM Repositories configured.

How can I register to channels of other architectures and releases?

Refer to https://linux.oracle.com/uln_faq.html for this and more such related FAQs.

After visiting some other page, I come back to "Setup Groups" page; I do not see the links to the jobs submitted. How can I get it back?

Click **Show** in the details column.

Package Information Job fails with "ERROR: No Package repository was found" or "Unknown Host" error. How do I fix it?

Package Repository you have selected is not good. Check if metadata files are created by running `yum-arch` and `createrepo` commands. The connectivity of the RPM Repository from OMS might be a cause as well.

Even after the deployment procedure finished its execution successfully, the Compliance report still shows my Group as non-compliant, why?

Compliance Collection is a job that runs once in every 24 hour. You should wait for the next cycle of the job for the Compliance report to update itself. Alternately, you can go to the **Jobs** tab and edit the job to change its schedule.

Package Information Job fails with "ERROR: No Package repository was found" or "Unknown Host" error. How do I fix it?

The package repository you have selected is not good. Check if the metadata files are created by running `yum-arch` and `createrepo` commands. The connectivity of the RPM Repository from OMS might be a cause as well.

I see a UI error message saying "Package list is too long". How do I fix it?

Deselect some of the selected packages. The UI error message tells you from which package to unselect.

Troubleshooting Linux Provisioning Issues

I cannot see my stage, boot server in the UI to configure them with the provisioning application?

Either Management Agents have not been installed on the Stage or Boot Server machine, or it is not uploading data to the OMS. Refer to the *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide* for troubleshooting information and known issues.

Bare metal machine is not coming up since it cannot locate the Boot file.

Verify the dhcp settings (/etc/dhcpd.conf) and tftp settings for the target machine. Check whether the services (dhcpd, xinetd, portmap) are running. Make the necessary setting changes if required or start the required services if they are down.

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.

Agent Installation fails after operating system has been provisioned on the bare metal box?

OR

No host name is assigned to the bare metal box after provisioning the operating system?

This might happen if the get - lease-hostnames entry in the dhcpd.conf file is set to true. Edit the dhcpd.conf file to set get - lease-hostnames entry to false.

Also, ensure that length of the host name is compatible with length of the operating system host name.

Bare metal machine hangs after initial boot up (tftp error/kernel error).

This may happen if the tftp service is not running. Enable the tftp service. Go to the /etc/xinetd.d/tftp file and change the disable flag to 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 step will fix the problem:

- Modify `ethtool -s eth0 duplex half` to `ethtool -s eth0 duplex full` in the kickstart file.

Bare metal machine cannot locate the kickstart file (Redhat screen appears for manually entering the values such as 'language', 'keyboard' etc).

This happens if STAGE_TOP_LEVEL_DIRECTORY is not mountable or not accessible. Make sure the stage top level is network accessible to the target machine. Though very rare but this might also happen because of any problem in resolving the stage server host-

name. Enter the IP address of the stage or the NAS server instead of hostname on which they are located, and try the provisioning operation again.

Bare metal machine does not go ahead with the silent installation (Redhat screen appears for manually entering the network details).

Verify that DNS is configured for the stage server host name, and that DHCP is configured to deliver correct DNS information to the target machine. If the problem persists, specify the IP address of the stage or NAS server instead of hostname, and try the provisioning operation again.

After provisioning, the machine is not registered in Enterprise Manager.

This happens if Enterprise Manager Agent is not placed in the `STAGE_TOP_LEVEL_DIRECTORY` before provisioning operation. Place the Enterprise Manager agent in this directory, and try the operation again. It might also happen if the OMS registration password provided for securing the agents is incorrect. Go to the agent oracle home on the target machine, and run the `emctl secure agent` command supplying the correct OMS registration password.

Check the time zone of the OMS and the provisioned operating system. Modify the time zone of the provisioned operating system to match with the OMS time zone.

With 64-bit OS provisioning, agent is not installed.

During OS provisioning, specify the full path of the agent RPM in the Advanced Operating System Properties page.

Provisioning operations cannot be initiated since either one or all of Stage Server, Boot Server, and RPM Repository have not been configured in the Infrastructure page.

Set up at least one stage server, boot server, and RPM repository to proceed with Linux Provisioning.

Submitting the deployment operations throws an error: "An unexpected error has occurred. Please check the log files for details." Logs have the corresponding message: "ComponentType with internal name BMPTYPE not found"

Set up Software Library from the Software Library console.

The deployment procedure fails with directory permission error.

This error occurs because of insufficient user privileges on the stage server machine. `STAGE_TOP_LEVEL_DIRECTORY` should have write permission for the stage server user. In case of NAS, the NAS directory should be mounted on the staging server. If the error appears while writing to the boot directory, then the boot server user must have the write permission.

Bare metal box fails to boot with "reverse name lookup failed" error.

Verify that the DNS has the entry for the IP address and the host name.

Fetching properties from reference machine throws the error: "Credentials specified does not have root access"

Verify if the credentials specified for the reference machine has sudo access.

Following Package/Package Group are not available in the RPM Repository. Either update the Package List or select the correct RPM Repository in the Deployment page.

Verify that the RPM packages mentioned in the error message are present in the repository, and that they are spelled correctly. If not, either copy the packages to the repository or do not install them.

Frequently Asked Questions on Linux Provisioning

What is PXE (Pre-boot Execution Environment)?

The Pre-boot Execution Environment (PXE, aka Pre-Execution Environment) is an environment to bootstrap computers using a network interface card independently of available data storage devices (like hard disks) or installed operating systems. See [Understanding PXE Booting and Kickstart Technology](#) for more information.

Can my boot server reside on a subnet other than the one on which the bare metal boxes will be added?

Yes. But it is a recommended best practice to have boot server in the same subnet on which the bare metal boxes will be added. If the network is subdivided into multiple virtual networks, and there is a separate DHCP/PXE boot server in each network, the Assignment must specify the boot server on the same network as the designated hardware server.

If one wants to use a boot server in a remote subnet then one of the following should be done:

-- Router should be configured to forward DHCP traffic to a DHCP server on a remote subnet. This traffic is broadcast traffic and routers do not normally forward broadcast traffic unless configured to do so. A network router can be a hardware-based router, such as those manufactured by the Cisco Corporation or software-based such as Microsoft's Routing and Remote Access Services (RRAS). In either case, you need to configure the router to relay DHCP traffic to designated DHCP servers.

-- If routers cannot be used for DHCP/BOOTP relay, set up a DHCP/BOOTP relay agent on one machine in each subnet. The DHCP/BOOTP relay agent relays DHCP and BOOTP message traffic between the DHCP-enabled clients on the local network and a remote DHCP server located on another physical network by using the IP address of the remote DHCP server.

Why is Agent rpm staged on the Stage server?

Agent rpm is used for installing the agent on the target machine after booting over the network using PXE. With operating system provisioning, agent bits are also pushed on the machine from the staging location specified in the Advanced Properties.

Can I use the Agent rpm for installing Agent on Stage and Boot Server?

This is true only if the operating system of the Stage or Boot Server machine is Red-Hat Linux 4.0, 3.1 or 3.0 or Oracle Linux 4.0 or later.

Can the yum repository be accessed by any protocol other than HTTP?

Though the rpm repository can be exposed via file:// or ftp:// as well, the recommended method is to expose it via http://. The latter is faster and more secure.

What is the significance of the Status of a directive? How can one change it?

Look at the following table to know the possible Status values and what they signify.

Table F-1 Status Values

Status	Description
Incomplete	This Status signifies that some step was not completed during the directive creation, for example uploading the actual script for the directive, or a user saved the directive while creating it and still some steps need to be performed to make complete the directive creation.
Ready	This signifies that the directive creation was successful and the directive is now ready to be used along with any component/image.
Active	A user can manually change the status of a Ready directive to Active to signify that it is ready for provisioning. Clicking Activate changes the Status to Active.

What is a Maturity Level of a directive? How can one change it?

See [Table F-2](#) to know the possible Status values and what they signify:

Table F-2 Maturity Levels

Maturity Level	Maturity Level Description
Untested	This signifies that the directive has not been tested and is the default maturity level that is assigned to the directive when it is created.
Beta	A directive can be manually promoted to Beta using the Promote button after testing the directive.
Production	A directive can be manually promoted to Production using the Promote button after a user is satisfied that the directive can be used for actual provisioning on production systems.

Can a same component be used in multiple deployments?

Yes. Components are reusable and a given component can be a part of multiple deployments at the same time.

Do I need to edit scheduled deployments associated with a component, if the component is edited?

Yes.

For creating the Linux OS component does the Reference Machine need to have a management agent running on it?

Yes. Reference Machine has to be one of the **managed targets** of the Enterprise Manager.

What is the significance of the Status of a component? How can one change it?

Status of a component is similar to that of a directive. Refer to [What is the significance of the Status of a directive? How can one change it?](#).

What is a Maturity Level of a component? How can one change it?

Maturity Level of a component is similar to that of a directive. Refer to [What is a Maturity Level of a directive? How can one change it?](#).

Refreshing Configurations

If you encounter issues and are expected to refresh the configurations in the host or the Oracle home, then follow the instructions outlined in the following sections:

- [Refreshing Host Configuration](#)
- [Refreshing Oracle Home Configuration](#)

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 Cloud Control, from the **Enterprise** menu, select **Configuration**, and then, click **Refresh Host Configuration**.
2. On the Refresh Host Configuration page, from the Available Hosts pane, select the hosts that the Deployment Procedure will use, and move them to the Selected Hosts pane.

Refresh Host Configuration...

Refresh Host Configuration: Select Hosts

Refresh Host Configuration is used to update the Oracle Enterprise Manager repository with the current configuration information for one or more hosts. Select the hosts whose configuration information you want to update in the Oracle Enterprise Manager repository.

Host Name

Available Hosts		Selected Hosts
ADC2101857.us.oracle.com		adc2100306.us.oracle.com
adc2100914.us.oracle.com		adc2101712.us.oracle.com
adc2110003		
adc2110529.us.oracle.com		
adc2110688.us.oracle.com		
adc2110865.us.oracle.com		
adc2110866.us.oracle.com		
adc2121058.us.oracle.com		
adc2121183.us.oracle.com		
adc2130480.us.oracle.com		

Related Link

[Host Configuration Collection Problems](#)

3. Click **Refresh Hosts**.

Refreshing Oracle Home Configuration

Although the Oracle Management Agent running on a host automatically refreshes the host configuration information every 24 hours, you can also manually refresh the host configuration information the host.



Note:


After patching the targets, refreshing the Oracle home configuration is handled internally by the deployment procedure. However if the refresh does not happen for some reason, then you can refresh the Oracle Home Configuration manually as described in this section.

To manually refresh the host configuration for one host:

1. In Cloud Control, from the **Targets** menu, select **All Targets**.
2. On the All Targets page, from the Refine Search section, click **Target Type** to expand the menu, and from the menu click **Others**, and then click **Oracle Home**.
On the right hand side of the page gets refreshed, and only the Oracle Home targets appear.
3. Click the Target name to select it.
4. On the <target_name> home page, from the **Oracle Home** menu, select **Configuration**, and then click **Last Collected**.

- On the latest Configuration:<target_name> page, from the **Actions** menu select **Refresh** to refresh the Oracle Home configuration for the host.

The following example describes the steps to refresh the Oracle home configuration for the target 11107CRSHome_1_slc00eii:



Reviewing Log Files

This section lists the log files you must review to resolve issues encountered while running a Deployment Procedure.

This section contains the following:

- [OMS-Related Log Files](#)
- [Management Agent-Related Log Files](#)
- [Advanced Options](#)

OMS-Related Log Files

The following are OMS-related log files.

Generic Enterprise Manager Trace File

<EM_INSTANCE_BASE>/em/<OMS_NAME>/sysman/log/emoms.trc

Generic Enterprise Manager Log File

<EM_INSTANCE_BASE>/em/<OMS_NAME>/sysman/log/emoms.log

Where, <EM_INSTANCE_BASE> is the OMS Instance Base directory. By default, the OMS Instance Base directory is gc_inst, which is present under the parent directory of the Oracle Middleware Home.

Management Agent-Related Log Files

The following are Management Agent-related log files:

EMSTATE/sysman/log/gcagent.log

EMSTATE/sysman/log/gcagent.trc

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

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.category.oracle.sysman.emdrep.jobs = parameter to DEBUG.

On the Management Agent Side

On the Management Agent side:

1. Open the following file:
EMSTATE/sysman/config/emd.properties
2. Make the following settings:
Logger.log4j.rootCategory=DEBUG, Rolling, Errors

Index

A

- accessing
 - compliance standards, [50-36](#)
 - compliance violations, [50-11](#)
- accessing Software Library console, [2-2](#)
- adding
 - host configuration, [43-20](#)
 - hosts targets, [43-13](#)
 - local host group, [43-10](#)
 - local user, [43-9](#)
- adding new target data collections, [49-51](#)
- administering pluggable databases, [14-46](#)
- administering, hosts, [43-1](#)
- administration tasks, [43-4](#)
- agent-side rules
 - compliance standard rules, [50-49](#)
 - definition, [50-2](#)
- AIX Installed Packages parser, [49-64](#)
- analyzing configuration health, [49-81](#)
- AND/OR logical operators
 - in comparison rules, [49-27](#)
- Apache HTTPD parser, [49-64](#)
- Application Data Model, [51-3](#)
- application signature, [51-2](#)
 - Oracle E-Business Suite, [51-2](#)
 - Oracle Fusion Applications, [51-2](#)
 - Oracle PeopleSoft Enterprise, [51-2](#)
- ASM, storage, [42-4](#)
- assigning tasks to group owners in change activity plans, [54-14](#)
- associating compliance standard to a target, [50-94](#)
- associations. See relationships, [49-77](#)
- auditing, how to set up compliance standards for, [50-36](#)
- automatic metadata discovery, [51-4](#)
- automating change activity planning, [54-20](#)
- Autosys parser, [49-64](#)

B

- bare metal provisioning
 - concepts
 - boot server, [38-4](#)

- bare metal provisioning (*continued*)
 - concepts (*continued*)
 - reference host, [38-4](#)
 - RPM repository, [38-4](#)
 - stage server, [38-4](#)
- overview, [38-2](#)
- setting up, [38-5](#)
- supported releases, [38-5](#)
- base parsers
 - columnar, [49-62](#)
 - format-specific, [49-57](#)
 - properties, [49-64](#)
- BEA Tuxedo parser. See Ubb Config parser, [49-58](#)
- Blue Martini DNA parser, [49-57](#)
- blueprints, for configuration extensions, [49-52](#)
- boot server, [38-4](#)
 - overview, [38-4](#)
 - setting up, [38-8](#)
- browsing
 - compliance frameworks, [50-33](#)
 - compliance standard rules, [50-74](#)
 - compliance standards, [50-43](#)

C

- change activity plans
 - assigning tasks to group owner, [54-14](#)
 - automating, [54-20](#)
 - changing owner, [54-14](#)
 - creating, [54-6](#)
 - plan like another plan, [54-12](#)
 - task definitions, [54-8](#)
 - task groups, [54-11](#)
 - deactivating, [54-13](#)
 - deleting, [54-13](#)
 - editing, [54-12](#)
 - examples, [54-20](#)
 - exporting, [54-14](#)
 - managing plans, [54-15](#)
 - operations on, [54-11](#)
 - overview, [54-1](#)
 - plans, [54-3](#)
 - printing, [54-14](#)
 - roles and privileges, [54-1](#)

- change activity plans (*continued*)
 - summary, [54-16](#)
 - task definitions, [54-3](#)
 - task groups, [54-5](#)
 - tasks, [54-5](#), [54-17](#)
 - terminology, [54-2](#)
 - using for patching, [54-22](#)
 - viewing tasks, [54-18](#)
- change management, [52-1](#)
- change plan, [52-16](#), [52-20](#)
 - creating change plan, [52-16](#)
 - submitting change plan, [52-20](#)
- change plans
 - overview, [52-15](#)
- changing owner of change activity plans, [54-14](#)
- clearing violations, [50-97](#)
- clone
 - cloning a running Oracle database replay client, [12-2](#)
 - cloning a running Oracle RAC instance, [9-3](#)
- cloning database, [15-10](#), [15-12](#), [15-14](#)
 - existing backup, [15-14](#)
 - RMAN backup, [15-10](#)
 - staging areas, [15-12](#)
- cloning database methods, [15-9](#)
- cloning pluggable databases, [14-18](#)
- cluster verification utility, [4-20](#)
- Coherence Node Provisioning
 - Deploying Coherence Nodes and Clusters
 - Troubleshooting, [35-16](#)
- Coherence Node Provisioning
 - Deploying Coherence Nodes and Clusters, [35-2](#)
 - Creating a Coherence Component, [35-3](#)
 - Deployment Procedure, [35-4](#)
 - Prerequisites, [35-2](#)
 - Getting Started, [35-1](#)
 - Supported Releases, [35-2](#)
- collected configurations for targets, [49-1](#)
- columnar
 - parser parameters, [49-63](#)
 - parsers, [49-62](#)
- commands
 - hosts, [43-11](#)
- comparing data, [52-23](#)
- comparison
 - rules, [1-4](#)
- comparison results, [49-36](#), [49-37](#)
 - standard target, [49-37](#)
 - synchronizing configuration files, [49-37](#)
- comparison template
 - creating or editing, [49-18](#)
 - deleting, [49-20](#)
 - exporting, [49-21](#)
 - importing, [49-21](#)
- comparison template (*continued*)
 - managing, [49-20](#), [49-45](#)
 - Member Settings tab, [49-18](#)
 - Rules for Matching tab, [49-20](#)
 - Template Settings tab, [49-18](#)
 - viewing, [49-20](#)
- comparison templates, [49-17](#)–[49-19](#)
 - creating templates, [49-18](#)
 - editing templates, [49-18](#)
 - property settings, [49-19](#)
 - rules for matching, [49-19](#)
 - template settings, [49-19](#)
- comparison wizard, [49-28](#)
- comparisons, [49-28](#)
- compliance
 - accessing, [50-4](#)
 - compliance frameworks, [50-27](#)
 - compliance standard rule folders, [50-48](#)
 - compliance standard rules, [50-49](#)
 - compliance standards, [50-36](#)
 - configuring, [50-27](#)
 - dashboard, [50-4](#), [50-8](#)
 - evaluating, [50-6](#)
 - evaluation
 - operations on real-time observations, [50-25](#)
 - results, [50-9](#), [50-10](#)
 - importance, [50-18](#)
 - investigating evaluation errors, [50-17](#)
 - library, [50-4](#)
 - managing, [50-27](#)
 - overview, [50-1](#)
 - privileges needed to use, [50-4](#)
 - real-time
 - monitoring facets, [50-76](#)
 - observations, [50-4](#)
 - reports, [50-17](#)
 - results, [50-4](#)
 - roles needed to use, [50-4](#)
 - score, [50-18](#)
 - statistics, accessing, [50-8](#)
 - summary information, [50-9](#)
 - terms used in, [50-2](#)
 - violations
 - accessing, [50-11](#)
 - examples of viewing, [50-14](#)
 - managing, [50-11](#)
 - manual rules, [50-11](#)
 - of a target, [50-14](#)
 - suppressing, [50-11](#)
 - unsuppressing, [50-11](#)
- Compliance Framework privilege, [50-5](#)
- compliance frameworks
 - about, [50-27](#)
 - accessing, [50-27](#)

- compliance frameworks (*continued*)
 - adding compliance standard to, [50-29](#)
 - advantages of using, [50-27](#)
 - compliance of database targets, [50-35](#)
 - compliance score, [50-20](#)
 - definition, [50-2](#)
 - editing importance, [50-29](#)
 - errors, [50-34](#), [50-35](#)
 - evaluation results, [50-34](#)
 - lifecycle status, [50-30](#)
 - operations on, [50-29](#)
 - browsing, [50-33](#)
 - creating, [50-29](#)
 - creating like, [50-31](#)
 - deleting, [50-32](#)
 - editing, [50-31](#)
 - exporting, [50-32](#)
 - importing, [50-33](#)
 - searching, [50-34](#)
 - provided by Oracle, [50-27](#)
 - reasons for using, [50-27](#)
 - user-defined, [50-27](#)
 - violations details, [50-14](#)
- compliance roles
 - EM_COMPLIANCE_DESIGNER, [50-4](#)
 - EM_COMPLIANCE_OFFICER, [50-4](#)
- compliance score
 - compliance framework, [50-20](#)
 - compliance standard, [50-20](#)
 - compliance standard rule, [50-19](#)
 - definition, [50-3](#)
 - parent node, [50-21](#)
 - real-time monitoring rules, [50-19](#)
- compliance standard rule folders
 - about, [50-48](#)
 - creating, [50-48](#)
 - definition, [50-3](#)
 - managing in compliance standard, [50-48](#)
- compliance standard rules
 - about, [50-49](#)
 - agent-side rules, [50-49](#)
 - definition, [50-2](#)
 - operations on, [50-51](#)
 - browsing, [50-74](#)
 - creating agent-side rules, [50-65](#), [50-88](#)
 - creating like, [50-72](#)
 - creating manual rules, [50-66](#), [50-88](#)
 - creating real-time monitoring rules, [50-58](#)
 - creating repository rules, [50-52](#)
 - creating WEbLogic Server Signature rules, [50-54](#)
 - deleting, [50-73](#)
 - editing, [50-73](#)
 - exporting, [50-74](#)
 - importing, [50-74](#)
- compliance standard rules (*continued*)
 - operations on (*continued*)
 - searching, [50-75](#)
 - repository rules, [50-50](#)
 - types, [50-49](#)
- compliance standards
 - about, [50-36](#)
 - accessing, [50-36](#)
 - adding to another compliance standard, [50-38](#)
 - adding to compliance framework, [50-29](#)
 - advantages of using, [50-36](#)
 - associating with targets, [50-45](#)
 - definition, [50-2](#)
 - errors, [50-44](#)
 - evaluation results, [50-43](#), [50-44](#)
 - investigating violations, [50-12](#)
 - operations on, [50-38](#)
 - browsing, [50-43](#)
 - creating, [50-38](#)
 - creating like, [50-41](#)
 - customizing, [50-41](#)
 - deleting, [50-41](#)
 - editing, [50-41](#)
 - exporting, [50-42](#)
 - importing, [50-42](#)
 - searching, [50-43](#)
 - security metrics, enabling, [50-47](#)
 - setting up for auditing use, [50-36](#)
 - violations details, [50-14](#)
- compliance violations, investigating, [50-12](#)
- configuration browser, [49-7](#), [49-8](#), [49-10](#), [49-11](#)
 - inventory and usage details, [49-11](#)
 - saved configurations, [49-10](#)
 - viewing configuration data, [49-8](#)
- Configuration Browser, viewing configurations, [49-7](#)
- configuration changes, [49-12](#)
- configuration collections, [49-40](#)
- configuration comparison
 - ignore rule, [49-23](#)
 - ignore rule example, [49-26](#)
 - matching rule, [49-22](#)
 - matching rule example, [49-25](#)
 - rule examples, [49-25](#)
 - rule language, [49-23](#)
 - rules, [49-21](#)
 - value constraint rule, [49-22](#)
 - wizard, [49-28](#)
- configuration data collections, [49-50](#)
- configuration data, extending collections, [49-50](#)
- configuration extension
 - blueprints, [49-52](#)
 - creating, [49-41](#)
 - credentials, [49-44](#)

- configuration extension (*continued*)
 - database roles, [49-44](#)
 - deleting, [49-47](#)
 - deploying, [49-48](#)
 - editing, [49-41](#)
 - editing deployment, [49-49](#)
 - enabling facet synchronization, [49-46](#)
 - encoding, [49-43](#)
 - exporting, [49-46](#)
 - Files ((amp)) Commands tab, [49-42](#)
 - importing, [49-46](#)
 - post-parsing rule, [49-73](#), [49-75](#), [49-77](#)
 - privileges, [49-47](#)
 - roles, [49-47](#)
 - rules, [49-45](#)
 - sample non-XML parsed file, [49-73](#)
 - sample parsed SQL query, [49-76](#)
 - sample XML parsed file, [49-72](#)
 - undeploying, [49-49](#)
 - versioning, [49-47](#)
 - viewing collection data, [49-50](#)
 - viewing specification details, [49-46](#)
 - XML parsed example (default), [49-56](#)
 - XML parsed example (generic), [49-57](#)
 - XML parsed example (modified), [49-57](#)
 - XPath, [49-45](#)
- configuration extension, creating, [50-88](#)
- configuration extensions, [49-40–49-44](#), [49-46–49-48](#)
 - commands tab, [49-42](#)
 - creating, [49-41](#)
 - deployment, [49-48](#)
 - enabling facet synchronization, [49-46](#)
 - privileges, [49-47](#)
 - setting up credentials, [49-44](#)
 - setting up rules, [49-44](#)
 - SQL tab, [49-43](#)
 - versioning, [49-47](#)
- Configuration Extensions privilege, [50-5](#)
- configuration history, [49-12](#), [49-13](#), [49-15](#)
 - accessing, [49-13](#)
 - annotating, [49-15](#)
 - creating notification list, [49-15](#)
 - scheduling, [49-15](#)
 - searching history, [49-13](#)
- configuration history, job activity, [49-16](#)
- configuration management, [49-1](#)
- configuration searches, [49-3](#), [49-5](#)
 - managing configuration searches, [49-3](#)
 - setting up, [49-5](#)
- Configuration Topology Viewer, [49-78](#)
- configuration topology, viewing, [49-80](#)
- configurations
 - hardware and software, [49-1](#)
 - history, [1-4](#)
- configurations (*continued*)
 - searching, [1-4](#), [49-3](#)
 - viewing, [49-7](#)
- configure Grid Infrastructure, [7-8](#), [7-21](#)
- configuring audit status
 - real-time monitoring rules, [50-58](#)
- configuring hosts, [43-1](#)
- Connect((colon))Direct parser, [49-57](#)
- consumption summary, storage, [42-4](#)
- container databases (CDBs)
 - prerequisites for creating, [5-4](#), [5-9](#), [5-13](#)
 - procedure, [5-5](#), [5-10](#), [5-14](#)
- controlling appearance of information on a graph, [49-86](#)
- CPU statistics
 - hosts monitoring, [42-1](#)
- Create Compliance Entity privilege, [50-5](#)
- creating
 - agent-side compliance standard rules, [50-65](#), [50-88](#)
 - change activity plans, [54-6](#)
 - like another plan, [54-12](#)
 - task definitions, [54-8](#)
 - task groups, [54-11](#)
 - comparison template, [49-18](#)
 - compliance frameworks, [50-29](#)
 - compliance manual rules, [50-88](#)
 - compliance standard example, [50-93](#)
 - compliance standard rule folders, [50-48](#)
 - compliance standards, [50-38](#)
 - compliance standards, considerations, [50-47](#)
 - configuration extension, [50-88](#)
 - custom target type, [49-41](#)
 - manual compliance standard rules, [50-66](#)
 - new relationships, [49-77](#)
 - real-time monitoring compliance standard rules, [50-58](#)
 - real-time monitoring facet folders, [50-82](#)
 - real-time monitoring facets, [50-80](#)
 - repository compliance standard rules, [50-52](#)
 - WebLogic Server Signature compliance standard rules, [50-54](#)
- creating change plans using external clients, [52-19](#)
- creating database provisioning entities, [4-15](#)
- creating database templates, [4-13](#)
- creating databases, [5-4](#), [5-5](#), [7-23](#)
 - prerequisites, [5-4](#)
 - procedure, [5-5](#)
- creating Disk Layout component, [38-16](#)
- creating installation media, [4-12](#)
- creating like
 - compliance frameworks, [50-31](#)
 - compliance standard rules, [50-72](#)
 - compliance standards, [50-41](#)

- creating like (*continued*)
 - real-time monitoring facets, [50-83](#)
- creating Oracle Database, [5-4](#)
 - prerequisites, [5-4](#)
- creating Oracle RAC One Node Database, [5-13](#)
 - prerequisites, [5-13](#)
 - procedure, [5-14](#)
- creating Oracle Real Application Clusters Database, [5-9](#)
 - prerequisites, [5-9](#)
 - procedure, [5-10](#)
- creating provisioning profiles, [4-8](#)
- creating relationships to a target, [49-85](#)
- creating repository rule based on custom configuration collections, [50-85](#)
 - examples, [50-85](#)
- creating users
 - designers, [2-8](#)
 - operators, [2-9](#)
- credentials
 - host, [40-3](#)
 - setting up monitoring, [40-4](#)
- credentials, for configuration extensions, [49-44](#)
- Cron Access parser, [49-62](#)
- Cron Directory parser, [49-62](#)
- CSV parser, [49-62](#)
- Custom CFG parser, [49-65](#)
- custom configuration
 - custom target type, [49-41](#)
- customization
 - changing error handling modes, [56-14](#)
 - customization types, [56-1](#)
 - directive workflow, [56-19](#)
 - overview, [56-1](#)
 - setting up email notifications, [56-15](#)
- customizing
 - hosts environment, [41-1](#)
- customizing topology views, [49-83](#)
- CVU, [4-20](#)

D

- data comparison, [52-2](#), [52-21](#)
 - overview, [52-21](#)
 - requirements, [52-21](#)
- data discovery, [51-1](#)
- data discovery job, [51-5](#)
- data discovery results, [51-6](#)
- Data Guard Rolling Upgrade, [19-26](#)
- Data Redaction, [51-2](#)
- database credentials, [49-44](#)
- database host readiness, [B-1](#)
 - configuring SSH, [B-2](#)
 - environment settings, [B-2](#)
 - kernel requirements, [B-3](#)

- database host readiness (*continued*)
 - environment settings (*continued*)
 - memory requirements, [B-4](#)
 - network and IP requirements, [B-5](#)
 - node time requirements, [B-4](#)
 - package requirements, [B-4](#)
 - installation directories and Oracle inventory, [B-7](#)
 - PDP setup, [B-2](#)
 - setting user accounts, [B-1](#)
 - shell limits, [B-2](#)
 - storage requirements, [B-6](#)
- database provisioning
 - administrator privileges, [4-6](#)
 - deployment procedures, [4-2](#)
 - getting started, [5-16](#)
 - host requirements, [4-6](#)
 - Oracle Database software, [5-28](#)
 - Oracle Databases with Oracle ASM, [5-23](#)
 - Oracle Grid Infrastructure and Oracle Database software, [6-8](#)
 - Oracle Grid Infrastructure and Oracle Databases with Oracle ASM, [6-2](#)
 - Oracle Grid Infrastructure for Oracle Real Application Clusters Databases, [7-1](#)
 - Oracle RAC database with file system on existing cluster, [7-13](#)
 - Oracle RAC database with file system on new cluster, [7-18](#)
 - Oracle Real Application Clusters One Node databases, [8-1](#)
 - prerequisites for designers, [4-6](#)
 - prerequisites for operators, [4-8](#)
 - provisioning and creating Oracle Databases, [5-17](#)
 - provisioning Oracle RAC, [9-1](#)
 - setup, [4-5](#)
 - supported targets, [4-2](#)
 - usecases, [4-3](#)
- database provisioning overview, [4-1](#)
- Database Query
 - parser, [49-58](#)
 - parser parameters, [49-59](#)
- Database Query Paired Column
 - parser, [49-58](#)
 - parser parameters, [49-59](#)
- database templates, [4-14](#)
 - uploading to software library, [4-14](#)
- databases
 - storage, [42-5](#)
- Db2 parser, [49-58](#)
- deactivating
 - change activity plans, [54-13](#)
- default system run level
 - in hosts, [43-6](#)

- delete Oracle RAC, [11-2](#)
 - prerequisites, [11-3](#)
- delete Oracle RAC nodes, [11-6](#)
- deleting
 - change activity plans, [54-13](#)
 - comparison template, [49-20](#)
 - compliance frameworks, [50-32](#)
 - compliance standard rules, [50-73](#)
 - compliance standards, [50-41](#)
 - configuration extension, [49-47](#)
 - custom topology views, [49-84](#)
 - real-time monitoring facets, [50-82](#)
 - relationships from a target, [49-85](#)
- deleting Oracle RAC, [11-2](#)
 - deleting core components, [11-2](#)
 - deleting entire Oracle RAC, [11-2](#)
- deleting pluggable databases, [14-33](#)
- Dell PowerEdge Linux hosts monitoring, [43-19](#)
- dependency analysis, [49-82](#)
- deployable patch plan, [44-6](#)
- Deploying Coherence Nodes and Clusters
 - Deployment Procedure
 - Adding a Coherence, [35-7](#)
 - Environment Variables, [35-9](#)
 - Sample Scripts, [35-10](#)
- deploying configuration extension, [49-48](#)
- deploying SOA composites, [36-9](#)
- Deploying, Undeploying or Redeploying Java EE Applications
 - Creating a Java EE Application Component, [34-3](#)
 - Deploying a Java EE Application, [34-5](#)
 - Getting Started, [34-1](#)
 - Java EE Applications Deployment Procedure, [34-5](#)
 - Prerequisites, [34-3](#)
 - Redeploying a Java EE Application, [34-9](#)
 - Undeploying a Java EE Application, [34-12](#)
- Deploying, Undeploying, or Redeploying Java EE Applications, [34-2](#)
- deployment procedures, [8-2](#)
 - editing the permissions, [55-18](#)
 - phases and steps, [55-8](#)
 - target list, [55-6](#)
 - tracking the status, [55-18](#)
 - User, roles and privileges, [55-3](#)
 - variables, [55-7](#)
 - viewing, editing, and deleting, [55-18](#)
- deployments
 - Management Repository, [49-1](#)
- Designer and Operator Roles, [4-1](#)
- determining configuration health compliance score, [49-81](#)
- dhcp server
 - setting up, [38-8](#)

- diagnosing
 - compliance violations, [50-12](#)

- Directory
 - parser, [49-58](#)
 - parser parameters, [49-60](#)

- discovering hosts, [3-1](#)
 - automatically, [3-1](#)
 - manually, [3-1](#)

- disks
 - statistics, hosts monitoring, [42-2](#)
 - storage, [42-5](#)

E

- E-Business Suite
 - parser, [49-58](#)
 - parser parameters, [49-60](#)

- e-mail notifications
 - configuring outgoing mail server, [56-16](#)
 - entering administrator e-mail and password, [56-17](#)
 - overview, [2-11](#)

- editing
 - change activity plans, [54-12](#)
 - comparison template, [49-18](#)
 - compliance frameworks, [50-31](#)
 - compliance standard rules, [50-73](#)
 - compliance standards, [50-41](#)
 - configuration extension deployment, [49-49](#)
 - host configuration, [43-20](#)
 - local host group, [43-10](#)
 - local user, [43-9](#)
 - real-time monitoring facet folders, [50-82](#)
 - real-time monitoring facets, [50-80](#)

- EM_ALL_OPERATOR
 - EM_ALL_VIEWER, [55-5](#)
 - EM_PATCH_OPERATOR, [55-5](#)
- EM_CAP_ADMINISTRATOR role, [54-1](#)
- EM_CAP_USER role, [54-1](#)
- em_catalog.zip file, [44-24](#)
- EM_COMPLIANCE_DESIGNER role, [49-47](#), [50-4](#)

- EM_COMPLIANCE_OFFICER role, [50-4](#)
- EM_LINUX_PATCHING_ADMIN role, [45-7](#)
- EM_PLUGIN_AGENT_ADMIN role, [49-47](#)
- EM_PLUGIN_OMS_ADMIN role, [49-47](#)

- EMCLI, [44-27](#), [A-1](#)
 - adding ATS service test, [A-52](#)
 - advantages, [A-1](#)
 - creating a new generic component by associating a zip file, [A-47](#)
 - creating properties file, [A-12](#)
 - launching a procedure with an existing saved procedure, [A-18](#)
 - limitations, [A-54](#)

EMCLI (continued)

- migrate and remove a software library storage location, [A-51](#)
 - overview, [A-1](#)
 - patching, [A-24](#)
 - creating a new properties file, [A-26](#)
 - using an existing properties file, [A-30](#)
 - patching verbs, [A-5](#)
 - patching WebLogic Server Target, [A-42](#)
 - prerequisites, [A-2](#)
 - provisioning Oracle WebLogic Server
 - using provisioning profile, [A-33](#)
 - provisioning Oracle Database software, [A-32](#)
 - provisioning Oracle WebLogic Server, [A-33](#)
 - scaling up or scaling out, [A-35](#)
 - provisioning pluggable databases, [A-19](#)
 - creating new pluggable databases, [A-19](#)
 - migrating databases as pluggable databases, [A-23](#)
 - using snapshot profiles, [A-21](#)
 - provisioning user defined deployment procedure, [A-41](#)
 - adding steps/phases, [A-41](#)
 - prerequisites, [A-41](#)
 - running the procedure, [A-42](#)
 - provisioning verbs, [A-11](#)
 - software library verbs, [A-9](#)
 - unplugging pluggable databases, [A-24](#)
 - using an existing properties file, [A-16](#)
 - verbs, [A-2](#)
- EMCLI verbs**
- create_pluggable_database, [A-20](#), [A-21](#)
 - migrate_noncdb_to_pdb, [A-23](#)
 - unplug_pluggable_database, [A-24](#)
 - upload_patches, [A-27](#)
- emctl partool utility, [C-1](#)**
- emctl partool options, [C-1](#)
 - exporting deployment procedure
 - creating PAR file, [C-4](#)
 - retrieving GUID, [C-3](#)
 - exporting deployment procedures, [C-3](#)
 - importing PAR files, [C-5](#)
 - using cloud control, [C-6](#)
 - using command line, [C-5](#)
 - overview, [C-1](#)
 - overview of PAR, [C-1](#)
 - software library, [C-3](#)
- enabling facet synchronization**
- configuration extensions, [A-46](#)
- Enterprise Data Governance dashboard, [51-3](#)**
- enterprise manager users**
- designers, [2-7](#)
 - operators, [2-7](#)
 - super administrators, [2-7](#)

- error handling modes
 - continue on error, [56-14](#)
 - inherit, [56-14](#)
 - skip target, [56-14](#)
 - stop on error, [56-14](#)
- errors
 - in compliance standards, [50-44](#)
- evaluating compliance, [50-6](#)
- evaluation errors
 - compliance, [50-17](#)
- evaluation results
 - compliance, [50-9](#), [50-10](#)
 - compliance standards, [50-43](#)
- examples
 - associating targets, [50-85](#)
 - creating compliance standards, [50-85](#)
 - creating custom configuration, [50-85](#)
 - creating custom-based repository rule based on custom configuration collection, [50-85](#)
 - using change activity plans, [54-20](#)
 - viewing compliance results, [50-85](#)
 - viewing compliance violations, [50-14](#)
- excluding relationships from custom topology views, [49-84](#)
- Execute Host Command
 - group, [43-16](#)
 - multiple hosts, [43-14](#)
 - single host, [43-17](#)
- execution history
 - host commands, [43-18](#)
- execution results
 - host command, [43-18](#)
- exporting
 - change activity plans, [54-14](#)
 - comparison template, [49-21](#)
 - compliance frameworks, [50-32](#)
 - compliance standard rules, [50-74](#)
 - compliance standards, [50-42](#)
 - configuration extension, [49-46](#)
 - real-time monitoring facets, [50-84](#)
- extend Oracle RAC, [10-1](#)
 - prerequisites, [10-2](#)
- extending Oracle RAC, [10-2](#)
 - prerequisites, [10-2](#)
 - procedure, [10-2](#)

F

- file synchronization, [49-37](#)
- file systems
 - storage, [42-5](#)
- finding See accessing, [50-11](#)
- format-specific parsers, [49-57](#)
- Full any Compliance Entity privilege, [50-5](#)

FULL_LINUX_PATCHING_SETUP privilege,
45-5, 45-7

G

Galaxy CFG
 parser, 49-58
 parser parameters, 49-60
generic system, and new relationships, 49-77
GNS settings, 7-22
gold image
 provisioning Oracle database replay client
 using gold image, 12-5
 provisioning Oracle RAC using gold image,
 9-10
GPG keys, 45-5, 45-8
GPG signatures, 45-8
group administration, 43-9, 43-10
groups
 in hosts, 41-2

H

hardware configuration, collecting information,
49-1
history
 statistics, 39-2
 storage, 42-7
history job activity, 49-16
history, of configuration changes, 49-12
history, of configurations, 1-4
Host command
 executing using sudo or PowerBroker, 43-12
 running, 43-13
hosts
 adding host targets, 43-13
 administering, 43-1
 administration, target setup, 40-4
 commands, 43-11
 execution history, 43-18
 execution results, 43-18
 configuration
 adding and editing, 43-20
 configuring, 43-1
 customizing environment, 41-1
 default system run level, 43-6
 Dell PowerEdge Linux monitoring, 43-19
 diagnosing problems on, 39-2
 group administration, 43-9
 groups, 41-2
 installing YAST on, 40-2
 log file alerts, 42-2
 lookup table
 host administration, 43-7
 metric collection errors, 42-2

hosts (*continued*)
 monitoring, 42-1
 monitoring setting up, 40-3
 NFS clients, 43-8
 overview, 39-1
 preferred credentials, 40-3
 running Host command, 43-13
 services, 43-5
 setting up credentials, 40-3
 statistics, 39-1
 CPU, 42-1
 disk, 42-2
 memory, 42-2
 program resource utilization, 42-2
 storage, 42-3
 tools, 43-11
 PowerBroker, 43-11
 Remote File Editor, 43-12
 sudo command, 43-11
 user administration, 43-9
 viewing targets on, 39-2
Hosts Access parser, 49-62

I

ignore rule example, in comparisons, 49-26
ignore rule, in comparisons, 49-23
impact analysis, 49-83
importance
 definition, 50-3
 editing in compliance standard, 50-29
importing
 comparison template, 49-21
 compliance frameworks, 50-33
 compliance standard rules, 50-74
 compliance standards, 50-42
 configuration extension, 49-46
 real-time monitoring facets, 50-84
incidents
 viewing details, 50-25
including relationships in custom topology views,
49-84
information map, 1-5
infrastructure requirements, 2-1
Introscope parser, 49-58
inventory and usage details, 49-11

J

Java Policy parser, 49-65
Java Properties parser, 49-65
job activity
 history, 49-16
Job System privilege, 50-5

K

Kernel Modules parser, [49-62](#)
 Ksplice for Oracle Linux, [42-9](#), [46-8](#)
 Ksplice Metrics, [46-8](#)

L

layers, storage, [42-8](#)
 LDAP parser, [49-65](#)
 library
 compliance, [50-4](#)
 lifecycle management
 overview, [1-1](#)
 solution areas, [1-1](#)
 change management, [1-3](#)
 compliance management, [1-4](#)
 configuration management, [1-4](#)
 discovery, [1-3](#)
 patching, [1-3](#)
 provisioning, [1-3](#)
 solution descriptions, [1-2](#)
 lifecycle status
 compliance frameworks, [50-30](#)
 Linux Directory List parser, [49-62](#)
 Linux hosts, installing YAST, [40-2](#)
 linux patching
 package compliance, [45-8](#)
 prerequisites, [45-3](#)
 registering with ULN, [45-6](#)
 setting up group, [45-7](#)
 setting up infrastructure, [45-2](#)
 setting up linux patching groups for compliance reporting, [45-7](#)
 setting up RPM repository, [45-3](#), [45-5](#)
 linux patching groups, [45-1](#)
 jobs, [45-8](#)
 local file systems, storage, [42-5](#)
 lock down, [55-23](#)
 locking down feature, [4-1](#)
 log file alerts, hosts monitoring, [42-2](#)
 logical operators
 AND/OR, [49-27](#)

M

Manage and Target Metric privilege, [50-5](#)
 Manage any Target Compliance privilege, [50-5](#)
 Management Repository, [49-1](#)
 managing
 change activity plans, [54-15](#)
 compliance standard rule folders, [50-48](#)
 mandatory infrastructure requirements
 creating user accounts, [2-6](#)
 setting up credentials, [2-4](#)

mandatory infrastructure requirements
 setting up software library, [2-2](#)
 manual rules
 definition, [50-3](#)
 violations, [50-11](#)
 matching rule, [49-25](#)
 examples, [49-25](#)
 matching rule example, for comparisons, [49-25](#)
 matching rule, in comparisons, [49-22](#)
 memory statistics
 hosts monitoring, [42-2](#)
 metadata discovery, [51-1](#)
 metadata discovery job, [51-3](#)
 metadata discovery results, [51-5](#)
 metadata XML files
 downloading files, [44-24](#)
 uploading files, [44-24](#)
 metric collection errors
 hosts monitoring, [42-2](#)
 middleware
 enabling as a service (MWaaS), [25-2](#), [26-2](#), [27-1](#), [28-1](#)
 Middleware as a Service (MWaaS)
 enabling, [25-2](#), [26-2](#), [27-1](#), [28-1](#)
 middleware provisioning
 coherence nodes and clusters, [24-22](#)
 deploying/redeploying/undeploying Java EE Applications, [24-21](#)
 deployment procedures, [24-1](#)
 introduction, [24-1](#)
 key concepts, [24-4](#)
 overview, [24-1](#)
 profiles, [24-1](#)
 scaling SOA, Service Bus, and WebLogic Servers, [24-15](#)
 service bus resources, [24-23](#)
 SOA artifacts, [24-22](#)
 WebLogic Domain and Oracle Home Provisioning, [24-7](#)
 Middleware Provisioning
 Middleware Provisioning and Scale Up / Scale Out Best Practices, [31-7](#)
 middleware provisioning console, [24-2](#)
 middleware provisioning solutions, [24-1](#)
 migrating databases as pluggable databases, [14-27](#)
 Mime Types parser, [49-65](#)
 monitoring
 credentials, setting up, [40-4](#)
 hosts, [42-1](#)
 NFS mounts, [42-6](#)
 moving
 task definitions in change activity plans, [54-8](#)
 MQ-Series
 parser, [49-58](#)

MQ-Series (*continued*)
 parser parameters, [49-61](#)

N

named credentials, host, [40-3](#)
 network cards
 configuring, [43-7](#)
 in hosts, [43-7](#)
 network file systems See NFS, [42-5](#)
 NFS (network file systems)
 clients
 adding and editing, [43-8](#)
 host administration, [43-8](#)
 monitoring mounts, [42-6](#)
 storage, [42-5](#)
 notifications, [2-11](#)

O

observations
 notifying user, [50-27](#)
 Odin parser, [49-58](#)
 OPlan, [44-44](#), [44-48](#)
 optional infrastructure requirements
 host configurations, [F-27](#)
 self update for provisioning, [2-10](#)
 setting up e-mail notifications, [2-11](#)
 Oracle clusterware clone, [4-18](#)
 Oracle Clusterware Clone, [4-19](#)
 Oracle Database Clone, [4-17](#)
 Oracle Database topology, [5-1](#)
 Oracle Flex ASM, [7-4](#)
 Oracle Flex Clusters, [7-4](#)
 Oracle Label Security, [51-2](#)
 Oracle ORA parser, [49-58](#)
 Oracle RAC database topology, [7-3](#)
 Oracle Real Application Clusters Database topology, [7-3](#)
 Oracle Service Bus, [37-1](#)
 OS script, load, [43-17](#)

P

PAM Configuration parser, [49-62](#)
 parsers
 AIX Installed Packages, [49-64](#)
 Apache HTTPD, [49-64](#)
 Autosys, [49-64](#)
 Blue Martini DNA, [49-57](#)
 columnar, [49-62](#)
 Connect((colon))Direct, [49-57](#)
 Cron Access, [49-62](#)
 Cron Directory, [49-62](#)

parsers (*continued*)
 CSV, [49-62](#)
 Custom CFG, [49-65](#)
 Database Query, [49-58](#)
 Database Query Paired Column, [49-58](#)
 Db2, [49-58](#)
 Directory, [49-58](#)
 E-Business Suite, [49-58](#)
 format-specific, [49-57](#)
 Galaxy CFG, [49-58](#)
 Hosts Access, [49-62](#)
 Introscope, [49-58](#)
 Java Policy, [49-65](#)
 Java Properties, [49-65](#)
 Kernel Modules, [49-62](#)
 LDAP, [49-65](#)
 Linux Directory List, [49-62](#)
 Mime Types, [49-65](#)
 MQ-Series, [49-58](#)
 Odin, [49-58](#)
 Oracle ORA, [49-58](#)
 PAM Configuration, [49-62](#)
 Process Local, [49-62](#)
 properties, [49-64](#)
 Radia, [49-65](#)
 Sectioned Properties, [49-65](#)
 Secure TTY, [49-62](#)
 Siebel, [49-58](#)
 SiteMinder Agent, [49-65](#)
 SiteMinder Registry, [49-65](#)
 SiteMinder Report, [49-65](#)
 SmWalker, [49-65](#)
 Solaris Installed Packages, [49-63](#)
 Sun ONE Magnus, [49-65](#)
 Sun ONE Obj, [49-65](#)
 Tuxedo, [49-65](#)
 UbbConfig, [49-58](#)
 Unix Config, [49-65](#)
 Unix Crontab, [49-63](#)
 Unix Directory List, [49-63](#)
 Unix Groups, [49-63](#)
 Unix GShadow, [49-63](#)
 Unix Hosts, [49-63](#)
 Unix INETD, [49-63](#)
 Unix Installed Patches, [49-58](#)
 Unix Login, [49-65](#)
 Unix Passwd, [49-63](#)
 Unix PROFTPD, [49-65](#)
 Unix Protocols, [49-63](#)
 Unix Recursive Directory List, [49-59](#)
 Unix Resolve, [49-65](#)
 Unix Services, [49-63](#)
 Unix Shadow, [49-63](#)
 Unix SSH Config, [49-65](#)
 Unix System, [49-65](#)

parsers (*continued*)

- Unix System Crontab, [49-63](#)
- Unix VSFTPD, [49-65](#)
- Unix XINETD, [49-65](#)
- WebAgent, [49-65](#)
- WebLogic (attribute-keyed), [49-55](#)
- WebSphere (attribute-keyed), [49-55](#)
- WebSphere (generic), [49-56](#)
- Windows Checksum, [49-65](#)
- XML (generic), [49-55](#)
- XML default (attribute-keyed), [49-54](#)

patch management solution

- accessing the screen, [44-3](#)
- conflict checks, [44-46](#)
- create plan wizard, [44-6](#)
- customizing deployment procedures, [44-43](#), [44-79](#)
- diagnosing and resolving patching issues, [44-68](#)
- downloading catalog file, [44-24](#)
- introduction, [44-2](#)
- knowledge articles, [44-33](#)
- overview, [44-1](#)
- patch recommendations, [44-30](#)
- patch templates, [44-50](#), [44-74](#)
- patchability reports, [44-28](#)
- patching modes
 - in-place mode, [44-14](#), [44-41](#)
 - offline mode, [44-13](#), [44-24](#)
 - online mode, [44-13](#), [44-22](#)
 - out-of-place mode, [44-14](#), [44-41](#)
 - parallel mode, [44-17](#), [44-41](#)
 - rolling mode, [44-17](#), [44-41](#)
- patching Oracle Data Guard targets, [44-57](#)
- patching Oracle Exadata, [44-53](#)
- patching Oracle Grid Infrastructure targets, [44-52](#)
- patching Oracle Identity Management targets, [44-62](#)
- patching Oracle Siebel targets, [44-62](#)
- patching workflow, [44-19](#)
- performing switchback, [44-49](#)
- preparing database targets, [44-39](#)
- registering proxy details, [44-22](#)
- resolving patch conflicts, [44-78](#)
- rolling back patches, [44-45](#), [44-73](#), [44-81](#)
- scheduling patch plans, [44-48](#)
- searching for patches, [44-33](#)
- setting up infrastructure, [44-19](#)
- supported targets, [44-9](#)
- uploading catalog file, [44-24](#)
- uploading patches to Software Library, [44-25](#)
- validating patch plans, [44-46](#)

patch plans

- accessing, [44-38](#)

patch plans (*continued*)

- adding patches, [44-40](#)
- adding targets, [44-40](#), [44-76](#)
- analyzing, [44-46](#)
- creating, [44-36](#)
- customizing deployment procedures, [44-43](#), [44-79](#)
- deleting, [44-75](#)
- deleting plans, [44-75](#)
- enabling notifications, [44-45](#)
- enabling or disabling conflict checks, [44-46](#)
- overview, [44-4](#)
- patch conflicts, [44-7](#)
- patch plan types, [44-6](#)
- preparing patch plans, [44-39](#)
- roles and privileges, [44-20](#)
- saving as patch templates, [44-50](#)
- scheduling, [44-48](#)
- specifying credentials, [44-43](#)
- specifying deployment options, [44-41](#)
- specifying plan information, [44-39](#)
- staging patches, [44-42](#)
- supported patch types, [44-4](#)
- switchback, [44-49](#)
- using create plan wizard, [44-6](#)
- validating, [44-46](#)
- validation and conflict resolution, [44-6](#)

patch recommendations, [44-30](#)

patch templates

- deleting, [44-76](#)
- deleting templates, [44-76](#)
- downloading patches, [44-75](#)
- modifying templates, [44-74](#)
- using edit template wizard, [44-8](#)
- viewing templates, [44-74](#)

patching

- analyzing the environment, [44-28](#)
- diagnosing issues, [44-70](#)
- identifying applicable patches
 - searching in software library, [44-34](#)
 - searching on MOS, [44-33](#)
 - using knowledge articles, [44-33](#)
 - using patch recommendations, [44-30](#)
- introduction, [44-1](#)
- linux patching, [45-1](#)
- patch management solution, [44-2](#)
- patchability reports, [44-28](#)
- patching linux hosts, [45-1](#)
- resolving issues, [44-72](#)
- using change activity plans, [54-22](#)

patching linux hosts

- concepts, [45-1](#)
- deployment procedures, [45-2](#)
- meeting prerequisites, [45-3](#)
- overview, [45-1](#)

- patching linux hosts (*continued*)
 - package compliance, [45-8](#)
 - registering with ULN, [45-6](#)
 - setting up infrastructure, [45-2](#)
 - setting up linux patching groups for compliance reporting, [45-7](#)
 - setting up RPM repository, [45-3](#), [45-5](#)
 - supported linux releases, [45-2](#)
- phases
 - parallel, [55-8](#)
 - rolling, [55-8](#)
- pluggable database administration
 - altering pluggable database state, [14-47](#)
 - opening/closing pluggable databases, [14-47](#)
 - switching between pluggable databases, [14-46](#)
- pluggable database jobs
 - create, [14-10](#), [14-18](#), [14-42](#)
 - creates, [14-27](#)
 - delete, [14-42](#), [14-45](#)
 - unplug, [14-38](#), [14-44](#)
- pluggable database management requirements, [14-3](#)
- pluggable database provisioning
 - cloning pluggable databases
 - using Full Clone, [14-18](#)
 - using Snap Clone, [14-19](#)
 - creating new pluggable databases, [14-4](#)
 - migrating databases as pluggable databases, [14-27](#)
 - plugging in pluggable databases, [14-10](#)
- pluggable database removal
 - deleting pluggable databases, [14-38](#)
 - unplugging pluggable databases, [14-33](#)
- pluggable databases
 - administering, [14-46](#)
 - getting started, [14-1](#)
 - jobs, [14-42](#)
 - overview, [14-2](#)
 - provisioning, [14-3](#)
 - removing, [14-33](#)
- post-parsing rules, [49-72](#)
- PowerBroker tool, [43-11](#)
 - executing host command, [43-12](#)
- preferred credential, hosts, [40-3](#)
- printing
 - change activity plans, [54-14](#)
- privilege delegation setting, [40-3](#)
- privileges
 - Compliance Framework, [50-5](#)
 - Configuration Extensions, [50-5](#)
 - Create Compliance Entity, [50-5](#)
 - Full any Compliance Entity, [50-5](#)
 - Job System, [50-5](#)
 - Manage any Target Compliance, [50-5](#)
- privileges (*continued*)
 - Manage any Target Metric, [50-5](#)
 - used in change activity plans, [54-1](#)
 - View any Compliance Framework, [50-5](#)
 - View any Target, [50-5](#)
- privileges, for configuration extensions, [49-47](#)
- problems
 - diagnosing on hosts, [39-2](#)
- Process Local parser, [49-62](#)
- program resource utilization statistics, hosts
 - monitoring, [42-2](#)
- properties parser constructs
 - delimited section, [49-71](#)
 - delimited structure, [49-70](#)
 - element cell, [49-71](#)
 - explicit property, [49-69](#)
 - implicit property, [49-69](#)
 - INI section, [49-71](#)
 - keyword name property, [49-68](#)
 - keyword property, [49-68](#)
 - reserved directive, [49-69](#)
 - reserved function, [49-69](#)
 - simple property, [49-68](#)
 - structure, [49-70](#)
 - XML structure, [49-70](#)
- properties parsers, [49-64](#)
 - advanced constructs, [49-68](#)
 - advanced parameters, [49-66](#)
 - basic parameters, [49-65](#)
- Protection Policy
 - Data Redaction, [51-2](#)
 - Oracle Label Security, [51-2](#)
 - Transparent Data Encryption, [51-2](#)
 - Virtual Private Database, [51-2](#)
- provision database
 - Grid Infrastructure and Oracle RAC Database, [7-5](#)
 - Oracle Grid Infrastructure and Oracle Real Application Clusters, [7-1](#)
- provision Linux, [38-1](#)
 - getting started, [38-1](#)
- provision Oracle RAC database
 - with file system on a new cluster, [7-18](#)
 - with file system on an existing cluster, [7-13](#)
- provision Oracle RAC databases, [7-18](#)
- provisioning
 - deleting or scaling down Oracle RAC (Real Application Cluster), [11-1](#)
 - extending Oracle RAC (Real Application Cluster), [10-1](#)
 - provisioning linux operating system, [38-1](#)
 - provisioning Oracle Application Server, [36-1](#)
 - provisioning Oracle database replay client, [12-1](#)

- provisioning (*continued*)
 - provisioning Oracle Service Bus resources, 37-1
 - storage, 42-4
- provisioning bare metal servers, 38-19
- provisioning database client
 - getting started, 12-1
- provisioning Oracle Database client, 12-9
- provisioning Oracle RAC, 9-10, 9-17, 9-25
 - archived software binaries, 9-17
 - gold image, 9-10
 - no root credentials, 9-25
- provisioning Oracle Real Application Clusters
 - One database, 8-2
- provisioning Oracle standby database, 13-1, 13-5, 13-8, 13-9
 - creating logical standby database, 13-5
 - creating physical standby database, 13-1
 - creating primary database backup, 13-9
 - managing existing standby database, 13-8
- provisioning Oracle standby databases, 13-1
- provisioning pluggable databases, 14-3
- provisioning profiles, 4-1
- provisioning SOA artifacts, 36-4
 - gold image, 36-7

R

- Radia parser, 49-65
- real-time monitoring facet folders, editing, 50-82
- real-time monitoring facets, 50-76
 - about, 50-77
 - changing base attributes, 50-84
 - creating, 50-80, 50-82
 - creating like, 50-83
 - definition, 50-3
 - deleting, 50-82
 - editing, 50-80
 - entity types, 50-77
 - exporting, 50-84
 - importing, 50-84
 - operations on, 50-79
 - patterns, 50-78
 - viewing library, 50-79
- real-time monitoring rules
 - compliance score, 50-19
 - definition, 50-2
 - target property filters, 50-60
 - using facets in, 50-58
- real-time monitoring, warnings, 50-46
- real-time observation audit status, definition, 50-3
- real-time observation bundle lifetimes, controlling, 50-58
- real-time observation bundles, definition, 50-3

- real-time observations
 - definition, 50-3, 50-4
 - investigating, 50-22
 - manually setting, 50-25
 - notifying user, 50-26
 - operations on, 50-25
 - rules, types of actions, 50-58
 - viewing, 50-22
- reference host, 38-4
- refresh, storage, 42-8
- relationships, 49-77
- Relocating Pluggable Databases, 14-31, 14-32
- Remote File Editor tool, 43-12
- removing pluggable databases, 14-33
- reports, compliance, 50-17
- repository rules
 - definition, 50-2
 - in compliance standards, 50-50
- results
 - compliance evaluation, 50-9, 50-10
- reverse transform, 49-39
- roles
 - EM_CAP_ADMINISTRATOR, 54-1
 - EM_CAP_USER, 54-1
 - for change activity plans, 54-1
- roles, for configuration extensions, 49-47
- Rolling Upgrade, 19-26
- Rolling Upgrade procedure, 19-28
- rollup options, 49-11
- root cause analysis. *See* dependency analysis, 49-82
- routing configuration, network cards, 43-7
- RPM packages, 45-2, 45-6
- RPM repository, 38-4, 45-2, 45-3
 - overview, 38-4
 - setting up, 38-10, 45-3
- rule examples, for comparisons, 49-25
- rule expressions, in comparisons, 49-23
- rules, 49-21, 49-22
 - include or exclude rule, 49-22
 - matching rule, 49-22
 - value constraint rule, 49-22
- rules expression, 49-23
- rules syntax, 49-23
- rules, in comparisons, 1-4, 49-21
- rules, in configuration extensions, 49-45

S

- save as draft, configuration extension, 49-47
- saved configurations, 49-10
- scale down Oracle RAC, 11-6
- Scaling Up / Scaling Out WebLogic Domains
 - Prerequisites, 31-1

- Scaling Up / Scaling Out WebLogic Domains (*continued*)
 - Running the Scale Up / Scale Out Middle-ware Deployment Procedure, [31-3](#)
- Scaling Up/Scaling Out SOA, Service Bus, and WebLogic Server Domains, [31-1](#)
- schema baseline, [52-2](#)
 - multiple versions, [52-4](#)
- schema baseline version, [52-4](#)
- schema baselines
 - export, [52-5](#)
 - import, [52-5](#)
 - overview, [52-2](#)
- schema change plans, [52-2](#)
- schema comparison, [52-2](#)
- schema comparison versions, [52-8](#)
- schema comparisons, [52-6](#), [52-7](#)
 - comparison options, [52-7](#)
 - overview, [52-6](#)
 - schema map, [52-7](#)
 - scope specification, [52-7](#)
- schema synchronization, [52-2](#)
 - overview, [52-9](#)
 - versions, [52-12](#)
- schema synchronization version, [52-12](#)
 - schema synchronization cycle, [52-12](#)
- schema synchronizations, [52-9–52-11](#)
 - schema map, [52-10](#)
 - scope specification, [52-9](#)
 - synchronization mode, [52-11](#)
 - synchronization options, [52-10](#)
- scope specification, [52-3](#)
- search configurations
 - predefined, [1-4](#), [49-3](#)
 - user-defined, [1-4](#), [49-3](#)
- searching
 - compliance frameworks, [50-34](#)
 - compliance standard rules, [50-75](#)
 - compliance standards, [50-43](#)
- Sectioned Properties parser, [49-65](#)
- Secure TTY parser, [49-62](#)
- security metrics
 - enabling, [50-47](#)
 - in compliance standards, [50-47](#)
- sensitive column type, [51-4](#)
- sensitive data discovery, [51-3](#)
- Service Bus provisioning, [25-1](#), [26-1](#), [27-1](#), [28-1](#)
- services
 - in hosts, [43-5](#)
- setting dependencies in task definitions, [54-8](#)
- setting up
 - environment to monitor hosts, [40-1](#)
 - host credentials, [40-3](#)
 - host monitoring, [40-3](#)
 - monitoring credentials, [40-4](#)
 - target for host administration, [40-4](#)
- Setting Up MOS, [2-10](#)
- Siebel
 - parser, [49-58](#)
 - parser parameters, [49-61](#)
- SiteMinder Agent parser, [49-65](#)
- SiteMinder Registry parser, [49-65](#)
- SiteMinder Report parser, [49-65](#)
- SmWalker parser, [49-65](#)
- SOA provisioning, [25-1](#), [26-1](#), [27-1](#), [28-1](#)
- software library
 - uploading patches, [44-25](#)
- Software Library Administration, [2-2](#), [2-3](#)
- Software Library console, [2-2](#)
- Solaris Installed Packages parser, [49-63](#)
- specify OS users, [7-19](#)
- specifying rules, [49-21](#)
- stage server, [38-4](#)
 - overview, [38-4](#)
- statistics
 - hosts, [39-1](#)
 - storage, [39-2](#)
- statistics, accessing compliance, [50-8](#)
- steps
 - action, [55-8](#)
 - computational, [55-9](#)
 - file transfer, [55-8](#)
 - host command, [55-10](#)
 - job, [55-8](#)
 - library component, [55-8](#)
 - library directive, [55-9](#)
 - manual, [55-8](#)
- storage
 - file systems, [42-5](#)
 - history, [42-7](#)
 - hosts monitoring, [42-3](#)
 - layers, [42-8](#)
 - network file systems, [42-5](#)
 - refresh, [42-8](#)
 - statistics, [39-2](#)
 - utilization, [42-3](#)
 - vendor distribution, [42-7](#)
 - volumes, [42-6](#)
- sudo command, [43-11](#)
 - executing host command, [43-12](#)
- summary
 - change activity plans, [54-16](#)
- Sun ONE Magnus parser, [49-65](#)
- Sun ONE Obj parser, [49-65](#)
- suppressing violations, [50-11](#), [50-95](#)
- synchronizing files, [49-37](#)
- system component structure, [49-81](#)

T

target setup
 host administration, [40-4](#)
 target type, custom, [49-41](#)
 targets
 associating compliance standard with, [50-45](#), [50-94](#)
 host, [39-2](#)
 task groups in change activity plans, [54-5](#)
 tasks
 in change activity plans, [54-17](#)
 TDE
 prerequisites, [23-2](#)
 tools
 hosts, [43-11](#)
 topology
 Oracle RAC database topology, [7-3](#)
 topology viewer
 controlling appearance of information on a graph, [49-86](#)
 creating
 relationships to a target, [49-85](#)
 customizing views, [49-83](#)
 deleting custom views, [49-84](#)
 deleting relationships from a target, [49-85](#)
 dependency analysis, [49-82](#)
 excluding relationships from custom views, [49-84](#)
 impact analysis, [49-83](#)
 including relationships in custom views, [49-84](#)
 tracking configuration changes, [49-12](#)
 Transparent Data Encryption, [51-2](#)
 meeting prerequisites, [23-2](#)
 troubleshooting missing property errors, [44-69](#)
 troubleshooting unsupported configuration errors, [44-70](#)
 Tuxedo parser, [49-65](#)

U

UbbConfig parser, [49-58](#)
 ULN channels, [45-1](#), [45-6](#)
 ULN configuration channel, [45-2](#)
 ULN custom channel, [45-2](#)
 Unbreakable Linux Network (ULN), [45-2](#)
 undeploying configuration extension, [49-49](#)
 Unix Config parser, [49-65](#)
 Unix Crontab parser, [49-63](#)
 Unix Directory List parser, [49-63](#)
 Unix Groups parser, [49-63](#)
 Unix GShadow parser, [49-63](#)
 Unix Hosts parser, [49-63](#)
 Unix INETD parser, [49-63](#)

Unix Installed Patches
 parser, [49-58](#)
 parser parameters, [49-61](#)
 Unix Login parser, [49-65](#)
 Unix Passwd parser, [49-63](#)
 Unix PROFTPD parser, [49-65](#)
 Unix Protocols parser, [49-63](#)
 Unix Recursive Directory List
 parser, [49-59](#)
 parser parameters, [49-61](#)
 Unix Resolve parser, [49-65](#)
 Unix Services parser, [49-63](#)
 Unix Shadow parser, [49-63](#)
 Unix SSH Config parser, [49-65](#)
 Unix System Crontab parser, [49-63](#)
 Unix System parser, [49-65](#)
 Unix VSFTPD parser, [49-65](#)
 Unix XINETD parser, [49-65](#)
 unplugging pluggable databases, [14-33](#)
 unsuppressing violations, [50-11](#)
 up2date patching tool, [45-3](#), [45-4](#), [45-7](#)
 upgrading a database instance, [19-21](#)
 upgrading database
 getting started, [19-1](#)
 upgrading databases, [19-1](#), [19-4](#), [19-6](#), [19-8](#), [19-13](#), [19-17](#), [19-21](#)
 database upgrade wizard, [19-21](#)
 deployment procedure, [19-4](#)
 prerequisites, [19-6](#)
 upgrading Oracle cluster database, [19-8](#)
 upgrading Oracle clusterware, [19-13](#)
 upgrading Oracle database instance, [19-17](#)
 user accounts
 overview, [2-6](#)
 user administration, [43-9](#), [43-10](#)
 User Defined Deployment Procedure, [55-21](#)
 UTF-8, encoding in configuration extensions, [49-43](#)

V

value constraint rule, in comparisons, [49-22](#)
 vendor distribution, storage, [42-7](#)
 Verifying Rolling Upgrade, [19-31](#)
 View any Compliance Framework privilege, [50-5](#)
 View any Target privilege, [50-5](#)
 viewing
 change activity plan tasks, [54-18](#)
 comparison template, [49-20](#)
 configuration data, [49-8](#)
 configuration extension specification details, [49-46](#)
 configuration health problem details, [49-82](#)
 incident details, [50-25](#)
 real-time monitoring facet library, [50-79](#)

- viewing (*continued*)
 - real-time observations, [50-22](#)
- violation details
 - compliance standard, [50-14](#)
- violations
 - clearing, [50-97](#)
 - compliance
 - managing, [50-11](#)
 - viewing examples, [50-14](#)
 - details
 - compliance framework, [50-14](#)
 - compliance standard, [50-14](#)
 - of a target, [50-14](#)
 - suppressing, [50-11](#), [50-95](#)
 - unsuppressing, [50-11](#)
- Virtual Private Database, [51-2](#)
- volumes, storage, [42-6](#)

W

- web-based enterprise management (WBEM)
 - fetchlet metrics, [43-19](#)

- WebAgent parser, [49-65](#)
- WebLogic parser (attribute-keyed), [49-55](#)
- WebLogic Server provisioning, [25-1](#), [26-1](#), [27-1](#), [28-1](#)
- WebSphere parser (attribute-keyed), [49-55](#)
- WebSphere parser (generic), [49-56](#)
- Windows Checksum parser, [49-65](#)

X

- XML default parser (attribute-keyed), [49-54](#)
- XML parser (generic), [49-55](#)
- XPath
 - conditions and expressions, [49-72](#)
 - configuration extension, [49-45](#)

Y

- YAST, installing on Linux hosts, [40-2](#)
- yum patching tool, [45-3](#), [45-7](#)